

Probability Theory

Yury Holubeu *

November 28, 2025

This note is not intended for distribution.

Probability theory and applications are discussed in detail. Links below show contents of [solved problems](#), and [literature](#).

Contents

Preface	14
I — Probability Theory in a Nutshell —	15
1 The Essence of Methods of Probability Theory	15
1.0.1 Probability Spaces	15
1.0.2 Conditional Probabilities, Independence of Events, Combinatorics	17
1.0.3 Test Sequences	17
1.0.4 Random Variables, Distributions	18
1.0.5 Numerical Characteristics of Random Variables	20
1.0.6 Chebyshev's Inequality, the Law of Large Numbers	20
1.0.7 Generating and Characteristic Functions	21
1.0.8 Limit Theorems	21
2 Other Very Useful or Basic Methods in a Nutshell	23
2.0.1 Other Measure-probability Properties	23
2.0.2 Fundamentals of Mathematical Statistics	24
2.0.3 Methods for Proofs of Probabilistic Inequalities (!?!)	24
2.0.4 Markov Chains	24
2.0.5 Probability Theory Famous Paradoxes (!?!?)	25
2.0.6 Basics of Random Processes in a Nutshell (!?!?)	27
II Fundamentals of Probability Theory	28
3 Fundamentals of Probability Theory	28
3.1 Fundamentals of Set Theory and Measure Theory	28
3.1.1 When and Why Do we Need Set Theory for Probability Tasks?	28
3.1.2 Set Theory for Probability Theory (??)	28
3.1.3 Calculus of Events (!?!?)	35
3.1.4 Characteristic Functions of Sets (!?!?!?)	37
3.2 Other Theorems About Set	37
3.2.1 The Borel-cantelli Lemma	37
3.2.2 Kolmogorov 's Inequality	39
3.2.3 The Formalism of Probability Theory (???)	40
3.2.4 Large Deviations Theory	43
3.3 The Simplest Probabilistic Models	44
3.3.1 Properties of the Classical Probabilistic Model	44
3.3.2 Combinatorics for Samples	48
3.3.3 geometric Probability	51
3.4 Conditional Probability.	52
3.4.1 Conditional Probability and Fundamental Theorems	52
3.4.2 Total Probability Formula	53
3.4.3 Bayes Formula.	54

*<https://yuriholubeu.github.io/>, yuri.holubev@gmail.com

3.4.4	Independence of Events.	55
3.5	Independent Test Sequences	59
3.5.1	Bernoulli Scheme	59
3.5.2	Typical Asymptotics of a Bernoulli Scheme	60
3.5.3	Geometric Distribution.	61
3.5.4	Poisson Distribution.	62
3.5.5	Exponential	65
3.5.6	Moivre-laplace Theorems	65
3.6	Random Variables in a Nutshell	68
3.6.1	Typical Random Variable Problems	68
3.6.2	Random Variable Design	68
3.6.3	Random Variable Distribution Function	71
3.6.4	Joint Distributions	73
3.6.5	Mathematical Expectation.	75
3.6.6	Variance	83
3.6.7	Power Moments	84
3.6.8	Covariation	85
3.6.9	Typical Inequalities in Probability Theory	89
3.6.10	About Other Structures	91
4	Introduction to Probability by Blitzstein, Hwang	91
4.0.1	Preface	91
4.1	Probability and Counting	92
4.1.1	1.1 Why Study Probability?	92
4.1.2	1.2 Sample Spaces and Pebble World	94
4.1.3	1.3 Naive Definition of Probability	97
4.1.4	1.4 How to Count	98
4.1.5	1.5 Story Proofs	109
4.1.6	1.6 Non-naive Definition of Probability	110
4.1.7	1.7 Recap	114
4.1.8	1.8R	116
4.1.9	1.9 Exercises	119
4.2	Conditional Probability	129
4.2.1	2.1 The Importance of Thinking Conditionally	129
4.2.2	2.2 Definition and Intuition	130
4.2.3	2.3 Bayes' Rule and the law of Total Probability	136
4.2.4	2.4 Conditional Probabilities are Probabilities	141
4.2.5	2.5 Independence of Events	144
4.2.6	2.6 Coherency of Bayes' Rule	147
4.2.7	2.7 Conditioning as a Problem-solving Tool	148
4.2.8	2.8 Pitfalls and Paradoxes	153
4.2.9	2.9 Recap	157
4.2.10	2.10 R	158
4.2.11	2.11 Exercises	161
4.3	Random Variables and Their Distributions	178
4.3.1	3.1 Random Variables	178
4.3.2	3.2 Distributions and Probability Mass Functions	181
4.3.3	3.3 Bernoulli and Binomial	185
4.3.4	3.4 Hypergeometric	187
4.3.5	3.5 Discrete Uniform	189
4.3.6	3.6 Cumulative Distribution Functions	191
4.3.7	3.7 Functions of Random Variables	193
4.3.8	3.8 Independence of r.v.s	198
4.3.9	3.9 Connections between Binomial and Hypergeometric	202

4.3.10	3.10 Recap	204
4.3.11	3.11r	205
4.3.12	3.12 Exercises	207
4.4	4 Expectation	214
4.4.1	4.1 Definition of Expectation	214
4.4.2	4.2 Linearity of Expectation	216
4.4.3	4.3 Geometric and Negative Binomial	220
4.4.4	4.4 Indicator r.v.s and the Fundamental Bridge	226
4.4.5	4.5 Law of the Unconscious Statistician (lotus)	231
4.4.6	4.6 Variance	232
4.4.7	4.7 Poisson	235
4.4.8	4.8 Connections between Poisson and Binomial	240
4.4.9	4.9 *Using Probability and Expectation to Prove Existence	243
4.4.10	4.10 Recap	247
4.4.11	4.11r	249
4.4.12	4.12 Exercises	250
4.5	5 Continuous Random Variables	267
4.5.1	5.1 Probability Density Functions	267
4.5.2	5.2 Uniform	272
4.5.3	5.3 Universality of the Uniform	275
4.5.4	5.4 Normal	281
4.5.5	5.5 Exponential	287
4.5.6	5.6 Poisson Processes	291
4.5.7	5.7 Symmetry of i.i.d. Continuous r.v.s	294
4.5.8	5.8 Recap	296
4.5.9	5.9r	298
4.5.10	5.10 Exercises	300
4.6	6 Moments	310
4.6.1	6.1 Summaries of a Distribution	310
4.6.2	6.2 Interpreting Moments	314
4.6.3	6.3 Sample Moments	317
4.6.4	6.4 Moment Generating Functions	319
4.6.5	6.5 Generating Moments with Mgfs	322
4.6.6	6.6 Sums of Independent r.v.s via Mgfs	326
4.6.7	6.7 *probability Generating Functions	326
4.6.8	6.8 Recap	330
4.6.9	6.9r	331
4.6.10	6.10 Exercises	334
4.7	Joint Distributions	337
4.7.1	7.1 Joint, Marginal, and Conditional	338
4.7.2	7.2 2D Lotus	355
4.7.3	7.3 Covariance and Correlation	356
4.7.4	7.4 Multinomial	361
4.7.5	7.5 Multivariate Normal	365
4.7.6	7.6 Recap	369
4.7.7	7.7r	371
4.7.8	7.8 Exercises	372
4.8	Transformations	390
4.8.1	8.1 Change of Variables	391
4.8.2	8.2 Convolutions	397
4.8.3	8.3 Beta	400
4.8.4	8.4 Gamma	406
4.8.5	8.5 Beta-gamma Connections	412

4.8.6	8.6 Order Statistics	414
4.8.7	8.7 Recap	417
4.8.8	8.8r	419
4.8.9	8.9 Exercises	421
4.9	9 Conditional Expectation	427
4.9.1	9.1 Conditional Expectation Given an Event	428
4.9.2	9.2 Conditional Expectation Given an r.v.	435
4.9.3	9.3 Properties of Conditional Expectation	437
4.9.4	9.4 *geometric Interpretation of Conditional Expectation	441
4.9.5	9.5 Conditional Variance	442
4.9.6	9.6 Adam and Eve Examples	445
4.9.7	9.7 Recap	448
4.9.8	9.8 R	450
4.9.9	9.9 Exercises	451
4.10	10 Inequalities and Limit Theorems	463
4.10.1	10.1 Inequalities	464
4.10.2	10.2 Law of Large Numbers	471
4.10.3	10.3 Central Limit Theorem	475
4.10.4	10.4 Chi-square and Student- t	479
4.10.5	10.5 Recap	482
4.10.6	10.6 R	484
4.10.7	10.7 Exercises	487
4.11	Markov Chains	494
4.11.1	11.1 Markov Property and Transition Matrix	495
4.11.2	11.2 Classification of States	499
4.11.3	11.3 Stationary Distribution	502
4.11.4	11.4 Reversibility	508
4.11.5	11.5 Recap	514
4.11.6	11.6 R	515
4.11.7	11.7 Exercises	518
4.12	12 Markov Chain Monte Carlo	528
4.12.1	12.1 Metropolis-hastings	529
4.12.2	12.2 Gibbs Sampling	538
4.12.3	12.3 Recap	543
4.12.4	12.4 R	544
4.12.5	12.5 Exercises	545
4.13	13 Poisson Processes	546
4.13.1	13.1 Poisson Processes in one Dimension	546
4.13.2	13.2 Conditioning, Superposition, and Thinning	548
4.13.3	13.3 Poisson Processes in Multiple Dimensions	557
4.13.4	13.4 Recap	559
4.13.5	13.5 R	559
4.13.6	13.6 Exercises	560
5	A Comprehensive Course of Probability Theory by Klenke	563
5.1	1 Basic Measure Theory	563
5.1.1	1.1 Classes of Sets	563
5.1.2	1.2 Set Functions	570
5.1.3	1.3 The Measure Extension Theorem	575
5.1.4	1.4 Measurable Maps	586
5.1.5	1.5 Random Variables	592
5.2	2 Independence	596
5.2.1	2.1 Independence of Events	596
5.2.2	2.2 Independent Random Variables	601

5.2.3	2.3 Kolmogorov's 0-1 Law	607
5.2.4	2.4 Example: Percolation	609
5.3	3 Generating Functions	618
5.3.1	3.1 Definition and Examples	618
5.3.2	3.2 Poisson Approximation	620
5.3.3	3.3 Branching Processes	622
5.4	4 The Integral	623
5.4.1	4.1 Construction and Simple Properties	624
5.4.2	4.2 Monotone Convergence and Fatou's Lemma	629
5.4.3	4.3 Lebesgue Integral Versus Riemann Integral	631
5.5	5 Moments and Laws of Large Numbers	634
5.5.1	5.1 Moments	634
5.5.2	5.2 Weak Law of Large Numbers	639
5.5.3	5.3 Strong Law of Large Numbers	641
5.5.4	5.4 Speed of Convergence in the Strong Lln	649
5.5.5	5.5 The Poisson Process	652
5.6	6 Convergence Theorems	657
5.6.1	6.1 Almost Sure and Measure Convergence	657
5.6.2	6.2 Uniform Integrability	661
5.6.3	6.3 Exchanging Integral and Differentiation	665
5.7	7 L^p -Spaces and the Radon-Nikodym Theorem	667
5.7.1	7.1 Definitions	667
5.7.2	7.2 Inequalities and the Fischer-riesz Theorem	668
5.7.3	7.3 Hilbert Spaces	672
5.7.4	7.4 Lebesgue's Decomposition Theorem	675
5.7.5	7.5 Supplement: Signed Measures	677
5.7.6	7.6 Supplement: Dual Spaces	683
5.8	8 Conditional Expectations	685
5.8.1	8.1 Elementary Conditional Probabilities	685
5.8.2	8.2 Conditional Expectations	687
5.8.3	8.3 Regular Conditional Distribution	693
5.9	9 Martingales	699
5.9.1	9.1 Processes, Filtrations, Stopping Times	699
5.9.2	9.2 Martingales	703
5.9.3	9.3 Discrete Stochastic Integral	706
5.9.4	9.4 Discrete Martingale Representation Theorem and the Crr Model	707
5.10	10 Optional Sampling Theorems	710
5.10.1	10.1 Doob Decomposition and Square Variation	710
5.10.2	10.2 Optional Sampling and Optional Stopping	713
5.10.3	10.3 Uniform Integrability and Optional Sampling	716
5.11	11 Martingale Convergence Theorems and Their Applications	718
5.11.1	11.1 Doob's Inequality	718
5.11.2	11.2 Martingale Convergence Theorems	719
5.11.3	11.3 Example: Branching Process	727
5.12	12 Backwards Martingales and Exchangeability	728
5.12.1	12.1 Exchangeable Families of Random Variables	728
5.12.2	12.2 Backwards Martingales	731
5.12.3	12.3 De Finetti's Theorem	733
5.13	13 Convergence of Measures	737
5.13.1	13.1 A Topology Primer	737
5.13.2	13.2 Weak and Vague Convergence	741
5.13.3	13.3 Prohorov's Theorem	747
5.13.4	13.4 Application: a Fresh Look at de Finetti's Theorem	755

5.14	14 Probability Measures on Product Spaces	756
5.14.1	14.1 Product Spaces	757
5.14.2	14.2 Finite Products and Transition Kernels	759
5.14.3	14.3 Kolmogorov's Extension Theorem	765
5.14.4	14.4 Markov Semigroups	769
5.15	15 Characteristic Functions and the Central Limit Theorem	772
5.15.1	15.1 Separating Classes of Functions	772
5.15.2	15.2 Characteristic Functions: Examples	777
5.15.3	15.3 Lévy's Continuity Theorem	782
5.15.4	15.4 Characteristic Functions and Moments	786
5.15.5	15.5 The Central Limit Theorem	790
5.15.6	15.6 Multidimensional Central Limit Theorem	797
5.16	16 Infinitely Divisible Distributions	798
5.16.1	16.1 Lévy-khinchin Formula	799
5.16.2	16.2 Stable Distributions	807
5.17	17 Markov Chains	813
5.17.1	17.1 Definitions and Construction	813
5.17.2	17.2 Discrete Markov Chains: Examples	818
5.17.3	17.3 Discrete Markov Processes in Continuous Time	821
5.17.4	17.4 Discrete Markov Chains: Recurrence and Transience	825
5.17.5	17.5 Application: Recurrence and Transience of Random Walks	828
5.17.6	17.6 Invariant Distributions	834
5.17.7	17.7 Stochastic Ordering and Coupling	839
5.18	18 Convergence of Markov Chains	842
5.18.1	18.1 Periodicity of Markov Chains	842
5.18.2	18.2 Coupling and Convergence Theorem	845
5.18.3	18.3 Markov Chain Monte Carlo Method	849
5.18.4	18.4 Speed of Convergence	855
5.19	19 Markov Chains and Electrical Networks	860
5.19.1	19.1 Harmonic Functions	860
5.19.2	19.2 Reversible Markov Chains	863
5.19.3	19.3 Finite Electrical Networks	864
5.19.4	19.4 Recurrence and Transience	868
5.19.5	19.5 Network Reduction	873
5.19.6	19.6 Random Walk in a Random Environment	880
5.20	20 Ergodic Theory	882
5.20.1	20.1 Definitions	883
5.20.2	20.2 Ergodic Theorems	885
5.20.3	20.3 Examples	888
5.20.4	20.4 Application: Recurrence of Random Walks	889
5.20.5	20.5 Mixing	891
5.20.6	20.6 Entropy	893
5.21	21 Brownian Motion	895
5.21.1	21.1 Continuous Versions	896
5.21.2	21.2 Construction and Path Properties	900
5.21.3	21.3 Strong Markov Property	904
5.21.4	21.4 Supplement: Feller Processes	907
5.21.5	21.5 Construction via l^2 -Approximation	909
5.21.6	21.6 The Space $c([0, \infty))$	914
5.21.7	21.7 Convergence of Probability Measures on $c([0, \infty))$	916
5.21.8	21.8 Donsker's Theorem	918
5.21.9	21.9 Pathwise Convergence of Branching Processes*	921
5.21.10	21.10 Square Variation and Local Martingales	926

5.22	22 Law of the Iterated Logarithm	935
5.22.1	22.1 Iterated Logarithm for the Brownian Motion	935
5.22.2	22.2 Skorohod's Embedding Theorem	937
5.22.3	22.3 Hartman-wintner Theorem	941
5.23	23 Large Deviations	943
5.23.1	23.1 Cramér's Theorem	943
5.23.2	23.2 Large Deviations Principle	947
5.23.3	23.3 Sanov's Theorem	950
5.23.4	23.4 Varadhan's Lemma and Free Energy	954
5.24	24 The Poisson Point Process	958
5.24.1	24.1 Random Measures	959
5.24.2	24.2 Properties of the Poisson Point Process	962
5.24.3	24.3 The Poisson-dirichlet Distribution*	968
5.24.4	The Chinese Restaurant Process	971
5.25	25 The Itô Integral	973
5.25.1	25.1 Itô Integral with Respect to Brownian Motion	973
5.25.2	25.2 Itô Integral with Respect to Diffusions	979
5.25.3	25.3 The Itô Formula	982
5.25.4	25.4 Dirichlet Problem and Brownian Motion	989
5.25.5	25.5 Recurrence and Transience of Brownian Motion	990
5.26	26 Stochastic Differential Equations	993
5.26.1	26.1 Strong Solutions	993
5.26.2	26.2 Weak Solutions and the Martingale Problem	1000
5.26.3	26.3 Weak Uniqueness via Duality	1005

III Problems 1011

6	Typical tasks	1011
6.1	Problems About the Simplest Probabilistic Schemes	1011
6.1.1	Problems About Combinatorial Probability Models (?)	1011
6.1.2	Problems About Geometric Probabilities	1027
6.1.3	Set Property Tasks	1032
6.1.4	Bone and Coin Problems (!!!)	1034
6.1.5	Map Tasks (!!!)	1036
6.2	Test Sequence Tasks (!!!)	1037
6.2.1	Problems About Conditional Probabilities	1037
6.2.2	Problems About the Independence of Events	1044
6.2.3	Problems About the Total Probability Formula	1045
6.2.4	Problems About Bernoulli's Scheme (!!!!!?)	1046
6.2.5	Problems About Poisson Events (??)	1050
6.2.6	Problems About Polynomial Schema	1052
6.3	Problems of Probability Theory in Life and Everyday Life	1053
6.3.1	Problems About Typical Simple Life Tasks (!!!!!)	1053
6.3.2	Service Sector Tasks with Many Customers (!!!)	1055
6.3.3	Distance Tasks	1056
6.3.4	Meeting Tasks	1057
6.4	Problems About Random Variables (!!!!)	1057
6.4.1	Problems About the Inequalities of Chebyshev and Markov	1057
6.4.2	Problems of Probability Distribution of Random Variables (!????)	1061
6.4.3	Problems About Mathematical Expectations, Covariance, Correlation	1080
6.4.4	Problems About Conditional Distributions	1089
6.4.5	Problems About Normal Distribution	1095
6.5	Problems About Limit Theorems That Produce and Characteristic Functions	1099

6.5.1	Problems About the law of Large Numbers and the Borel-cantelli Lemma (????)	1099
6.5.2	Limit Theorem Proof Problems	1102
6.5.3	Problems About Characteristic and Producing Functions	1106
6.5.4	Problems About Bonferroni Inequalities and Convergence to the Poisson Distribution (???)	1113
6.5.5	Problems About the Central Limit Theorem and the Method of Characteristic Functions (??!)	1117
6.6	Questions to Test Knowledge	1122
6.6.1	General Questions to Test Understanding	1122
6.6.2	Typical Technical Questions to Test Understanding	1122
6.6.3	Questions to Test Understanding of Details	1122
7	Other Problems	1123
7.1	Problems About Elementary Random Processes	1123
7.1.1	Problems About the Simplest Random Processes	1123
7.1.2	Problems About Poisson Processes	1134
7.1.3	Problems About Markov Chains (!!!)	1135
7.1.4	Problems About Markov Chains Applications	1145
7.2	Problems of Elementary Mathematical Statistics (????)	1147
7.2.1	Problems About the Parameters of Mathematical Statistics	1148
7.3	Problems Arising from the Application of Probability Theory	1152
7.3.1	Problems of Probability Theory in Physics	1152
8	Problems from Applications	1153
8.1	Problems From Applications in Physics	1153
8.1.1	Problems From Applications in Statistical Physics	1153
IV	— Exotic Topics in a Nutshell —	1155
9	Exotic Theoretical Topics in a Nutshell	1155
9.0.1	Quantum Probability Theory (???)	1155
10	Applications in a Nutshell	1155
10.1	Applications in Mathematics	1155
10.2	Applications in Physics	1155
10.2.1	Quantum Mechanics	1155
10.2.2	Thermodynamics	1155
10.2.3	QFT	1155
10.2.4	Physical Kinetics	1155
V	Other Topics	1156
11	Random Processes: Basics	1156
11.1	Basics of the Theory of Random Processes in a Nutshell	1156
11.1.1	Typical Random Processes	1156
11.1.2	Independent Incremental Processes - Poisson Process	1159
11.1.3	Markov Processes	1163
11.1.4	On the use of Markov Processes (?)	1163
11.1.5	Random Process Addendums	1164
11.2	methods for Setting Random Variables	1164
11.2.1	the Probability of Transition in the Presence of Random Processes	1164
11.2.2	Fundamentals of Brownian Motion	1167
11.3	Sequences of Past-dependent Tests	1168

11.3.1	Basics of Markov Chains	1168
11.3.2	Branching Processes	1178
11.4	Typical Applied Probabilistic Models (!!)	1181
11.4.1	Particle Walking Models	1181
11.4.2	Fokker-planck Equation	1181
12	Mathematical Statistics: Basics	1182
12.1	Fundamentals of Mathematical Statistics in a Nutshell	1182
12.1.1	Some Inequalities in Mathematical Statistics	1182
12.1.2	Central Limit Theorem	1182
12.1.3	Characteristic Functions (?!!!!!)	1185
12.1.4	law of Large Numbers	1201
12.1.5	types of Convergence of Sequences of Random Variables	1203
12.1.6	Some Problems About Mathematical Statistics	1208
12.2	Fundamentals of Mathematical Statistics	1208
12.2.1	Confidence Interval	1208
12.2.2	Asymptotic Confidence Intervals	1209
12.2.3	Linear Regression	1209
12.2.4	Statistical Hypothesis Testing	1209
12.2.5	Equally Most Powerful Criteria	1209
12.2.6	Testing Simple Hypotheses	1209
12.2.7	Testing Complex Hypotheses	1209
12.2.8	Verification of Linear Hypotheses in the Gaussian Regression Model	1209
12.2.9	Consent Criteria	1209
12.2.10	Kolmogorov's Goodness-of-fit Criterion.	1209
12.2.11	χ^2 Test.	1209
12.2.12	Other Mathematical Statistics (???)	1209
12.3	Typical Distributions	1210
12.3.1	Normal Distribution	1210
12.3.2	Homogeneous Distribution	1210
12.3.3	Demonstrative	1211
12.3.4	Cauchy Distribution	1211
12.3.5	Binomial (???)	1212
12.3.6	Geometric Distribution	1212
12.3.7	Hypergeometric Distribution	1212
12.3.8	Gamma Distribution	1213
12.3.9	Other Distributions: Overview	1213
12.4	Statistical Methods of Hypothesis Testing	1214
12.4.1	About Hypothesis Testing in Life (???)	1214
12.4.2	hypothesis Testing Theory	1214
13	Applications in Physics (!?!?!?)	1214
14	Analytics and Data Analysis with Probability Theory	1214
15	Other Applications in Programming (???)	1215
15.1	Typical Applications in Programming	1215
15.1.1	Visualization of Probabilistic Model Results	1215
15.2	Machine Learning	1216
15.3	Bayesian Machine Learning	1216

16 Special constructions	1217
16.0.1 Frobenius-perron Theorem	1217
16.1 Random Matrices	1217
16.2 diagramming Techniques	1218
16.3 Quantum Probability	1218
16.4 Special Models and Paradoxes	1218
16.4.1 Special Patterns of Ordinary Life	1218
16.5 Other Theories	1219
16.5.1 Decision Making Using Mathematics	1219
16.6 Applications to Financial Mathematics	1220
16.7 Game Theory	1220
16.7.1 Lotteries: Simple Cases	1220
16.8 Models of Situations in Life (!)	1220
16.8.1 Random Meeting	1220
17 Part I Methods by Alon, Spencer	1220
17.1 1 The Basic Method	1220
17.1.1 1.1 The Probabilistic Method	1220
17.1.2 1.2 Graph Theory	1222
17.1.3 1.3 Combinatorics	1224
17.1.4 1.4 Combinatorial Number Theory	1226
17.1.5 1.5 Disjoint Pairs	1227
17.1.6 1.6 Exercises	1228
17.2 2 Linearity of Expectation	1229
17.2.1 2.1 Basics	1230
17.2.2 2.2 Splitting Graphs	1230
17.2.3 2.3 Two Quickies	1232
17.2.4 2.4 Balancing Vectors	1233
17.2.5 2.5 Unbalancing Lights	1234
17.2.6 2.6 Without Coin Flips	1235
17.2.7 2.7 Exercises	1236
17.3 3 Alterations	1238
17.3.1 3.1 Ramsey Numbers	1238
17.3.2 3.2 Independent Sets	1239
17.3.3 3.3 Combinatorial Geometry	1240
17.3.4 3.4 Packing	1241
17.3.5 3.5 Recoloring	1242
17.3.6 3.6 Continuous Time	1244
17.3.7 3.7 Exercises	1247
17.4 4 The Second Moment	1248
17.4.1 4.1 Basics	1249
17.4.2 4.2 Number Theory	1250
17.4.3 4.3 More Basics	1252
17.4.4 4.4 Random Graphs	1253
17.4.5 4.5 Clique Number	1257
17.4.6 4.6 Distinct Sums	1258
17.4.7 4.7 The Rödl Nibble	1259
17.4.8 4.8 Exercises	1263
17.5 5 The Local Lemma	1266
17.5.1 5.1 The Lemma	1266
17.5.2 5.2 Property b and Multicolored Sets of Real Numbers	1268
17.5.3 5.3 Lower Bounds for Ramsey Numbers	1269
17.5.4 5.4 A Geometric Result	1270
17.5.5 5.5 The Linear Arboricity of Graphs	1271

17.5.6	5.6 Latin Transversals	1274
17.5.7	5.7 The Algorithmic Aspect	1275
17.5.8	5.8 Exercises	1277
17.6	6 Correlation Inequalities	1278
17.6.1	6.1 The Four Functions Theorem of Ahlswede and Daykin	1279
17.6.2	6.2 The Fkg Inequality	1281
17.6.3	6.3 Monotone Properties	1283
17.6.4	6.4 Linear Extensions of Partially Ordered Sets	1284
17.6.5	6.5 Exercises	1286
17.7	7 Martingales and Tight Concentration	1287
17.7.1	7.1 Definitions	1287
17.7.2	7.2 Large Deviations	1288
17.7.3	7.3 Chromatic Number	1290
17.7.4	7.4 Two General Settings	1292
17.7.5	7.5 Four Illustrations	1295
17.7.6	7.6 Talagrand's Inequality	1297
17.7.7	7.7 Applications of Talagrand's Inequality	1300
17.7.8	7.8 Kim-vu Polynomial Concentration	1301
17.7.9	7.9 Exercises	1302
17.8	8 The Poisson Paradigm	1304
17.8.1	8.1 The Janson Inequalities	1304
17.8.2	8.2 The Proofs	1306
17.8.3	8.3 Brun's Sieve	1308
17.8.4	8.4 Large Deviations	1311
17.8.5	8.5 Counting Extensions	1312
17.8.6	8.6 Counting Representations	1313
17.8.7	8.7 Further Inequalities	1316
17.8.8	8.8 Exercises	1317
17.9	9 Pseudorandomness	1318
17.9.1	9.1 The Quadratic Residue Tournaments	1319
17.9.2	9.2 Eigenvalues and Expanders	1321
17.9.3	9.3 Quasirandom Graphs	1326
17.9.4	9.4 Exercises	1331
18	Part II Topics by Alon, Spencer	1332
18.1	10 Random Graphs	1332
18.1.1	10.1 Subgraphs	1333
18.1.2	10.2 Clique Number	1335
18.1.3	10.3 Chromatic Number	1336
18.1.4	10.4 Zero-one Laws	1338
18.1.5	10.5 Exercises	1344
18.2	11 The Erdős-rényi Phase Transition	1346
18.2.1	11.1 AN Overview	1346
18.2.2	11.2 Three Processes	1348
18.2.3	11.3 The Galton-watson Branching Process	1349
18.2.4	11.4 Analysis of the Poisson Branching Process	1350
18.2.5	11.5 The Graph Branching Model	1352
18.2.6	11.6 The Graph and Poisson Processes Compared	1353
18.2.7	11.7 The Parametrization Explained	1354
18.2.8	11.8 The Subcritical Regimes	1355
18.2.9	11.9 The Supercritical Regimes	1355
18.2.10	11.10 The Critical Window	1358
18.2.11	11.11 Analogies to Classical Percolation Theory	1360
18.2.12	11.12 Exercises	1363

18.3	12 Circuit Complexity	1365
18.3.1	12.1 Preliminaries	1365
18.3.2	12.2 Random Restrictions and Bounded-depth Circuits	1367
18.3.3	12.3 More on Bounded-depth Circuits	1370
18.3.4	12.4 Monotone Circuits	1372
18.3.5	12.5 Formulae	1374
18.3.6	12.6 Exercises	1375
18.4	13 Discrepancy	1376
18.4.1	13.1 Basics	1377
18.4.2	13.2 SIX Standard Deviations Suffice	1378
18.4.3	13.3 Linear and Hereditary Discrepancy	1381
18.4.4	13.4 Lower Bounds	1383
18.4.5	13.5 The Beck-fiala Theorem	1385
18.4.6	13.6 Exercises	1386
18.5	14 Geometry	1387
18.5.1	14.1 The Greatest Angle Among Points in Euclidean Spaces	1388
18.5.2	14.2 Empty Triangles Determined by Points in the Plane	1389
18.5.3	14.3 Geometrical Realizations of Sign Matrices	1390
18.5.4	14.4 ϵ -nets and Vc-dimensions of Range Spaces	1392
18.5.5	14.5 Dual Shatter Functions and Discrepancy	1395
18.5.6	14.6 Exercises	1398
18.6	15 Codes, Games and Entropy	1399
18.6.1	15.1 Codes	1399
18.6.2	15.2 Liar Game	1401
18.6.3	15.3 Tenure Game	1403
18.6.4	15.4 Balancing Vector Game	1404
18.6.5	15.5 Nonadaptive Algorithms	1406
18.6.6	15.6 Half Liar Game	1406
18.6.7	15.7 Entropy	1407
18.6.8	15.8 Exercises	1412
18.7	16 Derandomization	1414
18.7.1	16.1 The Method of Conditional Probabilities	1414
18.7.2	16.2 D -wise Independent Random Variables in Small Sample Spaces	1417
18.7.3	16.3 Exercises	1421
18.8	17 Graph Property Testing	1422
18.8.1	17.1 Property Testing	1422
18.8.2	17.2 Testing Colorability	1423
18.8.3	17.3 Szemerédi's Regularity Lemma	1426
18.8.4	17.4 Testing Triangle-freeness	1429
18.8.5	17.5 Characterizing the Testable Graph Properties	1430
18.8.6	17.6 Exercises	1432
18.8.7	Appendix A: Bounding of Large Deviations	1434
18.9	Appendix B Paul Erdős	1446
18.10	References	1451

VI Appendix 1460

A	Introduction and Overview	1460
A.1	Extra Motivation	1460
A.2	The Mindset of Professionals	1460
A.2.1	Attitude to Probability Theory	1460
A.2.2	Features of Effective Learning	1461
A.3	Acknowledgements	1461

A.4	Literature	1462
A.4.1	Main Literature	1462
A.4.2	Extra Literature	1462
A.4.3	Literature About Applications	1463
A.4.4	History in a Nutshell	1464
A.5	Puzzles for Different Situations	1464
A.5.1	Puzzles for Interesting Discussion	1464
A.5.2	Puzzles About Everyday Life	1464
A.5.3	Selected Fun Typical Problems	1464
B	Mathematics for Probability Theory	1464
B.1	Another Mathematics for Probability Theory	1464
B.1.1	Constructions of the Theory of Measure	1464
B.1.2	Convergence Theorems	1464
B.1.3	Calculation of Integrals	1466
B.2	Elements of Discrete Mathematics	1466
B.2.1	Math.-phys.	1466
B.2.2	Combinatorics	1466
B.3	Programming for Probability Theory (!?!?!)	1467
B.4	Math	1467
B.4.1	A. 1 Sets	1467
B.4.2	A. 2 Functions	1470
B.4.3	A. 3 Matrices	1474
B.4.4	A. 4 Difference Equations	1476
B.4.5	A. 5 Differential Equations	1477
B.4.6	A. 6 Partial Derivatives	1477
B.4.7	A. 7 Multiple Integrals	1478
B.4.8	A. 8 Sums	1479
B.4.9	A. 9 Pattern Recognition	1481
B.4.10	A. 10 Common Sense and Checking Answers	1481
B.5	R	1482
B.5.1	B. 1 Vectors	1482
B.5.2	B. 2 Matrices	1482
B.5.3	B. 3 Math	1483
B.5.4	B. 4 Sampling and Simulation	1483
B.5.5	B. 5 Plotting	1483
B.5.6	B. 6 Programming	1483
B.5.7	B. 7 Summary Statistics	1484
B.5.8	B. 8 Distributions	1484
B.6	Table of Distributions	1485
B.7	References	1485
C	Bibliography	1487

Preface

Основная мотивация

Очень много приложений в обычной жизни

(раскрою позже)

Головоломки теории вероятностей для мотивации

Фокус с угадыванием конечного полученного числа в последовательности

(впишу, по сути то, что Маес показал. а почему так? вот - чтобы ответить, мотивация к изучению теорвера)

(сам пока не знаю, почему)

Part I

Probability Theory in a Nutshell

1 The Essence of Methods of Probability Theory

1.0.1 Probability Spaces

Математическая схематизация случайных явлений

Пространство элементарных событий

Случайные события,

Аксиомы теории вероятностей,

Следствия из аксиом

Примеры вероятностных пространств .Классическая схема Дискретное вероятностное пространство Геометрические вероятности Абсолютно непрерывные вероятностные пространства

Вероятностное пространство и вероятность

Пространство элементарных событий и события Ω
алгебра и события

(??)

Понятия теории множеств

Суммой множеств событий A и B называется их объединение $A \cup B$.

пересечение $A \cap B$ - произведением, $A \cap B \equiv AB$.

дополнение $A^c = \bar{A} = \Omega \setminus A$ - противоположным событием.

Если A и B - события, то говорят, что при каждом $\omega \in A \cap B$ события A и B происходят одновременно.

Если $A \cap B = \emptyset$, то события A и B называют несовместными.

Если $A \subseteq B$, то говорят, что событие A влечет событие B .

$\forall A, B, C \subseteq \Omega$

$$\begin{aligned} AB &= BA, & A \cup B &= B \cup A \\ (AB)C &= A(BC), & (A \cup B) \cup C &= A \cup (B \cup C) \\ A(B \cup C) &= AB \cup AC, & A \cup (BC) &= (A \cup B)(A \cup C) \end{aligned}$$

Также: (???)

$$\begin{aligned} \text{если } A \subset B, \text{ то } \bar{B} \subset \bar{A} \\ \overline{AB} = \bar{A} \cup \bar{B}, \\ \overline{A \cup B} = \bar{A} \cap \bar{B}, \\ A \cup B = A \cup (\bar{A} \cap B), \\ A \setminus B = A \setminus (A \cap \bar{B}) \end{aligned}$$

$$\overline{\bigcap_{k \geq 1} A_k} = \bigcup_{k \geq 1} \bar{A}_k, \quad \bigcup_{k \geq 1} A_k = \bigcap_{k \geq 1} \bar{A}_k$$

Алгебра множеств

Совокупность $\mathcal{A}$ подмножеств Ω называется алгеброй множеств, если выполнены следующие условия:

- 1) $\emptyset \in \mathcal{A}, \quad \Omega \in \mathcal{A}$
- 2) если $A \in \mathcal{A}$, то и $\Omega \setminus A = A^c \in \mathcal{A}$;
- 3) если $A, B \in \mathcal{A}$, то $A \cup B \in \mathcal{A}$ и $A \cap B \in \mathcal{A}$.

В условия 1) и 3) вторые части являются следствиями первых и условия 2).

Примеры алгебр множеств:

- а) совокупность всех подмножеств конечного множества Ω ,
- б) совокупность всех конечных объединений содержащихся в $(0, 1]$ замкнутых справа полуинтервалов.

σ -алгебра

Алгебра $\mathcal{F}$ подмножеств множества Ω называется σ -алгеброй, если для любой конечной или счетной совокупности множеств $A_1, A_2, \dots \in \mathcal{F}$

$$\bigcup_{n \geq 1} A_n \in \mathcal{F} \quad \text{и} \quad \bigcap_{n \geq 1} A_n \in \mathcal{F}$$

Множество всех подмножеств Ω , очевидно, является и алгеброй, и σ -алгеброй. Его выбирают в качестве множества событий $\mathcal{F}$, если пространство элементарных событий Ω конечно или счетно.

Пусть $\mathcal{F}_k, k \in K$, - конечная или бесконечная совокупность алгебр (или σ - алгебр) подмножеств множества Ω .

Пересечением алгебр (соответственно, σ -алгебр) $\mathcal{F}_k, k \in K$, называется совокупность всех таких подмножеств $A \subseteq \Omega$, что $A \in \mathcal{F}_k$ при любом $k \in K$.

(??) объединение алгебр может не быть алгеброй, в отличие от пересечений

Задача: проверить, что пересечение алгебр (σ -алгебр) является алгеброй (σ - алгеброй) и что обычное теоретико-множественное объединение алгебр (σ -алгебр) может не быть алгеброй (σ -алгеброй).

(хз, как проверить, потом сделаю это)

Борелевская σ -алгебра

Определение. Пусть M - совокупность подмножеств множества Ω .

Пересечение всех алгебр подмножеств множества Ω , которые содержат M , называется наименьшей алгеброй, порожденной M . Пересечение всех σ -алгебр

подмножеств множества Ω , которые содержат M , называется наименьшей σ -алгеброй, или борелевской σ -алгеброй, порожденной M .

Мера и типичные её свойства

Пусть $\mathcal{F}$ - алгебра или σ -алгебра подмножеств Ω .

Функция $\mu : \mathcal{F} \rightarrow \mathbb{R}$ называется **мерой**, если она аддитивна, т.е. $\mu(A \cup B) = \mu(A) + \mu(B)$ для любых таких $A, B \in \mathcal{F}$, что $A \cap B = \emptyset$.

Мера μ называется **счётно-аддитивной**, если $\mu(\cup_{k=1}^{\infty} A_k) = \sum_{k=1}^{\infty} \mu(A_k)$ для любой последовательности попарно не пересекающихся множеств $A_1, A_2, \dots \in \mathcal{F}$.

1.0.2 Conditional Probabilities, Independence of Events, Combinatorics

Перестановки, сочетания, размещения

Число всевозможных перестановок из n элементов равно $n!$:

$$P_n = n!$$

Доказывается по индукции, $P_k = kP_{k-1}$.

Условные вероятности

Вероятность произведения событий

Формула полной вероятности Формулы Байеса

Независимость событий

Примеры применения формул абсолютной вероятности Случайные блуждания Работа телефонной линии

1.0.3 Test Sequences

Конечные последовательности испытаний

Последовательность независимых испытаний

Предельные теоремы в схеме Бернулли

Бесконечные последовательности испытаний

1.0.4 Random Variables, Distributions

Определения и примеры случайных величин

Определение 4.1. Пусть $(\Omega, \mathcal{F}), (E, \mathcal{E})$ - измеримые пространства и $f : \Omega \rightarrow E$. Функция f называется $(\mathcal{F}|\mathcal{E})$ -измеримой, если

$$\forall A \in \mathcal{E} : f^{-1}(A) \in \mathcal{F}$$

Если $E = \mathbb{R}, \mathcal{E} = \mathcal{B}(\mathbb{R})$, то измеримая функция f называется случайной величиной. Если $E = \mathbb{R}^n, \mathcal{E} = \mathcal{B}(\mathbb{R}^n)$, то измеримая функция f называется случайным вектором. В случае, когда $(\Omega, \mathcal{F}) = (\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$, а $(E, \mathcal{E}) = (\mathbb{R}^k, \mathcal{B}(\mathbb{R}^k))$, то $(\mathcal{F}|\mathcal{E})$ -измеримые функции называются борелевскими.

Утверждение 4.1.1. Пусть $\mathcal{F}_\xi = \{A \in \mathcal{F} : \exists B \in \mathcal{B}(\mathbb{R}^k) : A = \xi^{-1}(B)\}$. Тогда $\mathcal{F}_\xi$ - σ -алгебра. Эта σ -алгебра называется σ -алгеброй, порожденной случайной величиной ξ . (Иногда она обозначается $\sigma(\xi)$)

Теорема 4.1. Критерий измеримости Пусть даны $(\Omega, \mathcal{F}), (E, \mathcal{E})$ - измеримые пространства, $M \subset \mathcal{E}, \sigma(M) = \mathcal{E}$. Тогда для $(\mathcal{F}|\mathcal{E})$ - измеримости функции $f : \Omega \rightarrow E$ необходимо и достаточно, чтобы $\forall B \in M f^{-1}(B) \in \mathcal{F}$.

Следствие 4.1.1. Для того чтобы функция $\xi : \Omega \rightarrow \mathbb{R}$ была случайной величиной, необходимо и достаточно, чтобы $\forall x \in \mathbb{R}$

$$\{\omega : \xi(\omega) < x\} \in \mathcal{F}$$

или

$$\{\omega : \xi(\omega) \leq x\} \in \mathcal{F}.$$

Доказательство следует из того, что каждая из систем множеств

$$\mathcal{E}_1 = \{x : x < c, c \in \mathbb{R}\}$$

$$\mathcal{E}_2 = \{x : x \leq c, c \in \mathbb{R}\}$$

порождают σ -алгебру $\mathcal{B}(\mathbb{R})$.

Определение 4.2. Распределением случайного n -мерного вектора ξ называется такая функция $P_\xi : \mathcal{B}(\mathbb{R}^n) \rightarrow [0, 1]$, что $P_\xi(B) = P(\xi \in B)$

Свойства функций распределения

Совместные распределения нескольких случайных величин

Независимость случайных величин

Определение 4.3. Пусть ξ - n -мерный случайный вектор, η - k -мерный. Тогда они называются независимыми (Обозначение: $\xi \perp \eta$), если $\forall B_1 \in \mathcal{B}(\mathbb{R}^n), B_2 \in \mathcal{B}(\mathbb{R}^k) : P(\xi \in B_1, \eta \in B_2) = P(\xi \in B_1) P(\eta \in B_2)$

Случайные вектора $\xi_1, \xi_2, \dots, \xi_n$ называются попарно независимыми, если $\forall i \neq j \xi_i \perp \xi_j$

Определение 4.4. Пусть Ω - некоторое пространство. Система $\mathcal{P}$ подмножеств Ω называется π системой, если она замкнута относительно взятия конечных пересечений: $\forall A_1, \dots, A_n \in \mathcal{P} : \bigcap_{1 \leq k \leq n} A_k \in \mathcal{P}, n \geq 1$.

Система $\mathcal{L}$ подмножеств Ω называется λ -системой, если

1. $\Omega \in \mathcal{L}$
2. $(A, B \in \mathcal{L} \text{ и } A \subseteq B) \Rightarrow (B \setminus A \in \mathcal{L})$
3. $(A_n \in \mathcal{L}, n \geq 1, \text{ и } A_n \uparrow A) \Rightarrow (A \in \mathcal{L})$

Теорема 4.3. Критерий независимости. Пусть $\xi_1, \dots, \xi_n$ - случайные величины, и $\xi = (\xi_1, \dots, \xi_n)$. Пусть $F_\xi(x_1, \dots, x_n) = P(\xi_1 \leq x_1, \dots, \xi_n \leq x_n)$ - функция распределения случайного вектора ξ , а $F_{\xi_i}(x)$ - функция распределения ξ_i .

Тогда для того, чтобы случайные величины $\xi_1, \dots, \xi_n$ были независимы, необходимо и достаточно, чтобы для всех векторов $(x_1, \dots, x_n) \in \mathbb{R}^n$

$$F_\xi(x_1, \dots, x_n) = F_{\xi_1}(x_1) \dots F_{\xi_n}(x_n)$$

Функции от случайных величин

Многомерные распределения

Определение 3.6. $\mathcal{B}(\mathbb{R}^n)$ - наименьшая σ -алгебра, такая, что:

$$\begin{aligned} \mathcal{B}(\mathbb{R}^n) &= \sigma(\{(-\infty, a_1] \times (-\infty, a_2] \times \dots \times (-\infty, a_n] ; a_i \in \mathbb{R}\}) = \\ &= \sigma(\{B_1 \times B_2 \times \dots \times B_n ; \forall i B_i \in \mathcal{B}(\mathbb{R})\}) \end{aligned}$$

Определение 3.7. Если P - вероятностная мера на $\mathcal{B}(\mathbb{R}^n)$, то P называется n -мерным распределением вероятностей.

Определение 3.8. $F : \mathbb{R}^n \rightarrow [0, 1]$ называется функцией распределения, соответствующей распределению P , если

$$F(x_1, \dots, x_n) = P((-\infty, x_1] \times \dots \times (-\infty, x_n])$$

Теорема 3.3. Функция распределения, соответствующая распределению P , обладает следующими свойствами:

1. $\lim_{x_1 \rightarrow +\infty} F(x) = 1, \forall i \in \{1, \dots, n\} \lim_{x_i \rightarrow -\infty} F(x) = 0, x_n \rightarrow +\infty$
2. Пусть $x = (x_1, \dots, x_n) \geq (y_1, \dots, y_n) = y \iff \forall i \in \{1, \dots, n\} x_i \geq y_i$,

$$x^{(k)} \downarrow x \iff \left(x^{(k+1)} \leq x^{(k)} \wedge \forall i \in \{1, \dots, n\} x_i^{(k)} \rightarrow x_i \right). \text{ Тогда:}$$

$$\lim_{k \rightarrow +\infty} F(x^{(k)}) = F(x)$$

3. $\Delta_{a_1 b_1} \dots \Delta_{a_n b_n} F(x) \geq 0$

Алгоритм определения составной случайной величины по известным распределениям ее частей

1.0.5 Numerical Characteristics of Random Variables

Математическое ожидание Определения и примеры

Свойства математического ожидания

Дисперсия

Ковариация Коэффициент корреляции

Условные распределения и условные математические ожидания

1.0.6 Chebyshev's Inequality, the Law of Large Numbers

Неравенство Чебышева

Неравенство Чебышева (для неотрицательных случайных величин). Пусть случайная величина $\xi \geq 0$ и имеет конечное математическое ожидание $E\xi < \infty$. Тогда $\forall \varepsilon > 0$

$$P(\xi \geq \varepsilon) \leq \frac{E\xi}{\varepsilon}$$

Доказательство.

$$\xi \equiv \xi_1 + \xi_2$$

$$\xi_1 = \begin{cases} \xi, & \text{если } \xi \geq \varepsilon, \\ 0, & \text{если } \xi < \varepsilon, \end{cases} \quad \xi_2 = \begin{cases} 0, & \text{если } \xi \geq \varepsilon, \\ \xi, & \text{если } \xi < \varepsilon. \end{cases}$$

Поскольку случайная величина $\xi \geq 0$, то и $\xi_1 \geq 0$, и $\xi_2 \geq 0$. Тогда

$$E\xi = E\xi_1 + E\xi_2 \geq E\xi_1 \geq \varepsilon \cdot P(\xi \geq \varepsilon)$$

Разделив это неравенство на $\varepsilon > 0$, получим неравенство Чебышева.

Неравенство Чебышева (в общем виде). Пусть случайная величина ξ имеет конечную дисперсию $D\xi < \infty$. Тогда для любого $\varepsilon > 0$ выполнено неравенство

$$P(|\xi - E\xi| \geq \varepsilon) \leq \frac{D\xi}{\varepsilon^2}$$

Доказательство. Из неравенства Чебышева для неотрицательных случайных величин получаем

$$P(|\xi - E\xi| \geq \varepsilon) = P((\xi - E\xi)^2 \geq \varepsilon^2) \leq \frac{E(\xi - E\xi)^2}{\varepsilon^2} = \frac{D\xi}{\varepsilon^2}$$

Закон больших чисел

Закон больших чисел. Если у случайных величин

$\xi_1, \xi_2, \dots, \xi_n, \dots$ существуют дисперсии и существует предел

$$\lim_{n \rightarrow \infty} \frac{D\xi_1 + D\xi_2 + \dots + D\xi_n}{n^2} = 0$$

то к случайным величинам

$$\xi_1, \xi_2, \dots, \xi_n, \dots$$

применим закон больших чисел, а именно, справедливо равенство

$$\lim_{n \rightarrow \infty} P \left\{ \left| \frac{\xi_1 + \xi_2 + \dots + \xi_n}{n} - \frac{E\xi_1 + E\xi_2 + \dots + E\xi_n}{n} \right| < \varepsilon \right\} = 1$$

Усиленный закон больших чисел. Если случайные величины $\xi_1, \xi_2, \dots, \xi_n, \dots$ одинаково распределены, попарно независимы и имеют конечное математическое ожидание $E\xi_1 = a$, то выполнено равенство

$$P \left\{ \lim_{n \rightarrow \infty} \frac{\xi_1 + \xi_2 + \dots + \xi_n}{n} = a \right\} = 1$$

1.0.7 Generating and Characteristic Functions

Производящие функции Определение и свойства

Характеристические функции Определения и свойства Задачи к главе

1.0.8 Limit Theorems

Закон больших чисел

Центральная предельная теорема

Classical CLT

Let $\{X_1, \dots, X_n\}$ be a sequence of i.i.d. random variables having a distribution with expected value given by μ and finite variance given by σ^2 . Suppose we are interested in the sample average

$$\bar{X}_n \equiv \frac{X_1 + \dots + X_n}{n}.$$

By the law of large numbers, the sample average converges almost surely (and therefore also converges in probability) to the expected value μ as $n \rightarrow \infty$.

The classical central limit theorem describes the size and the distributional form of the stochastic fluctuations around the deterministic number μ during this convergence. More precisely, it states that as n gets larger, the distribution of the difference between the sample average $\bar{X}_n$ and its limit μ , when multiplied by the factor $\sqrt{n}$ - that is, $\sqrt{n}(\bar{X}_n - \mu)$ - approaches the normal distribution with mean 0 and variance σ^2 . For large enough n , the distribution of $\bar{X}_n$ gets arbitrarily close to the normal distribution with mean μ and variance σ^2/n .

The usefulness of the theorem is that the distribution of $\sqrt{n}(\bar{X}_n - \mu)$ approaches normality regardless of the shape of the distribution of the individual X_i . Formally, the theorem can be stated as follows:

Lindeberg-Lévy CLT - Suppose $X_1, X_2, X_3, \dots$ is a sequence of i.i.d. random variables with $E[X_i] = \mu$ and $\text{Var}[X_i] = \sigma^2 < \infty$. Then, as n approaches infinity, the random variables $\sqrt{n}(\bar{X}_n - \mu)$ converge in distribution to a normal $\mathcal{N}(0, \sigma^2)$ [4]

$$\sqrt{n}(\bar{X}_n - \mu) \xrightarrow{d} \mathcal{N}(0, \sigma^2).$$

In the case $\sigma > 0$, convergence in distribution means that the cumulative distribution functions of $\sqrt{n}(\bar{X}_n - \mu)$ converge pointwise to the cdf of the $\mathcal{N}(0, \sigma^2)$ distribution: for every real number z ,

$$\lim_{n \rightarrow \infty} \mathbb{P}[\sqrt{n}(\bar{X}_n - \mu) \leq z] = \lim_{n \rightarrow \infty} \mathbb{P}\left[\frac{\sqrt{n}(\bar{X}_n - \mu)}{\sigma} \leq \frac{z}{\sigma}\right] = \Phi\left(\frac{z}{\sigma}\right),$$

where $\Phi(z)$ is the standard normal cdf evaluated at z . The convergence is uniform in z in the sense that

$$\lim_{n \rightarrow \infty} \sup_{z \in \mathbb{R}} \left| \mathbb{P}[\sqrt{n}(\bar{X}_n - \mu) \leq z] - \Phi\left(\frac{z}{\sigma}\right) \right| = 0,$$

where sup denotes the least upper bound (or supremum) of the set. [5]

Lyapunov CLT

In this variant of the central limit theorem the random variables X_i have to be independent, but not necessarily identically distributed. The theorem also requires that random variables $|X_i|$ have moments of some order $(2 + \delta)$, and that the rate of growth of these moments is limited by the Lyapunov condition given below. Lyapunov CLT [6] - Suppose $\{X_1, \dots, X_n, \dots\}$ is a sequence of independent random variables, each with finite expected value μ_i and variance σ_i^2 . Define

$$s_n^2 = \sum_{i=1}^n \sigma_i^2$$

If for some $\delta > 0$, Lyapunov's condition

$$\lim_{n \rightarrow \infty} \frac{1}{s_n^{2+\delta}} \sum_{i=1}^n E[|X_i - \mu_i|^{2+\delta}] = 0$$

is satisfied, then a sum of $\frac{X_i - \mu_i}{s_n}$ converges in distribution to a standard normal random variable, as n goes to infinity:

$$\frac{1}{s_n} \sum_{i=1}^n (X_i - \mu_j) \xrightarrow{d} \mathcal{N}(0, 1).$$

In practice it is usually easiest to check Lyapunov's condition for $\delta = 1$. If a sequence of random variables satisfies Lyapunov's condition, then it also satisfies Lindeberg's condition. The converse implication, however, does not hold.

Lindeberg CLT

The Lyapunov condition can be replaced with the following weaker one:

$$\forall \varepsilon > 0 \quad \lim_{n \rightarrow \infty} \frac{1}{s_n^2} \sum_{i=1}^n \mathbb{E} \left[(X_i - \mu_i)^2 \cdot \mathbf{1}_{\{|X_i - \mu_i| > \varepsilon s_n\}} \right] = 0 \quad \Rightarrow$$

$$\frac{1}{s_n} \sum_{i=1}^n (X_i - \mu_i) \rightarrow \mathcal{N}(0, 1)$$

Lyapunov CLT

In this variant of the central limit theorem the random variables X_i have to be independent, but not necessarily identically distributed. The theorem also requires that random variables $|X_i|$ have moments of some order $(2 + \delta)$, and that the rate of growth of these moments is limited by the Lyapunov condition given below. Lyapunov CLT^[6] - Suppose $\{X_1, \dots, X_n, \dots\}$ is a sequence of independent random variables, each with finite expected value μ_i and variance σ_i^2 . Define

$$s_n^2 = \sum_{i=1}^n \sigma_i^2$$

If for some $\delta > 0$,

$$\lim_{n \rightarrow \infty} \frac{1}{s_n^{2+\delta}} \sum_{i=1}^n \mathbb{E} \left[|X_i - \mu_i|^{2+\delta} \right] = 0 \quad (\text{Lyapunov's condition})$$

$\Rightarrow$ sum of $\frac{X_i - \mu_i}{s_n}$ converges in distribution to a standard normal random variable, as n goes to infinity:

$$\frac{1}{s_n} \sum_{i=1}^n (X_i - \mu_j) \xrightarrow{d} \mathcal{N}(0, 1).$$

In practice it is usually easiest to check Lyapunov's condition for $\delta = 1$. If a sequence of random variables satisfies Lyapunov's condition, then it also satisfies Lindeberg's condition. The converse implication, however, does not hold.

2 Other Very Useful or Basic Methods in a Nutshell

2.0.1 Other Measure-probability Properties

Some Inequalities

Boole's inequality: for a countable set of events $A_1, A_2, A_3, \dots$, we have

$$\mathbb{P} \left(\bigcup_{i=1}^{\infty} A_i \right) \leq \sum_{i=1}^{\infty} \mathbb{P}(A_i)$$

In measure-theoretic terms, Boole's inequality follows from the fact that a measure (and certainly any probability measure) is σ -subadditive.

(?? idea of proof??)

(?? applications??)

Bonferroni inequalities:

$$S_1 := \sum_{i=1}^n \mathbb{P}(A_i), \quad S_2 := \sum_{1 \leq i_1 < i_2 \leq n} \mathbb{P}(A_{i_1} \cap A_{i_2}), \quad \dots, \quad S_k := \sum_{1 \leq i_1 < \dots < i_k \leq n} \mathbb{P}(A_{i_1} \cap \dots \cap A_{i_k})$$

for all integers k in $\{1, \dots, n\}$.

$$\sum_{j=1}^K (-1)^{j-1} S_j \geq \mathbb{P}\left(\bigcup_{i=1}^n A_i\right) = \sum_{j=1}^n (-1)^{j-1} S_j \quad K \leq n \text{ is odd}$$

$$\sum_{j=1}^K (-1)^{j-1} S_j \leq \mathbb{P}\left(\bigcup_{i=1}^n A_i\right) = \sum_{j=1}^n (-1)^{j-1} S_j \quad K \leq n \text{ is even}$$

(?? idea of proof??)

(?? applications??)

2.0.2 Fundamentals of Mathematical Statistics

Задачи математической статистики Понятие выборки

Оценка неизвестных параметров распределения по выборке Точечные оценки Метод наибольшего правдоподобия для нахождения оценок параметров Метод моментов Интервальные оценки

Статистическая проверка гипотез Проверка гипотез о законе распределения Выбор из двух гипотез

2.0.3 Methods for Proofs of Probabilistic Inequalities (!?!)

(есть много задач про это, но пока не до этого)

2.0.4 Markov Chains

Определение Основные свойства

Уравнение для вероятностей перехода:

Стационарное распределение Теорема о предельных вероятностях

О Марковских цепях с непрерывным временем

Вводим (??)

Основное - управляющее уравнение

$$\vec{\pi}(n+1) = \hat{P} \cdot \vec{\pi}(n).$$

Чаще всего известно распределение вероятностей в начальный момент времени $\vec{\pi}(0)$, тогда

$$\vec{\pi}(n) = \hat{P}^n \cdot \vec{\pi}(0).$$

Отсюда

$$\begin{aligned}\pi_1(n) &= \frac{q}{p+q} + \frac{p}{p+q}(1-p-q)^n \pi_1(0) - \frac{q}{p+q}(1-p-q)^n \pi_2(0) \\ \pi_2(n) &= \frac{p}{p+q} - \frac{p}{p+q}(1-p-q)^n \pi_1(0) + \frac{q}{p+q}(1-p-q)^n \pi_2(0).\end{aligned}$$

Это потому что мы представляем матрицу в виде $\hat{P} = \hat{Q}\hat{D}\hat{Q}^{-1}$, где $\hat{D}$ - диагональная матрица, на диагонали которой стоят собственные числа матрицы $\hat{P}$, $\hat{Q}$ - матрица, составленная из собственных столбцов матрицы $\hat{P}$. Решая уравнение $\det(\hat{P} - \lambda\hat{I}) = 0$ находим пару собственных чисел $\lambda_1 = 1$ и $\lambda_2 = 1 - p - q$, которым соответствуют собственные столбцы $\vec{a}_1 = (1/p, 1/q)^T$, $\vec{a}_2 = (1, -1)^T$, по определению удовлетворяющие условию $\hat{P}\vec{a} = \lambda\vec{a}$. Используя эти результаты, запишем

$$\begin{aligned}\hat{P}^n &= \hat{Q}\hat{D}^n\hat{Q}^{-1} = \begin{pmatrix} \frac{1}{p} & 1 \\ \frac{1}{q} & -1 \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ 0 & 1-p-q \end{pmatrix}^n \cdot \begin{pmatrix} \frac{pq}{p+q} & \frac{pq}{p+q} \\ \frac{p}{p+q} & -\frac{q}{p+q} \end{pmatrix} \\ &= \begin{pmatrix} \frac{q}{p+q} + \frac{p}{p+q}(1-p-q)^n & \frac{q}{p+q} - \frac{q}{p+q}(1-p-q)^n \\ \frac{p}{p+q} - \frac{p}{p+q}(1-p-q)^n & \frac{p}{p+q} + \frac{q}{p+q}(1-p-q)^n \end{pmatrix}\end{aligned}$$

О Марковских цепях с дискретным временем

О марковских цепях в последовательностях случайных событий

(чисто для задач теорвера подготовка)

О других приложениях марковских цепей

(их много, потом укажу)

О Марковских цепях в кинетике

О марковских цепях в системах массового обслуживания (??)

О другом про марковские цепи

(наверняка там многое еще есть.)

2.0.5 Probability Theory Famous Paradoxes (!!!?)

(продумаю потом их, интересно!)

- Bertrand's box paradox: A paradox of conditional probability closely related to the Boy or Girl paradox.

- Bertrand's paradox: Different common-sense definitions of randomness give quite different results.
- Birthday paradox: In a random group of only 23 people, there is a better than 50/50 chance two of them have the same birthday.
- Borel's paradox: Conditional probability density functions are not invariant under coordinate transformations.
- Boy or Girl paradox: A two-child family has at least one boy. What is the probability that it has a girl?
- False positive paradox: A test that is accurate the vast majority of the time could show you have a disease, but the probability that you actually have it could still be tiny.
- Grice's paradox: Shows that the exact meaning of statements involving conditionals and probabilities is more complicated than may be obvious on casual examination.
- Intransitive dice: One can have three dice, called A, B, and C, such that A is likely to win in a roll against B, B is likely to win in a roll against C, and C is likely to win in a roll against A.
- Monty Hall problem, also known as the Monty Hall paradox. ^[2] An unintuitive consequence of conditional probability.
- Necktie paradox: A wager between two people seems to favour them both. Very similar in essence to the Two-envelope paradox.
- Proebsting's paradox: The Kelly criterion is an often optimal strategy for maximizing profit in the long run. Proebsting's paradox apparently shows that the Kelly criterion can lead to ruin.
- Sleeping Beauty problem: A probability problem that can be correctly answered as one half or one third depending on how the question is approached.

- Three Prisoners problem, also known as the Three Prisoners paradox: ^[3] A variation of the Monty Hall problem.

- Two-envelope paradox: You are given two indistinguishable envelopes, each of which contains a positive sum of money. One envelope contains twice as much as the other. You may pick one envelope and keep whatever amount it contains. You pick one envelope at random but before you open it you are given the chance to take the other envelope instead.

2.0.6 Basics of Random Processes in a Nutshell (!!!?)

Пуассоновский процесс (!!!)

Обзор методов случайных процессов

Решения через усреднения

(так у Белана домашка и делалась)

Решения через управляющее уравнение Маркова

Такую модель можно составить.

Решения через интеграл по траекториям

К нему не сложно прийти, а потом мб что-то можно с ним сделать.
(потом разберусь с этим методом.)

Part II

Fundamentals of Probability Theory

Приведем подробную теорию вероятностей на необходимом уровне строгости.

3 Fundamentals of Probability Theory

Остановимся на концепциях, которые будем дальше использовать.
(все строго в рамках приложений к теории вероятностей)

3.1 Fundamentals of Set Theory and Measure Theory

3.1.1 When and Why Do we Need Set Theory for Probability Tasks?

(answer will be written here)

3.1.2 Set Theory for Probability Theory (??)

(тут куча всего, что мало нужно, однако этим занимаются при изучении теорвера)
Обсудим формализм теории множеств.

применение теории множеств для теории вероятностей

Построение теоретико-множественной модели теории вероятностей связано с введением пространства Ω элементарных событий, называемых исходами. В результате каждого проведения случайного эксперимента может наблюдаться только один элементарный исход ω , при этом любое событие A ассоциируется с некоторым набором исходов, то есть подмножеством в Ω . Например, Ω также рассматривается как событие, его называют достоверным событием, поскольку оно происходит всегда, а $\emptyset$ связывают с невозможным событием.

Теоретико-множественные операции и отношения соответствуют логическим операциям над событиями. Если $A \in B$, то это означает, что наступление события A влечет наступление события B . Событие $\bar{A} = \Omega \setminus A$ называется отрицанием события A .

Объединению $A \cup B$ соответствует событие, которое происходит в случае, когда произошло A или B .

Пересечению $A \cap B \equiv AB$ соответствует событие, которое происходит лишь в случае, когда происходят одновременно и A и B . Аналогичный смысл имеют объединение и пересечение любого числа событий.

Разность $A \setminus B$ означает событие, когда A произошло, но B не произошло.

Далее вероятность будет определена как мера на множествах событий, поэтому научимся работать с множествами, после - с мерой, а потом займемся теорией вероятностей.

Типичные операции над множествами

Приведем свойства операций над множествами, которые часто используются. Все эти свойства легко понять с помощью кругов Эйлера.

Операции объединения и пересечения множеств обладают свойством дистрибутивности:

$$A \cap \left(\bigcup_{i \in I} B_i \right) = \bigcup_{i \in I} (A B_i), \quad A \cup \left(\bigcap_{i \in I} B_i \right) = \bigcap_{i \in I} (A \cup B_i)$$

Для дополнений выполняются правила де Моргана:

$$\overline{\bigcup_{i \in I} A_i} = \bigcap_{i \in I} \overline{A_i}, \quad \overline{\bigcap_{i \in I} A_i} = \bigcup_{i \in I} \overline{A_i}$$

пределы множеств (????)

Пределы множеств определяются так:

$$A^* = \overline{\lim}_{n \rightarrow \infty} A_n = \bigcap_{n=1}^{\infty} \bigcup_{k \geq n} A_k, \quad A_* = \underline{\lim}_{n \rightarrow \infty} A_n = \bigcup_{n=1}^{\infty} \bigcap_{k \geq n} A_k$$

Если $A_1, A_2, \dots$ интерпретировать как события, то верхний предел A^* выражает то, что произошло бесконечно много событий последовательности, а нижний предел A_* соответствует тому, что произошли все события последовательности, за исключением, быть может, конечного числа.

В случае $A_1 \subset A_2 \subset \dots$ последовательность $\{A_n\}$ называется монотонно возрастающей и тогда $A^* = A_* = \bigcup_{n=1}^{\infty} A_n$. Аналогично, в случае $A_1 \supset A_2 \supset \dots$ последовательность $\{A_n\}$ называется монотонно убывающей и $A^* = A_* = \bigcap_{n=1}^{\infty} A_n$. В общем случае последовательность $\{A_n\}$ называется сходящейся, если $A^* = A_*$.

Типы множеств

Часто используются другие понятия теории множеств, с которыми тут познакомимся.

Definition 3.1. Семейство $\mathcal{F}$ подмножеств множества X называется *дизъюнктным*, если элементы $\mathcal{F}$ попарно не пересекаются, то есть для любых $A, B \in \mathcal{F}$ из $A \cap B \neq \emptyset$ следует $A = B$

(тут много дополнять нужно)

Алгебры множеств

(потом к каждому определению пояснения добавлю)

Определим понятия теории множеств, с которым будем далее работать.

Обзор понятий теории множеств

(позже напишу общую картину всего, пока не усвоено это)

кольца

Definition 3.2. кольцом называется семейство $\mathcal{R}$ подмножеств множества X (подмножеств), если

(??)

1. $\emptyset \in \mathcal{R}$

2. $\forall A, B \in \mathcal{R}$ подмножества $A \cup B, A \cap B$ и $A \setminus B \in \mathcal{R}$.

(оно нужно?)

Алгебра множеств

Кольцо расширяется до алгебры множеств.

Definition 3.3. (алгебра множеств) Алгеброй множеств называется множество $\mathcal{A}$ подмножеств Ω , если $\Omega \in \mathcal{A}$ и оно замкнуто относительно операций объединения, пересечения и взятия дополнения.

Примером кольца подмножеств R , не являющегося алгеброй, является семейство всех ограниченных подмножеств R . (?)

Чтобы понять, алгебра ли данная система подмножеств, иногда более удобно использовать эквивалентное определение:

Lemma 1. Множество $\mathcal{A}$ подмножеств Ω называется алгеброй множеств, если выполняются

- (i) $\Omega \in \mathcal{A}$
- (ii) Если $A \in \mathcal{A}$, то $\bar{A} \in \mathcal{A}$
- (iii) Если $A, B \in \mathcal{A}$, то $A \cup B \in \mathcal{A}$

Заметим, что здесь нужно объединять только конечных множеств.

Доказательство

(пока что скопированный текст из горяйного, потом пересмотрю)

Нам нужно показать, что $\mathcal{A}$ замкнуто относительно всех теоретико множественных операций.

Из условий (ii) и (iii) следует, что класс $\mathcal{A}$ замкнут относительно операции взятия дополнения и объединения. Остается показать, что класс $\mathcal{A}$ также замкнут относительно операции пересечения.

Пусть A и B из $\mathcal{A}$. Тогда в силу условия (ii) дополнения A , B принадлежат $\mathcal{A}$. Представление

$\mathcal{A}$ замкнут относительно операции взятия дополнения и объединения. Остает- Пусть A и B из $\mathcal{A}$.

Тогда в силу условия (ii) доолнения $\bar{A}, \bar{B}$ принадлежат

$$AB = \overline{\bar{A} \cup \bar{B}}$$

Часто мы работаем не с конечным пересечением, а со счетным, в этом случае говорят о сигма-алгебре.

Definition 3.4. σ -алгебра (сигма алгебра) - алгебра, в которой любое конечное или счетное семейство подмножеств даст также элемент из этой алгебры:

$$\bigcup_{n \geq 1} A_n \in \mathcal{F} \quad \text{и} \quad \bigcap_{n \geq 1} A_n \in \mathcal{F}$$

Например, для страховой компании можно посчитать вероятность того, что она никогда не разорится, придумав какую-то вероятностную модель в рамках которой эта компания работает. Введем событие "один" - компания не разорится никогда, событие "два" - разорится завтра, "три" - послезавтра, т.д. В итоге у нас имеется счетное число событий, и их объединение должно быть достоверным событием, для таких задач используется σ -алгебра.

(пример, где алгебра не сигма алгебра - ну я хз вообще, пока забьем, но вообще для понимания важно было бы понять это)

утв: на конечном множестве алгебра то же самое что и сигма алгебра. (???)

Алгебр и сигма-алгебр может быть много, а мы хотим создавать простейшие. Для этого нам понадобятся борелевские сигма-алгебры.

Definition 3.5. Борелевская σ -алгебра, порожденной M - совокупностью подмножеств множества Ω - это пересечение всех алгебр подмножеств множества Ω , которые содержат M .

Также она называется наименьшей алгеброй, порожденной M .

ну и самое интересное - примеры.

Definition 3.6. события - подмножества в борелевской сигма алгебре

в общем, это все просто сильные слова, на самом деле за ними нет ничего кроме как понимание, что мы можем делать с этими элементами, ну в точности то же, что и подсказывает нам интуиция.

вывод можно сделать только такой, что типа по сути все - подмножества и все типа. по сути это только введение.

ну и здесь есть оффтоп - связь с теорией меры, позже обязательно свяжу, сейчас нет возможности.

как измеримость и работы лебега в этом всем находятся?

разбиения

(зачем вообще они???)

В множестве можно выделять подмножества. Если они покрывают все множество, то это очень хорошо и это называется разбиением.

Definition 3.7. (Разбиения)

Совокупность событий $D = D_1, \dots, D_n$ будем называть разбиением, если $D_i D_j = \emptyset$ при $i \neq j$ и $\cup_{i=1}^n D_i = W$. Множества D_i называются в этом случае атомами разбиения D .

Таким образом, разбиения для теории вероятностей

Definition 3.8. Пусть $(\Omega, \mathcal{F}, \mathbf{P})$ - вероятностное пространство. Систему событий $A_1, A_2, \dots \in \mathcal{F}$ называют разбиением пространства Ω , если эти события попарно несовместны, а их объединение совпадает с Ω :

$$A_i \cap A_j = \emptyset (1 \leq i < j), \quad A_1 \cup A_2 \cup \dots = \Omega$$

Справедливо и обратное утверждение. Теорема. Каждая конечная алгебра множеств порождается конечным разбиением.

Доказательство

(скопировано, пока что пофиг)

Доказательство. Пусть $\mathcal{F}$ - конечная алгебра подмножеств Ω . Для каждого $\omega \in \Omega$ ПОЛОЖИМ

$$A_\omega = \bigcap_{\{A \in \mathcal{F} : \omega \in A\}} A$$

Покажем, что совокупность Γ всех попарно различных множеств A_ω образует разбиение Ω , порождающее алгебру $\mathcal{F}$, т. е. что

$$\bigcup_{\omega \in \Omega} A_\omega = \Omega$$

если $\omega_1, \omega_2 \in \Omega$, То A_{ω_1} и A_{ω_2} Либо совпадают, либо не пересекаются,

Равенство (9) очевидно, так как по определению $\omega \in A_\omega$ для каждого $\omega \in \Omega$ и совокупность $\{A_\omega\}_{\omega \in \Omega}$ отличается от Γ лишь повторами элементов.

Для доказательства утверждения (10) заметим, что при любых $\omega \in \Omega$ и $B \in \mathcal{F}$

$$\omega \in B \Rightarrow A_\omega \subseteq B$$

Покажем, что если $A_\omega \cap A_\kappa \neq \emptyset$, то $A_\omega = A_\kappa$. Пусть $\beta \in A_\omega \cap A_\kappa$. В силу (12) тогда

$$A_\beta \subseteq A_\omega \text{ и } A_\beta \subseteq A_\kappa$$

Допустим, что $\omega \notin A_\beta$. Тогда $\omega \in A_\omega \setminus A_\beta$ и в силу (12)

$$A_\omega \subseteq A_\omega \setminus A_\beta \subset A_\omega \setminus \{\beta\}$$

что невозможно, так как $\beta \in A_\omega$; Значит, $A_\beta \cap A_\omega \neq \emptyset$. Следовательно, $\omega \in A_\beta$. Но тогда $A_\omega \subseteq A_\beta$ в силу (12), а отсюда и из (13) вытекает, что $A_\omega = A_\beta$. Так же доказывается, что $A_\kappa = A_\beta$. Тем самым (10) доказано. Из (9) и (10) следует, что совокупность Γ образует разбиение Ω . Это разбиение конечно, так как алгебра $\mathcal{F}$ конечна и любое $A_\omega \in \mathcal{F}$. Из (12) следует, что $\omega \in A_\omega \subseteq B$ для любых $\omega \in B, B \in \mathcal{F}$, поэтому

$$B = \bigcup_{\omega \in B} \omega \subseteq \bigcup_{\omega \in B} A_\omega, \quad \bigcup_{\omega \in B} A_\omega \supseteq B \Rightarrow B = \bigcup_{\omega \in B} A_\omega$$

т. е. конечное разбиение Γ порождает алгебру $\mathcal{F}$. Теорема доказана.

независимые алгебры и разбиения

Definition 3.9. Пусть $(\Omega, \mathcal{F}, \mathbf{P})$ - вероятностное пространство. Независимыми разбиения

$$\{A_{11}, A_{12}, \dots\}, \{A_{21}, A_{22}, \dots\}, \dots, \{A_{m1}, A_{m2}, \dots\}$$

пространства Ω на множества $A_{ij} \in \mathcal{F}$ являются тогда, когда

$$\mathbf{P}(A_{1i_1} \cap A_{2i_2} \cap \dots \cap A_{mi_m}) = \mathbf{P}(A_{1i_1}) \mathbf{P}(A_{2i_2}) \dots \mathbf{P}(A_{mi_m})$$

для любых допустимых значений $i_1, \dots, i_m \geq 1$.

Definition 3.10. Независимыми алгебрами или σ -алгебрами событий $\mathcal{F}_1, \dots, \mathcal{F}_m \subseteq \mathcal{F}$ являются тогда, когда

$$\mathbf{P}(A_1 \cap \dots \cap A_m) = \mathbf{P}(A_1) \dots \mathbf{P}(A_m)$$

для любых событий $A_k \in \mathcal{F}_k, \quad k = 1, \dots, m$

(и еще тут теорему зубкова про все это. хз, нужна ли она??)

Основы теории меры

(только основные следствия теории меры, которые будем дальше использовать)

Definition 3.11. (

определение*меры по Гусеву)***Definition 3.12.****Definition 3.13.****вероятностные меры**

(отдельно теория меры для вероятностных мер)

Для теории вероятностей определяют на алгебрах событий меру.

Definition 3.14. Мера μ на σ -алгебре F подмножеств Ω называется вероятностной, если все ее значения лежат в отрезке $[0, 1]$ и $\mu(\Omega) = 1$

здесь пара заметок из тем гусева, очень даже нужны они тут, чтобы понимать, что и зачем мы сделали?

(а также из карасевника пара выписок)

посмотрим на свойства вероятностных мер

Теорема

Теорема Каратеодори.

Если на алгебре A подмножеств пространства Ω определена вероятностная мера μ , обладающая свойством счетной аддитивности, то эту меру можно продолжить на наименьшую σ -алгебру $F = \sigma(A)$, порожденную алгеброй A с сохранением свойства счетной аддитивности.

(????)

продолжения мер

(зачем они???)

Независимость разбиений, алгебр и σ -алгебр**Суть****О случаях применения****Теория**Определение. Пусть $\gamma = \{\gamma_1, \gamma_2, \dots\}$, $\gamma_k \subseteq \Omega$, $k = 1, 2, \dots$ - совокупность подмножеств Ω . Наименьшая алгебра (σ -алгебра) множеств $\mathcal{F}(\gamma)$, содержащая γ , называется алгеброй (σ -алгеброй), порожденной γ .Определение. Пусть $(\Omega, \mathcal{F}, \mathbf{P})$ - вероятностное пространство. Систему событий $A_1, A_2, \dots \in \mathcal{F}$ будем называть разбиением пространства Ω , если эти события попарно несовместны, а их объединение совпадает с Ω :

$$A_i \cap A_j = \emptyset (i \neq j), \quad \bigcup_{n \geq 1} A_n = \Omega$$

Разбиения называют еще полными системами несовместных событий. Замечание. Если $\gamma = \{\gamma_1, \gamma_2, \dots, \gamma_n\}$ - конечное разбиение Ω , то алгебра $\mathcal{F}(\gamma)$ порожденная разбиением γ , является конечной и состоит из пустого множества $\emptyset$ и всех множеств вида

$$\gamma_{i_1} \cup \gamma_{i_2} \cup \dots \cup \gamma_{i_k}, \quad k = 1, \dots, n, \quad 1 \leq i_1 < \dots < i_k \leq n$$

Из этого замечания следует, что если в разбиении $\gamma = \{\gamma_1, \gamma_2, \dots, \gamma_n\}$ все n множеств не пустые, то $\mathcal{F}(\gamma)$ состоит из 2^n элементов. Справедливо и обратное утверждение. Теорема. Каждая конечная алгебра множеств порождается конечным разбиением.

Доказательство. Пусть $\mathcal{F}$ - конечная алгебра подмножеств Ω . Для каждого $\omega \in \Omega$ положим

$$A_\omega = \bigcap_{\{A \in \mathcal{F} : \omega \in A\}} A$$

Покажем, что совокупность Γ всех попарно различных множеств A_ω образует разбиение Ω , порождающее алгебру $\mathcal{F}$, т. е. что

$$\bigcup_{\omega \in \Omega} A_\omega = \Omega$$

если $\omega_1, \omega_2 \in \Omega$, то A_{ω_1} и A_{ω_2} либо совпадают, либо не пересекаются, любое множество $B \in \mathcal{F}$ можно представить в виде $B = \bigcup_{\omega \in B} A_\omega$.

Равенство (9) очевидно, так как по определению $\omega \in A_\omega$ для каждого $\omega \in \Omega$ и совокупность $\{A_\omega\}_{\omega \in \Omega}$ отличается от Γ лишь повторами элементов.

Для доказательства утверждения (10) заметим, что при любых $\omega \in \Omega$ и $B \in \mathcal{F}$

$$\omega \in B \Rightarrow A_\omega \subseteq B$$

Покажем, что если $A_\omega \cap A_\kappa \neq \emptyset$, то $A_\omega = A_\kappa$. Пусть $\beta \in A_\omega \cap A_\kappa$. В силу (12) тогда

$$A_\beta \subseteq A_\omega \text{ и } A_\beta \subseteq A_\kappa$$

Допустим, что $\omega \notin A_\beta$. Тогда $\omega \in A_\omega \setminus A_\beta$ и в силу (12)

$$A_\omega \subseteq A_\omega \setminus A_\beta \subset A_\omega \setminus \{\beta\}$$

что невозможно, так как $\beta \in A_\omega$; значит, $A_\beta \cap A_\omega \neq \emptyset$. Следовательно, $\omega \in A_\beta$. Но тогда $A_\omega \subseteq A_\beta$ в силу (12), а отсюда и из (13) вытекает, что $A_\omega = A_\beta$. Так же доказывается, что $A_\kappa = A_\beta$. Тем самым (10) доказано. Из (9) и (10) следует, что совокупность Γ образует разбиение Ω . Это разбиение конечно, так как алгебра $\mathcal{F}$ конечна и любое $A_\omega \in \mathcal{F}$. Из (12) следует, что $\omega \in A_\omega \subseteq B$ для любых $\omega \in B, B \in \mathcal{F}$, поэтому

$$B = \bigcup_{\omega \in B} \omega \subseteq \bigcup_{\omega \in B} A_\omega, \quad \bigcup_{\omega \in B} A_\omega \supseteq B \Rightarrow B = \bigcup_{\omega \in B} A_\omega$$

т. е. конечное разбиение Γ порождает алгебру $\mathcal{F}$. Теорема доказана.

Задача. Доказать, что если пространство элементарных событий Ω счетно, то любая σ -алгебра его подмножеств порождается некоторым разбиением. Задача. Может ли бесконечная σ -алгебра быть счетным множеством? Определение. Пусть $(\Omega, \mathcal{F}, \mathbf{P})$ - вероятностное пространство. Разбиения

$$\{A_{11}, A_{12}, \dots\}, \{A_{21}, A_{22}, \dots\}, \dots, \{A_{m1}, A_{m2}, \dots\}$$

пространства Ω на множества $A_{ij} \in \mathcal{F}$ называются независимыми, если

$$\mathbf{P}(A_{1i_1} \cap A_{2i_2} \cap \dots \cap A_{mi_m}) = \mathbf{P}(A_{1i_1}) \mathbf{P}(A_{2i_2}) \dots \mathbf{P}(A_{mi_m})$$

для любых допустимых значений $i_1, \dots, i_m \geq 1$. Алгебры или σ -алгебры событий $\mathcal{F}_1, \dots, \mathcal{F}_m \subseteq \mathcal{F}$ называются независимыми, если

$$\mathbf{P}(A_1 \cap \dots \cap A_m) = \mathbf{P}(A_1) \dots \mathbf{P}(A_m)$$

для любых событий $A_k \in \mathcal{F}_k, k = 1, \dots, m$. Теорема. Если $(\Omega, \mathcal{F}, \mathbf{P})$ - вероятностное пространство и σ -алгебры $\mathcal{F}_1, \dots, \mathcal{F}_m \subseteq \mathcal{F}$ порождаются конечными или счетными разбиениями Ω , то σ -алгебры $\mathcal{F}_1, \dots, \mathcal{F}_m$ независимы тогда и только тогда, когда независимы порождающие их разбиения.

Доказательство. Если σ -алгебры $\mathcal{F}_1, \dots, \mathcal{F}_m$ независимы, то порождающие их разбиения тоже независимы, поскольку элементы разбиений принадлежат алгебрам. Пусть теперь независимы разбиения

$$\{A_{11}, A_{12}, \dots\}, \{A_{21}, A_{22}, \dots\}, \dots, \{A_{m1}, A_{m2}, \dots\}$$

порождающие σ -алгебры $\mathcal{F}_1, \dots, \mathcal{F}_m$, и $B_1 \in \mathcal{F}_1, \dots, B_m \in \mathcal{F}_m$. Покажем, что $\mathbf{P}(B_1 B_2 \dots B_m) = \mathbf{P}(B_1) \mathbf{P}(B_2) \dots \mathbf{P}(B_m)$. Действительно, так как σ -алгебра $\mathcal{F}_k$ порождается разбиением $\{A_{k1}, A_{k2}, \dots\}$, то каждому событию $B_k \in \mathcal{F}_k$ соответствует такое множество натуральных чисел $\{i_k(1), i_k(2), \dots\}$, что

$$B_k = \bigsqcup_{v \geq 1} A_{ki_k(v)}, \quad k = 1, \dots, m$$

Тогда

$$B_1 B_2 \dots B_m = \bigcap_{k=1}^m \left(\bigsqcup_{v \geq 1} A_{ki_k(v)} \right) = \bigsqcup_{v_1 \geq 1} \dots \bigsqcup_{v_m \geq 1} A_{1i_1(v_1)} \dots A_{mi_m(v_m)}$$

поскольку пересечения элементов разбиений в правой части попарно не пересекаются: если $(v_1, \dots, v_m) \neq (u_1, \dots, u_m)$, то существует $k \in \{1, \dots, m\}$, при котором $v_k \neq u_k$; тогда $A_{ki_k(v_k)} \cap A_{ki_k(u_k)} = \emptyset$ и

$$(A_{1i_1(v_1)} \dots A_{mi_m(v_m)}) \cap (A_{1i_1(u_1)} \dots A_{mi_m(u_m)}) \subseteq (A_{ki_k(v_k)} \cap A_{ki_k(u_k)}) = \emptyset$$

Так как вероятность объединения несовместных событий равна сумме вероятностей этих событий, а разбиения $\mathcal{F}_1, \dots, \mathcal{F}_m$ независимы, то

$$\begin{aligned} \mathbf{P}\{B_1 B_2 \dots B_m\} &= \sum_{v_1 \geq 1} \dots \sum_{v_m \geq 1} \mathbf{P}\{A_{1i_1(v_1)} \dots A_{mi_m(v_m)}\} = \\ &= \sum_{v_1 \geq 1} \dots \sum_{v_m \geq 1} \mathbf{P}\{A_{1i_1(v_1)}\} \dots \mathbf{P}\{A_{mi_m(v_m)}\} = \\ &= \prod_{k=1}^m \left(\sum_{v \geq 1} \mathbf{P}\{A_{ki_k(v)}\} \right) = \prod_{k=1}^m \mathbf{P}\{B_k\} \end{aligned}$$

Теорема доказана.

3.1.3 Calculus of Events (!?!?)

(очень важный раздел!)

Определения

Случайное событие A , связанное с опытом S , - это такое событие, которое может произойти или не произойти в результате опыта S , причём заранее, до проведения опыта, неизвестно, произойдёт оно или нет. Всюду в дальнейшем при рассмотрении случайных событий мы будем опускать слово «случайное». Достоверным событием, связанным с опытом S , называется такое событие Ω , которое обязательно произойдёт в результате опыта S . Невозможным событием, связанным с опытом S , называется такое событие $\emptyset$, которое обязательно не произойдёт в результате опыта S .

Над событиями A и B , связанными с одним и тем же опытом S , определены следующие операции.

Событие A влечёт за собой событие B (или событие A Вложено в событие B), если каждое появление события A сопровождается появлением события B . Это обозначается как $A \subseteq B$. События A и B называют эквивалентными, если $A \subseteq B$ и $B \subseteq A$. Эквивалентность обозначается как $A = B$. Объединением (или суммой) событий A и B называется событие $A \cup B$ (или $A + B$), которое наступает всегда, когда наступает либо событие A , либо событие B .

Пересечением (или произведением) событий A и B называется событие $A \cap B$ (или AB), которое наступает всегда, когда события A и B наступают одновременно.

Дополнением события B до события A (или разностью событий A и B) называется событие $A \setminus B$, которое наступает всегда, когда наступает событие A , и при этом не наступает событие B . Противоположным событию A называется событие $\bar{A} = \Omega \setminus A$ (читается «не A »), которое наступает всегда, когда событие A не наступает.

События A и B называются несовместными, если $A \cap B = \emptyset$, т. е. если в результате опыта события A и B не могут наступить одновременно.

Говорят, что события $H_1, H_2, \dots, H_n$ образуют полную группу, если они попарно несовместны ($H_i \cap H_j = \emptyset$, $i \neq j$) и их объединение эквивалентно достоверному событию ($H_1 \cup H_2 \cup \dots \cup H_n = \Omega$).

Случайное событие ω , связанное с опытом S , которое невозможно представить как объединение или пересечение более простых событий, связанных с тем же опытом, называется элементарным событием. Очевидно, достоверное событие $\Omega = \{\omega\}$ - это множество всех элементарных событий (поэтому Ω называют ещё пространством элементарных событий), а невозможное событие $\emptyset$ - это пустое множество. Любое событие, связанное с опытом S , можно представить как некоторое подмножество достоверного события Ω , т. е. как множество некоторых элементарных событий.

Для наглядного представления событий, операций над событиями и отношений между ними используются диаграммы Вьенна - Эйлера (рис. 1.1). На этих диаграммах достоверное событие Ω изображается в виде некоторой области на плоскости, элементарные события ω_i - точками внутри области, соответствующей Ω . При этом любому случайному событию A будет соответствовать некоторая геометрическая фигура внутри области, соответствующей Ω (рис. 1.1а). Объединение $A \cup B$ событий A и B состоит из всех элементарных событий, принадлежащих, по крайней мере, одному из событий A и B (рис. 1.1б). Пересечение $A \cap B$ событий A и B состоит из всех элементарных событий, принадлежащих одновременно обоим событиям A и B (рис. 1.1в). Дополнение $A \setminus B$ события B до события A состоит из всех элементарных событий, принадлежащих событию A и при этом не принадлежащих событию B (рис. 1.1г). Событие $\bar{A}$, противоположное событию A , состоит из всех элементарных событий, не принадлежащих событию A (рис. 1.1д). Несовместные события не имеют общих элементарных событий (рис. 1.1е). Полная группа событий представлена на рис. 1.1ж. Событие A влечёт за собой событие B , если все элементарные события, входящие в A , входят и в B (рис. 1.1з). Операции над событиями обладают следующими свойствами:

коммутативность объединения событий: $A \cup B = B \cup A$,

ассоциативность объединения событий:

$$(A \cup B) \cup C = A \cup (B \cup C),$$

$$A \cup A = A, \quad A \cup \emptyset = A, \quad A \cup \Omega = \Omega,$$

$$A \cup \bar{A} = \Omega,$$

коммутативность пересечения событий: $A \cap B = B \cap A$, ассоциативность пересечения событий: $(A \cap B) \cap C = A \cap (B \cap C)$, дистрибутивность пересечения событий относительно объединения:

$$(A \cup B) \cap C = (A \cap C) \cup (B \cap C),$$

$$A \cap A = A,$$

$$A \cap \bar{A} = \emptyset,$$

$$A \cap \emptyset = \emptyset,$$

$$A \cap \Omega = A,$$

$$A \setminus B = A \cap \bar{B},$$

$$\bar{\Omega} = \emptyset$$

$$\bar{\emptyset} = \Omega$$

правила де Моргана: $\overline{A \cup B} = \bar{A} \cap \bar{B}$, $\overline{A \cap B} = \bar{A} \cup \bar{B}$.

пример несовместных случайных событий.

Известно, что $A \subseteq B$. Найти: а) $A \cup B$; б) $A \cap B$.

3.1.4 Characteristic Functions of Sets (!?!?!?!?)

3.2 Other Theorems About Set

3.2.1 The Borel-cantelli Lemma

(зубков)

Пусть на вероятностном пространстве $(\Omega, \mathcal{F}, P)$ задана последовательность событий $\{A_1, A_2, \dots\}$ и

$$A^* = \left\{ \omega \in \Omega : \sum_{n \geq 1} \mathbb{I}(A_n) = \infty \right\}$$

где $\mathbb{I}(A_n)$ - индикатор события A_n .

а) Если $\sum_{n \geq 1} P\{A_n\} < \infty$, то $P\{A^*\} = 0$

б) Если события $A_1, A_2, \dots$ попарно независимы и $\sum_{n \geq 1} P\{A_n\} = \infty$, то $P\{A^*\} = 1$

Доказательство

Пусть $\nu_T = \sum_{1 \leq n \leq T} \mathbb{I}(A_n)$ - число одновременно происходящих событий в совокупности $A_1, \dots, A_T$. Последовательность ν_T монотонно не убывает, поэтому при каждом $\omega \in \Omega$ существует $\lim_{T \rightarrow \infty} \nu_T \stackrel{\text{def}}{=} \nu = \sum_{n \geq 1} \mathbb{I}\{A_n\} \leq \infty$ Тогда

$$M\nu_T = \sum_{k=1}^T P\{A_k\}, \quad M\nu = \sum_{k \geq 1} P\{A_k\} \quad \text{и} \quad A^* = \{\nu = \infty\}$$

а) Если ряд $\sum_{n \geq 1} P\{A_n\}$ сходится, то $M\nu < \infty$ и поэтому $P\{\nu < \infty\} = 1$, т. е. $P\{A^*\} = 0$ б)
Покажем, что $D\nu_T \leq M\nu_T$. По общей формуле для дисперсии суммы случайных величин

$$D\nu_T = D \sum_{n=1}^T \mathbb{I}(A_n) = \sum_{n=1}^T D\mathbb{I}(A_n) + \sum_{\substack{k,n=1 \\ k \neq n}}^T \text{cov}(\mathbb{I}(A_k), \mathbb{I}(A_n))$$

В силу попарной независимости событий A_n их индикаторы тоже попарно независимы, значит, $\text{cov}(\mathbb{I}(A_k), \mathbb{I}(A_n)) = 0$ при $k \neq n$. Кроме того, $D\mathbb{I}(A) = P\{A\}(1 - P\{A\}) \leq P\{A\}$. Следовательно

$$D\nu_T = \sum_{k=1}^T D\mathbb{I}(A_k) = \sum_{k=1}^T P\{A_k\}(1 - P\{A_k\}) \leq M\nu_T$$

Используя неравенство Чебышева и эту оценку, находим, что при любом $a < M\nu_T$

$$\begin{aligned} P\{\nu_T < a\} &\leq P\{\{\nu_T < a\} \cup \{\nu_T > M\nu_T + (M\nu_T - a)\}\} = \\ &= P\{|\nu_T - M\nu_T| > M\nu_T - a\} \leq \frac{D\nu_T}{(M\nu_T - a)^2} \leq \frac{M\nu_T}{(M\nu_T - a)^2} \end{aligned}$$

В случае б) $\nu_T \uparrow \nu$ и $M\nu_T \uparrow \infty$ при $T \uparrow \infty$, значит, при любом $a < \infty$

$$P\{\nu < a\} \leq \lim_{T \rightarrow \infty} P\{\nu_T < a\} = 0 \quad \text{и} \quad P\{\nu \geq a\} = 1 - P\{\nu < a\} = 1$$

Наконец,

$$P\{\nu = \infty\} = P\left\{\bigcap_{a=1}^{\infty} \{\nu \geq a\}\right\} = 1$$

как вероятность пересечения счетной совокупности событий, вероятность каждого из которых равна 1. Лемма доказана.

О важности условия попарной независимости

Второе утверждение леммы Бореля-Кантелли без условия попарной независимости (или какого-то заменяющего его условия) может быть неверным. В качестве примера можно рассмотреть вероятностное пространство $(\Omega, \mathcal{F}, P)$, в котором $\Omega = [0, 1]$, $\mathcal{F}$ — σ -алгебра борелевских множеств и P — мера Лебега, и положить $A_k = (0, a_k]$, $k = 1, 2, \dots$. Если $a_k \rightarrow 0$ при $n \rightarrow \infty$, то для любого $\omega \in [0, 1]$ множество таких k , что $\omega \in A_k$, конечно, поэтому $A^* = \emptyset$ и $P\{A^*\} = 0$ независимо от того, сходится или расходится ряд $\sum_{k \geq 1} a_k$

О случаях применения (!?!?!?)

В доказательствах (???)

(?? пока хз)

Лемма Бореля-Кантелли оказывается полезным инструментом при доказательстве предельных теорем для последовательностей случайных величин.

Из Леммы Б-К следует, что в одномерной цепочке если есть вероятность, что какое-то звено будет изъято, то она не будет одной линией? (???)

(на статфизе говорили это, я хз, как это следует отсюда??)

3.2.2 Kolmogorov 's Inequality

Неравенство Колмогорова. Пусть $\xi_1, \dots, \xi_n$ - независимые случайные величины, которые имеют конечные математические ожидания и дисперсии, и $S_n = \xi_1 + \dots + \xi_n$. Тогда

$$\mathbf{P} \left\{ \max_{1 \leq k \leq n} |S_k - \mathbf{M}S_k| \geq x \right\} \leq \frac{\mathbf{D}S_n}{x^2}$$

Доказательство

Так как $S_k - \mathbf{M}S_k = \sum_{m=1}^k (\xi_m - \mathbf{M}\xi_m)$ и

$$\mathbf{M}(\xi_m - \mathbf{M}\xi_m) = 0, \quad \mathbf{D}(\xi_m - \mathbf{M}\xi_m) = \mathbf{D}\xi_m$$

то без ограничения общности можно считать, что $\mathbf{M}\xi_1 = \dots = \mathbf{M}\xi_n = 0$, и доказывать, что при этом условии

$$\mathbf{P} \left\{ \max_{1 \leq k \leq n} |S_k| \geq x \right\} \leq \frac{\mathbf{D}S_n}{x^2}$$

Введем случайную величину

$$\tau = \begin{cases} k, & \text{если } S_1^2, \dots, S_{k-1}^2 < x^2, S_k^2 \geq x^2 \\ n+1, & \text{если } S_1^2, \dots, S_n^2 < x^2 \end{cases}$$

Так как события $\{\tau = k\}, k = 1, \dots, n+1$, попарно несовместны, то

$$S_n^2 \geq S_n^2 \sum_{k=1}^n \mathbb{I}\{\tau = k\}$$

и в силу монотонности и аддитивности математического ожидания

$$\begin{aligned} \mathbf{D}S_n &= \mathbf{M}S_n^2 \geq \sum_{k=1}^n \mathbf{M}S_n^2 \mathbb{I}\{\tau = k\} = \\ &= \sum_{k=1}^n \mathbf{M}(S_k + S_n - S_k)^2 \mathbb{I}\{\tau = k\} \geq \\ &\geq \sum_{k=1}^n \mathbf{M}S_k^2 \mathbb{I}\{\tau = k\} + 2 \sum_{k=1}^n \mathbf{M}S_k (S_n - S_k) \mathbb{I}\{\tau = k\} \end{aligned}$$

(отброшенное слагаемое $\sum_{k=1}^n \mathbf{M}(S_n - S_k)^2 \mathbb{I}\{\tau = k\}$ неотрицательно). Сделаем два замечания. Во-первых,

$$\mathbf{P}\{S_k^2 \mathbb{I}\{\tau = k\} \geq x^2\} = \mathbf{P}\{S_k^2 \mathbb{I}\{\tau = k\} \geq 0\} = \mathbf{P}\{\tau = k\}$$

т. е. $\mathbf{M}S_k^2 \mathbb{I}\{\tau = k\} \geq x^2 \mathbf{P}\{\tau = k\}, k = 1, \dots, n$. Во-вторых, при любом фиксированном k индикатор $\mathbb{I}\{\tau = k\}$ есть функция от $\xi_1, \dots, \xi_k$, и поэтому разность $S_n - S_k = \xi_{k+1} + \dots + \xi_n$ не зависит от значений $\mathbb{I}\{\tau = k\}$ и S_k . Значит, при любом $k = 1, \dots, n$

$$\begin{aligned} \mathbf{M}S_k (S_n - S_k) \mathbb{I}\{\tau = k\} &= \mathbf{M}S_k \mathbb{I}\{\tau = k\} \cdot \mathbf{M}(S_n - S_k) = \\ &= \mathbf{M}S_k \mathbb{I}\{\tau = k\} \cdot \mathbf{M}(\xi_{k+1} + \dots + \xi_n) = 0 \end{aligned}$$

Объединяя эти замечания с оценкой (53), получаем:

$$\begin{aligned} \mathbf{D}S_n &\geq \sum_{k=1}^n \mathbf{M}S_k^2 \mathbb{I}\{\tau = k\} \geq \sum_{k=1}^n x^2 \mathbf{P}\{\tau = k\} = \\ &= x^2 \mathbf{P}\{\tau \leq n\} = x^2 \mathbf{P} \left\{ \max_{1 \leq k \leq n} |S_k| \geq x \right\} \end{aligned}$$

Деля обе части неравенства на x^2 , завершаем доказательство неравенства Колмогорова.

О применениях неравенства Колмогорова

(?? хз)

3.2.3 The Formalism of Probability Theory (???)

Обсудим вкратце формализм теории вероятностей, приведем часто используемые теоремы.

Вероятностное пространство и вероятность

(тут идея его прописать на достаточном уровне подробно)

Определение пространства элементарных событий и события

Исходным объектом вероятностной модели является пространство элементарных событий. Его принято обозначать прописной греческой буквой Ω , произвольную точку этого пространства называют элементарным событием и обозначают строчной буквой $\omega : \omega \in \Omega$.

алгебра и события

(пока хз)

(тут пара слов про алгебру, начинаем использовать ее понятия и язык, а подробно со всеми свойствами - в следующей главе, там с алгебры и всяких в деталях свойств все вообще и будет начинаться)

вот наши аксиомы

вот наше вероятностное пространство.

Ω Пространство элементарных событий

Основными понятиями теории вероятностей являются вероятностное пространство и событие.

Исходным объектом вероятностной модели является пространство элементарных событий, которое принято обозначать прописной греческой буквой Ω . Любой элемент этого пространства называют элементарным событием и обозначают строчной буквой $\omega : \omega \in \Omega$.

Если считать, что мы строим вероятностную модель конкретного случайного явления, то пространство элементарных событий Ω можно представлять себе как множество всех возможных реализаций ω этого явления.

Пространство элементарных событий для одной и той же задачи можно выбирать разными способами. (?? поч??) Например, в качестве пространства элементарных событий в модели бросания игрального кубика можно выбрать $\Omega = [0, 1)$ и считать, что выпадает k очков, если $\omega \in [\frac{k-1}{6}, \frac{k}{6})$.

Общий алгоритм решения вероятностных задач (??!!!)

Удачный выбор пространства элементарных событий может упростить решение задачи.

Следующий шаг в построении вероятностного пространства состоит в выделении из множества всех подмножеств Ω некоторой совокупности $\mathcal{F}$ подмножеств, которые далее будут называться событиями. Каждому событию сопоставляется число из отрезка $[0, 1]$,

которое называют вероятностью события. Множество событий должно быть замкнуто относительно простейших теоретико-множественных операций: объединения, пересечения и дополнения, а вероятности событий удовлетворять некоторым естественным условиям.

Понятия теории множеств

В теории вероятностей принята следующая множественная терминология.

Суммой множеств событий A и B называется их объединение $A \cup B$.

пересечение $A \cap B$ - произведением, Для сокращения длины формул часто вместо $A \cap B$ пишут просто AB . (!!!)

дополнение $A^c = \bar{A} = \Omega \setminus A$ - противоположным событием.

Если A и B - события, то говорят, что при каждом $\omega \in A \cap B$ события A и B происходят одновременно.

Если $A \cap B = \emptyset$, то события A и B называют несовместными.

Если $A \subseteq B$, то говорят, что событие A влечет событие B .

Наиболее используемые свойства $\cup$ и $\cap$ (???)

(потом оставлю именно такие)

Операции $\cup$ и $\cap$ коммутативны, ассоциативны и транзитивны: для любых подмножеств $A, B, C \subseteq \Omega$

$$\begin{aligned} AB &= BA, & A \cup B &= B \cup A \\ (AB)C &= A(BC), & (A \cup B) \cup C &= A \cup (B \cup C) \\ A(B \cup C) &= AB \cup AC, & A \cup (BC) &= (A \cup B)(A \cup C) \end{aligned}$$

Также: (???)

$$\begin{aligned} \text{если } A \subset B, \text{ то } \bar{B} &\subset \bar{A} \\ \overline{AB} &= \bar{A} \cup \bar{B}, \\ \overline{A \cup B} &= \bar{A} \cap \bar{B}, \\ A \cup B &= A \cup (\bar{A} \cap B), \\ A \setminus B &= A \setminus (A \cap \bar{B}) \end{aligned}$$

$$\overline{\bigcap_{k \geq 1} A_k} = \bigcup_{k \geq 1} \bar{A}_k, \quad \bigcup_{k \geq 1} A_k = \bigcap_{k \geq 1} \bar{A}_k$$

Алгебра множеств

Совокупность $\mathcal{A}$ подмножеств Ω называется алгеброй множеств, если выполнены следующие условия:

- 1) $\emptyset \in \mathcal{A}, \quad \Omega \in \mathcal{A}$
- 2) если $A \in \mathcal{A}$, то и $\Omega \setminus A = A^c \in \mathcal{A}$;
- 3) если $A, B \in \mathcal{A}$, то $A \cup B \in \mathcal{A}$ и $A \cap B \in \mathcal{A}$.

В условия 1) и 3) вторые части являются следствиями первых и условия 2).

Примеры алгебр множеств:

- а) совокупность всех подмножеств конечного множества Ω ,
- б) совокупность всех конечных объединений содержащихся в $(0, 1]$ замкнутых справа полуинтервалов.

В алгебре множеств можно (разными способами) выделять конечную или бесконечную совокупность «элементарных» множеств, из которых операциями

объединения, пересечения и взятия дополнения можно получить все множества, составляющие алгебру.

Говорят, что такая совокупность подмножеств порождает алгебру множеств.

σ -алгебра

Алгебра $\mathcal{F}$ подмножеств множества Ω называется σ -алгеброй, если для любой конечной или счетной совокупности множеств $A_1, A_2, \dots \in \mathcal{F}$

$$\bigcup_{n \geq 1} A_n \in \mathcal{F} \quad \text{и} \quad \bigcap_{n \geq 1} A_n \in \mathcal{F}$$

Множество всех подмножеств Ω , очевидно, является и алгеброй, и σ -алгеброй. Его выбирают в качестве множества событий $\mathcal{F}$, если пространство элементарных событий Ω конечно или счетно.

Пусть $\mathcal{F}_k, k \in K$, - конечная или бесконечная совокупность алгебр (или σ -алгебр) подмножеств множества Ω .

Пересечением алгебр (соответственно, σ -алгебр) $\mathcal{F}_k, k \in K$, называется совокупность всех таких подмножеств $A \subseteq \Omega$, что $A \in \mathcal{F}_k$ при любом $k \in K$.

(??) объединение алгебр может не быть алгеброй, в отличие от пересечений

Задача: проверить, что пересечение алгебр (σ -алгебр) является алгеброй (σ -алгеброй) и что обычное теоретико-множественное объединение алгебр (σ -алгебр) может не быть алгеброй (σ -алгеброй).

(хз, как проверить, потом сделаю это)

Борелевская σ -алгебра

Определение. Пусть M - совокупность подмножеств множества Ω .

Пересечение всех алгебр подмножеств множества Ω , которые содержат M , называется наименьшей алгеброй, порожденной M . Пересечение всех σ -алгебр подмножеств множества Ω , которые содержат M , называется наименьшей σ -алгеброй, или борелевской σ -алгеброй, порожденной M .

Мера и типичные её свойства

(всё, что про неё нужно, в один пар-ф сюда сношу)

Определение. Пусть $\mathcal{F}$ - алгебра или σ -алгебра подмножеств Ω .

Функция $\mu : \mathcal{F} \rightarrow \mathbb{R}$ называется мерой, если она аддитивна, т.е. $\mu(A \cup B) = \mu(A) + \mu(B)$ для любых таких $A, B \in \mathcal{F}$, что $A \cap B = \emptyset$.

Мера μ называется σ -счётно-аддитивной, если $\mu(\bigcup_{k=1}^{\infty} A_k) = \sum_{k=1}^{\infty} \mu(A_k)$ для любой последовательности попарно не пересекающихся множеств $A_1, A_2, \dots \in \mathcal{F}$. Примеры мер на $\mathbb{R}$: а) мера Лебега, б) $\mu(A)$ - количество целых чисел, принадлежащих A .
Определение. Мера μ на σ -алгебре $\mathcal{F}$ подмножеств Ω называется вероятностной, если все ее значения лежат в отрезке $[0, 1]$ и $\mu(\Omega) = 1$.

Теорема Каратеодори

Обычно меры определяют их значениями на совокупности множеств, порождающих σ -алгебру. Теорема Каратеодори указывает условия, при которых можно продолжить счетно-аддитивную меру с алгебры на σ -алгебру.

Согласно теореме Каратеодори, если на алгебре $\mathcal{A}$ подмножеств пространства Ω определена вероятностная мера μ , обладающая свойством счетной аддитивности, то эту меру можно продолжить на наименьшую σ -алгебру $\mathcal{F} = \sigma(\mathcal{A})$, порожденную алгеброй $\mathcal{A}$ с сохранением свойства счетной аддитивности.

(в теории меры доказательство и пояснения, тут потом вставлю про теорвер дополнения)

Свойства вероятности

(вкратце об основных свойствах)

3.2.4 Large Deviations Theory

Суть (??)

$$P(M_N > x) \approx \exp(-NI(x))$$

Теория по вики

Consider a sequence of independent tosses of a fair coin. The possible outcomes could be heads or tails. Let us denote the possible outcome of the i -th trial by X_i , where we encode head as 1 and tail as 0. Now let M_N denote the mean value after N trials, namely

$$M_N = \frac{1}{N} \sum_{i=1}^N X_i$$

Then M_N lies between 0 and 1. From the law of large numbers it follows that as N grows, the distribution of M_N converges to $0.5 = E[X]$ (the expected value of a single coin toss). Moreover, by the central limit theorem, it follows that M_N is approximately normally distributed for large N . The central limit theorem can provide more detailed information about the behavior of M_N than the law of large numbers. For example, we can approximately find a tail probability of M_N , $P(M_N > x)$, that M_N is greater than x , for a fixed value of N . However, the approximation by the central limit theorem may not be accurate if x is far from $E[X_i]$ unless N is sufficiently large. Also, it does not provide information about the convergence of the tail probabilities as $N \rightarrow \infty$. However, the large deviation theory can provide answers for such problems. Let us make this statement more precise. For a given value $0.5 < x < 1$, let us compute the tail probability $P(M_N > x)$. Define

$$I(x) = x \ln x + (1 - x) \ln(1 - x) + \ln 2$$

Note that the function $I(x)$ is a convex, nonnegative function that is zero at $x = \frac{1}{2}$ and increases as x approaches 1. It is the negative of the Bernoulli entropy with $p = \frac{1}{2}$; that it's appropriate for coin tosses follows from the asymptotic equipartition property applied to a Bernoulli trial. Then by Chernoff's inequality, it can be shown that $P(M_N > x) < \exp(-NI(x))$.^[2] This bound is rather sharp, in the sense that $I(x)$ cannot be replaced with a larger number which would yield a strict inequality for all positive N .^[3] (However, the exponential bound can still be reduced

by a subexponential factor on the order of $1/\sqrt{N}$; this follows from the Stirling approximation applied to the binomial coefficient appearing in the Bernoulli distribution.) Hence, we obtain the following result:

$$P(M_N > x) \approx \exp(-NI(x))$$

The probability $P(M_N > x)$ decays exponentially as $N \rightarrow \infty$ at a rate depending on x . This formula approximates any tail probability of the sample mean of i.i.d. variables and gives its convergence as the number of samples increases.

О связи энтропией в статфизе (??)

(остается вопрос, там же иначе все вроде мы получали.)

Возьмем логарифм от $P(M_N > x) \approx \exp(-NI(x))$ и переобозначим вероятность, получим

$$\frac{1}{N} \ln W = -I$$

(??? дальше какое-то рассуждение про это, пока мало что знаю.)

3.3 The Simplest Probabilistic Models

(тут школа и около школьные все задачи и модели, самое простейшее. мб помещу раздел куда-то в другое место)

3.3.1 Properties of the Classical Probabilistic Model

Теория вероятностей работает с классической вероятностной моделью, поэтому познакомимся с ней и посмотрим то, что с помощью нее можно сделать?

(вроде бы она вообще всегда будет работать)

вероятностной моделью работает с вероятностным пространством.

Теорема

(Основная часть аксиоматики Колмогорова).

Вероятностным пространством называется тройка (Ω, F, P) , в котором Ω - пространство элементарных событий, F - σ -алгебра подмножеств Ω , называемых событиями, $P : F \rightarrow [0, 1]$ - функция, называемая вероятностью, удовлетворяющая следующим свойствам:

1. $P(A) \geq 0 \quad \forall A \in F$, неотрицательность
2. $P(\Omega) = 1$, нормированность.
3. $P(\bigcup_{n \geq 1} A_n) = \sum_{n \geq 1} P(A_n)$ - счетная аддитивность.

Таким образом мы строго определили вероятность. Далее посмотрим на ее свойства.

свойства вероятности

- 1) $P(A^c) = 1 - P(A)$, в частности, $P\{\emptyset\} = 0$ и $P(A) \leq 1$ для любого $A \in \mathcal{F}$.

Доказательство

Первое равенство следует из аксиомы 3° , так как $A \cup A^c = \Omega$, $A \cap A^c = \emptyset$. При $A = \Omega$ получаем второе равенство. Наконец, $P(A) \leq 1$, так как должно выполняться условие:

$$P(A^c) = 1 - P(A) \geq 0$$

- 2) Если $A \subseteq B$, то $P\{A\} \leq P\{B\}$ и $P\{B \setminus A\} = P\{B\} - P\{A\}$.

Доказательство

Действительно, так как $B = A \cup (B \setminus A)$ и $A \cap (B \setminus A) = \emptyset$, то по аксиоме 3° примененной к последовательности $A, B \setminus A$,

$$\mathbf{P}\{B\} = \mathbf{P}\{A\} + \mathbf{P}\{B \setminus A\}$$

что эквивалентно доказываемому равенству. Неравенство $\mathbf{P}\{A\} \leq \mathbf{P}\{B\}$ следует из соотношения (1) и аксиомы 1°

3) Для любых событий $A_1, A_2, \dots, A_n \subset \Omega$

$$\mathbf{P}\left\{\bigcup_{k=1}^n A_k\right\} \leq \sum_{k=1}^n \mathbf{P}\{A_k\}$$

Доказательство

Положим $B_1 = A_1, B_k = A_k \setminus (A_1 \cup \dots \cup A_{k-1}), k = 2, \dots, n$. Докажем, что при Любом $m = 1, \dots, n$

$$\bigcup_{k=1}^m A_k = \bigcup_{k=1}^m B_k$$

Из включений $B_k \subset A_k, k = 1, \dots, n$, следует, что $\bigcup_{k=1}^m B_k \subset \bigcup_{k=1}^m A_k$. С другой стороны, для каждого элемента $\omega \in \bigcup_{k=1}^m A_k$ определено значение $k(\omega) =$

$$\min\{k \geq 1 : \omega \in A_k\} \leq m, \text{ и тогда } \omega \in B_{k(\omega)} \in \bigcup_{k=1}^m B_k$$

События $B_1, \dots, B_n$ попарно не пересекаются, так как

$$B_m \cap \left(\bigcup_{k=1}^{m-1} B_k\right) = \left(A_m \setminus \left(\bigcup_{k=1}^{m-1} A_k\right)\right) \cap \left(\bigcup_{k=1}^{m-1} A_k\right) = \emptyset, \quad m = 2, \dots, n$$

Используем теперь аксиому 3° и свойство 2)

$$\mathbf{P}\left(\bigcup_{k=1}^n A_k\right) = \sum_{k=1}^n \mathbf{P}(B_k) \leq \sum_{k=1}^n \mathbf{P}(A_k)$$

4) Для любых событий A и B

$$\mathbf{P}(A \cup B) = \mathbf{P}(A) + \mathbf{P}(B) - \mathbf{P}(A \cap B)$$

Доказательство

Так как $A \cup B = (A \setminus AB) \cup (B \setminus AB) \cup AB$ и события $A \setminus AB, B \setminus AB$ и AB попарно не пересекаются, то по аксиоме 3°

$$\mathbf{P}(A \cup B) = \mathbf{P}(A \setminus AB) + \mathbf{P}(B \setminus AB) + \mathbf{P}(AB)$$

Кроме того, по той же аксиоме 3°

$$\mathbf{P}(A) = \mathbf{P}(A \setminus AB) + \mathbf{P}(AB) \text{ и } \mathbf{P}(B) = \mathbf{P}(B \setminus AB) + \mathbf{P}(AB)$$

Из и (4) следует (2)

5) От объединений можно переходить к пересечениям с помощью теоремы сложения вероятностей.

Теорема о сложении вероятностей

Для любых событий $A_1, A_2, \dots, A_n$

$$\begin{aligned} \mathbf{P} \left\{ \bigcup_{k=1}^n A_k \right\} &= \sum_{m=1}^n (-1)^{m-1} \sum_{1 \leq i_1 < \dots < i_m \leq n} \mathbf{P} \{A_{i_1} A_{i_2} \dots A_{i_m}\} \equiv \\ &\equiv \sum_{i=1}^n \mathbf{P}(A_i) - \sum_{1 \leq i_1 < i_2 \leq n} \mathbf{P}(A_{i_1} A_{i_2}) + \dots + (-1)^{n-1} \mathbf{P}(A_1 \cdot \dots \cdot A_n) \quad (3.1) \end{aligned}$$

Доказательство

При $n = 2$ утверждение совпадает с 4). Поэтому докажем по индукции: допустим, что при некотором $n \geq 2$ утверждение справедливо, и покажем, что тогда оно справедливо и при $n + 1$. Пусть $A_1, \dots, A_{n+1}$ - произвольные события. Мы должны показать, что

$$\mathbf{P} \left\{ \bigcup_{k=1}^{n+1} A_k \right\} = \sum_{m=1}^{n+1} (-1)^{m-1} \sum_{1 \leq i_1 < \dots < i_m \leq n+1} \mathbf{P} \{A_{i_1} A_{i_2} \dots A_{i_m}\}$$

Если $B = A_1 \cup \dots \cup A_n$, то $\bigcup_{k=1}^{n+1} A_k = B \cup A_{n+1}$ и в силу 4) и предположения индукции

$$\mathbf{P} \left\{ \bigcup_{k=1}^{n+1} A_k \right\} = \mathbf{P}\{B\} + \mathbf{P}\{A_{n+1}\} - \mathbf{P}\{B \cap A_{n+1}\}$$

где $\mathbf{P}\{B\}$ есть правая часть (5). Правая часть (6) отличается от правой части (5) только слагаемыми, содержащими A_{n+1} , и нам осталось доказать, что

$$\mathbf{P}\{A_{n+1}\} - \mathbf{P}\{B \cap A_{n+1}\} = \sum_{m=1}^{n+1} (-1)^{m-1} \sum_{1 \leq i_1 < \dots < i_m = n+1} \mathbf{P}\{A_{i_1} A_{i_2} \dots A_{i_m}\}$$

Слагаемое в правой части, соответствующее $m = 1$, есть $\mathbf{P}\{A_{n+1}\}$. Остальные слагаемые соответствуют $\mathbf{P}\{B \cap A_{n+1}\}$:

$$\begin{aligned} \mathbf{P}\{B \cap A_{n+1}\} &= \mathbf{P}\left\{\left(\bigcup_{k=1}^n A_k\right) \cap A_{n+1}\right\} = \mathbf{P}\left\{\bigcup_{k=1}^n (A_k \cap A_{n+1})\right\} = \\ &= \sum_{m=1}^n (-1)^{m-1} \sum_{1 \leq i_1 < \dots < i_m \leq n} \mathbf{P}\left\{\bigcap_{j=1}^m (A_{i_j} \cap A_{n+1})\right\} = \\ &= - \sum_{m=1}^n (-1)^{(m+1)-1} \sum_{1 \leq i_1 < \dots < i_m < i_{m+1} = n+1} \mathbf{P}\{A_{i_1} \dots A_{i_m} A_{i_{m+1}}\} \end{aligned}$$

О случаях применения (???)**Теорема о непрерывности вероятностной меры**

Для любой неубывающей

$$\mathbf{P} \left\{ \bigcup_{n=1}^{\infty} A_n \right\} = \lim_{n \rightarrow \infty} \mathbf{P} \{A_n\}$$

$$\mathbf{P} \left\{ \bigcap_{n=1}^{\infty} B_n \right\} = \lim_{n \rightarrow \infty} \mathbf{P} \{B_n\}$$

В частности, если $B_1 \supseteq B_2 \supseteq \dots \cup \bigcap_{n=1}^{\infty} B_n = \emptyset$, то

$$\lim_{n \rightarrow \infty} \mathbf{P} \{B_n\} = 0$$

Доказательство

Доказательство. 1. Аналогично доказательству свойства 3) по неубывающей последовательности событий $\{A_n\}$ построим последовательность несовместных событий $B_1 = A_1, B_2 = A_2 \setminus A_1, B_3 = A_3 \setminus A_2$ и т.д. Так же, как в доказательстве свойства 3), $\cup_{n=1}^{\infty} A_n = \sqcup_{n=1}^{\infty} B_n$, кроме того, $\mathbf{P}\{B_n\} = \mathbf{P}\{A_n\} - \mathbf{P}\{A_{n-1}\}$ при $n \geq 2$ по свойству 2). Тогда в силу аксиомы счетной аддитивности

$$\mathbf{P}\{\cup_{n=1}^{\infty} A_n\} = \mathbf{P}\{\sqcup_{n=1}^{\infty} B_n\} = \sum_{n=1}^{\infty} \mathbf{P}\{B_n\}$$

Сумма ряда с неотрицательными членами равна пределу его частичных сумм, но так как $A_N = \sqcup_{n=1}^N B_n$ при любом $N \geq 1$, то

$$\mathbf{P}\{B_1\} + \mathbf{P}\{B_2\} + \dots + \mathbf{P}\{B_N\} = \mathbf{P}\{A_N\}$$

()

$$\mathbf{P}\{\cup_{n=1}^{\infty} A_n\} = \sum_{n=1}^{\infty} \mathbf{P}\{B_n\} = \lim_{N \rightarrow \infty} \mathbf{P}\{A_N\}$$

что и требовалось доказать.

$A_n = \overline{B_n} = \Omega \setminus B_n, n \geq 1$, то $A_1 \subseteq A_2 \subseteq \dots$ и по доказанному в п.1 и по свой- crBy 1)

$$\mathbf{P}\{\cup_{n=1}^{\infty} A_n\} = \lim_{N \rightarrow \infty} \mathbf{P}\{A_N\} = \lim_{N \rightarrow \infty} (1 - \mathbf{P}\{B_N\}) = 1 - \lim_{N \rightarrow \infty} \mathbf{P}\{B_N\}$$

С другой стороны,

$$\begin{aligned} \mathbf{P}\{\cup_{n=1}^{\infty} A_n\} &= \mathbf{P}\{\omega : \exists k : \omega \in A_k\} = \mathbf{P}\{\Omega \setminus \{\omega : \omega \in A_k \forall k\}\} = \\ &= \mathbf{P}\{\overline{\cap_{n=1}^{\infty} \overline{A_n}}\} = \mathbf{P}\{\overline{\cap_{n=1}^{\infty} B_n}\} = 1 - \mathbf{P}\{\cap_{n=1}^{\infty} B_n\} \end{aligned}$$

Из этих двух равенств вытекает первое утверждение п.2. Второе утверждение следует из первого и из свойства 1). Теорема доказана.

свойства вероятностного пространства

(6 параграф зубков)

О случаях применения (???)**примеры**

поясним примерами классическую вероятностную модель.

посмотрим теперь, используя вероятностное пространство, на известные нам уже модели.

Example 1. Конечные вероятностные пространства.

Пусть в тройке $(\Omega, \mathcal{F}, \mathbf{P})$ $\Omega = \{\omega\}$ - пространство элементарных событий - конечное множество, состоящее из $|\Omega|$ элементов, σ -алгебра (в данном случае - алгебра) $\mathcal{F}$ - совокупность всех подмножеств Ω . Вероятность $\mathbf{P}$ можно задать, определив для каждого $\omega \in \Omega$ неотрицательное число $\mathbf{P}(\omega)$ так, чтобы

$$\sum_{\omega \in \Omega} \mathbf{P}(\omega) = 1$$

Для произвольного события $A \subseteq \Omega$ положим

$$P(A) = \sum_{\omega \in A} P(\omega)$$

Выполнение аксиом Колмогорова 1° – 3° в этом случае очевидно.

Частным случаем является так называемое классическое

определение

вероятности, когда $P(\omega) = \frac{1}{|\Omega|}$ для всех $\omega \in \Omega$. В этом случае

$$P(A) = \frac{|A|}{|\Omega|} \text{ для любого } A \subseteq \Omega$$

Конструкция классической вероятности

Приведем типичную конструкцию вероятности, которая чаще всего имеется в виду.
(?? самое нужное тут пропишу как раз, в самом начале)

Конструкция вероятности на пальцах

(тут крайне не строгое понимание всего, которое часто дается в вводных курсах)

Определение классической вероятности

Пусть при некотором испытании множество всевозможных событий Ω состоит из конечного числа элементарных событий $\omega_1, \dots, \omega_n$.

Говоря на бытовом языке, классическая вероятность $P(A)$ есть число, сопоставляемое каждому множеству событий:

$$P(A) = \frac{|A|}{|\Omega|}$$

От этого числа требуется только нормированность и чтобы вероятность нулевого была равна 0.

3.3.2 Combinatorics for Samples

Именно этим параграфом и решаются все простые задачи классической вероятностной модели.

Основное правило комбинаторики

Все формулы простейшей комбинаторики выводятся из основного правило комбинаторики.

(????)

Lemma 2. (Основное правило комбинаторики) Пусть имеется m наборов элементов, причем i -тый набор состоит из k_i элементов. Общее число N способов, которыми можно произвести выбор по одному элементу из каждой группы, определяется равенством

$$N = k_1 \cdot k_2 \cdot \dots \cdot k_m.$$

(вроде достаточно очевидно)

(потом посмотрю подробнее, мб докажу)

Комбинаторные указания (!!!!!)

(!!!! вот это супер важно для решений задач!!!)

Частным случаем классической вероятностной схемы является урновая схема: в урне содержится $(K + L)$ шаров, среди которых K белых и L чёрных; из урны наугад без возвращения извлекаются $(k + l)$ шаров, тогда вероятность $\mathbf{P}_{K,L}(k, l)$ того, что в выборке содержится ровно k белых шаров и l чёрных, вычисляется по формуле гипергеометрической вероятности:

$$\mathbf{P}_{K,L}(k, l) = \frac{C_K^k C_L^l}{C_{K+L}^{k+l}}$$

Размещения по ящикам

(такое тоже есть, подробно - в записи по комбинаторике)

Размещение незанумерованных шаров по занумерованным ящикам

Размещение занумерованных шаров по незанумерованным ящикам

Размещение занумерованных шаров по занумерованным ящикам

О других размещениях (???)

(пока не знаю, но могут быть ведь мб?)

(??? если у нас больше ящиков, чем коробок?)

(?? если у нас больше коробок, чем ящиков??)

(пока просто не думал, комбинаторику нужно подтянуть.)

Выборки из ящиков

Приведем типичные задачи про выборки, решаемые простой комбинаторикой.

(по идее многие простые задачи должны решаться этой главой. мб потом сделаю unindent, если будет прямо много на это задач)

Обзор количества комбинаций при разных выборках

Многие (?? какие??) комбинаторные вычисления укладываются в следующую схему. Допустим, что мы имеем n занумерованных шаров: $a_1, \dots, a_n$. Осуществляется выборка объема k из этой совокупности шаров. Нас интересует количество способов, которыми можно осуществить эту выборку.

При этом возникает четыре различные ситуации: выборка может быть упорядоченной или неупорядоченной и с возвращением или без возвращения.

(тут итоги приведу)

Число возможных комбинаций для данного типа выборки указано в таблице.

Выборки	с возвращением	без возвращения
упорядоченные	n^k	$A_n^k = \frac{n!}{(n-k)!}$
неупорядоченные	$C_{n+k-1}^k = \frac{(n+k-1)!}{k!(n-1)!}$	$C_n^k = \frac{n!}{k!(n-k)!}$

Выборка с возвращением одинаковых шаров разных цветов

(если без нумерации, то что?)

О применениях выборки с возвращением одинаковых шаров разных цветов (??)

(пока хз)

Выборка с возвращением занумерованных шаров разных цветов (??)

Пусть есть урна, содержащая N шаров, занумерованных числами $1, 2, \dots, N$, причем шары с номерами от 1 до M белые, а остальные $N - M$ шаров черные. Из урны последовательно извлекают, узнают номер, возвращают обратно t шаров. Вероятность, что среди t шаров, появилось s белых есть:

$$\mathbf{P}(A_s) = C_t^s \left(\frac{M}{N}\right)^s \left(1 - \frac{M}{N}\right)^{t-s}$$

(?? пока не вникал как следует в док-во, не актуально, но оно ниже)

Действительно.

Возможные исходы t извлечений шаров (элементарные события) можно отождествить с наборами $(k_1, k_2, \dots, k_t)$ чисел из множества $\{1, \dots, N\}$.

Поэтому в качестве пространства элементарных событий в данном случае можно выбрать множество $\Omega = \{(k_1, \dots, k_t) : k_1, \dots, k_t \in \{1, \dots, N\}\} = 1, \dots, N^t$

Количество наборов $(k_1, k_2, \dots, k_t)$, образующих такое пространство элементарных событий Ω , равно N^t

При классическом определении вероятности каждому набору $(k_1, \dots, k_t) \in \Omega$ сопоставляется вероятность $\frac{1}{N^t}$

(тут комбинаторика должна работать!!!! пока я криво очень уж ее написал, так что я хз, почему так задача решается)

Количество наборов чисел $(k_1, k_2, \dots, k_t) \in \{1, \dots, N\}^t$, содержащих ровно s чисел из множества $\{1, \dots, M\}$ и $t - s$ чисел из $\{M + 1, \dots, N\}$, равно $C_t^s M^s (N - M)^{t-s}$, поэтому

$$\mathbf{P}(A_s) = C_t^s \left(\frac{M}{N}\right)^s \left(1 - \frac{M}{N}\right)^{t-s}$$

О применениях выборки с возвращением занумерованных шаров разных цветов (??)

Можно посмотреть на применения такого примера.

Например, случай при $N = 2$, $M = 1$ можно интерпретировать как t -кратное бросание идеальной симметричной монеты, или как двоичная последовательность.

при этом исходы бросаний называют «гербом» и «решеткой». Другая интерпретация этого же случая состоит в том, что белый шар отождествляется с 0, а черный — с 1;

тогда выборка соответствует равновероятной случайной двоичной последовательности и событие $B_s =$ равновероятная случайная двоичная последовательность длины t содержит ровно s единиц имеет вероятность

$$\mathbf{P}\{B_s\} = \frac{1}{2^t} C_t^s$$

Выборка без возвращения

Пусть есть урна, содержащая N шаров, занумерованных числами $1, 2, \dots, N$, причем шары с номерами от 1 до M белые, а остальные $N - M$ шаров черные. Пусть из урны

извлекают $t \leq N$ шаров без возвращения. Вероятность A_s^* того, что при выборе t шаров без возвращения появилось s белых шаров есть:

$$\mathbf{P}(A_s^*) = C_t^s \frac{M^{[s]}(N-M)^{[t-s]}}{N^{[t]}}$$

(???)

Действительно.

Теперь в качестве элементарных событий следует рассматривать наборы попарно различных чисел $(k_1, k_2, \dots, k_t) \in \{1, \dots, N\}^t$.

В качестве пространства элементарных событий можно выбрать

$$\Omega = \{(k_1, \dots, k_t) : k_1, \dots, k_t \in \{1, \dots, N\}, k_i \neq k_j, 1 \leq i < j \leq t\}.$$

Число элементов этого пространства элементарных событий Ω вычисляется по другой формуле:

$$|\Omega| = N(N-1)\dots(N-t+1) = N^{[t]}$$

Вероятность каждого элементарного события в схеме выбора t шаров без возвращения равна $\frac{1}{N^{[t]}}$.

Поэтому число элементарных событий A_s^* равно

$$N(A_s^*) = C_t^s M^{[s]}(N-M)^{[t-s]}$$

и находим вероятность

$$\mathbf{P}(A_s^*) = C_t^s \frac{M^{[s]}(N-M)^{[t-s]}}{N^{[t]}}$$

О применениях выборки...

Другие выборки (?)

(если еще что-то будет, то сюда буду вписывать. пока хз, какие еще есть выборки)

3.3.3 geometric Probability

Обсудим геометрическую вероятность сполна.

(думаю, эта вся теория в этом разделе и поместится.)

Конструкция геометрической вероятности

Определение геометрической вероятности

Если пространство элементарных событий Ω представляет собой область в $\mathbb{R}^n, n = 1, 2, \dots$, и все элементарные события считаются равновероятными, вероятность события $A \subset \Omega$ вычисляется по формуле так называемой геометрической вероятности:

$$\mathbf{P}(A) = \frac{\mu(A)}{\mu(\Omega)}$$

где под $\mu(A)$ понимается мера Лебега множества A , хотя для простых множеств достаточно и меры Жордана.

(??? конкретики мало, потом добавлю её, досмотрю про это лекции)

О случаях применимости геометрической вероятности (!!)

(потом четко напишу!!!)

Об ошибках в понимании геометрической вероятности

Ошибка думать, что тут задачи только про падение точки в область, ведь геометрическая вероятность встречается в самых абсолютно разных случаях.

(конкретно - пока я не усвоил)

Почему в геометрической вероятности пространство Евклидово?

Если плоскость x, t , то казалось бы, имеем пространство Минковского????

А если те же задачи, но со скоростями близкими к скорости света решить????

вот тут я просто не знаю.

остался вопрос.

О способах решения задач на геометрическую вероятность (?!?)

Вроде просто рисуем картинку, считаем нужную меру множества. Чаще всего это просто площадь.

По идее иногда это мог бы быть и объем большей фигуры.

О применениях геометрической вероятности в негеометрических задачах (!!?)

(пока мало опыта, напишу потом.)

О задачах про встречи через геометрическую вероятность

(типичные задачи, укажу на их решение.)

3.4 Conditional Probability.**3.4.1 Conditional Probability and Fundamental Theorems****Конструкция условной вероятности в двух словах**

(всё тут, не вижу смысла делить на пар-фы)

Рассмотрим ограничение вероятностного пространства на множество $B : (\Omega_B, \mathcal{F}_B, \mathbf{P}_B^*)$, где $\Omega_B = B$, $\mathcal{F}_B = \{C \in \mathcal{F} : C \subseteq B\}$ и $\mathbf{P}_B^*\{C\} = \mathbf{P}\{C\}$ для всех множеств $C \subset B, C \in \mathcal{F}$.

Чтобы $\mathbf{P}_B^*$ - было вероятностной мерой, нужно ее нормировать числом $\mathbf{P}\{B\}$, иначе было бы $\mathbf{P}_B^*\{\Omega_B\} = \mathbf{P}\{B\} < 1$.

Definition 3.15. Условная вероятность события $A \in \mathcal{F}$ в вероятностном пространстве $(\Omega, \mathcal{F}, \mathbf{P})$ при условии наступления события B называется величина

$$\mathbf{P}(A|B) = \frac{\mathbf{P}(AB)}{\mathbf{P}(B)}$$

Свойства условной вероятности следующие.

Условные вероятности независимых событий A, B совпадают с безусловными. Если $\mathbf{P}\{A\} > 0, \mathbf{P}\{B\} > 0$ и $\mathbf{P}\{AB\} = \mathbf{P}\{A\}\mathbf{P}\{B\}$, то

$$\mathbf{P}\{A|B\} = \mathbf{P}\{A|\bar{B}\} = \mathbf{P}\{A\}, \mathbf{P}\{\bar{A}|B\} = \mathbf{P}\{\bar{A}|\bar{B}\} = \mathbf{P}\{\bar{A}\}$$

(?? какие еще???)

О случаях применения условной вероятности (?!)

Часто нужно перейти от всего вероятностного пространства к его подмножеству нужно воспользоваться условной вероятностью. Например, это нужно, если мы ищем среди события B исходы благоприятствующие событию A .

Теорема умножения

Из определения следует формула

$$P(AB) = P(B)P(A|B)$$

Обобщим эту формулу для n событий.

Теорема умножения

Для любых событий A таких что $P(A_1 A_2 \dots A_{n-1}) \neq 0$, справедливо равенство

$$P\{A_1 A_2 \dots A_n\} = P\{A_1\} P\{A_2|A_1\} P\{A_3|A_1 A_2\} \dots P\{A_n|A_1 A_2 \dots A_{n-1}\}$$

Доказательство

(пока скопировано, не важно оно)

Так как $A_1 \dots A_{n-1} \subseteq A_1 \dots A_k$ при любом $k \leq n-1$, то из условия $P\{A_1 A_2 \dots A_{n-1}\} > 0$ следует, что $P\{A_1 A_2 \dots A_k\} > 0$ при любом $k = 1, 2, \dots, n-1$ т.е. что все условные вероятности в правой части (19) определены. При $n = 2$ формула (19) совпадает с (18). Далее воспользуемся методом математической индукции. Допустим, что при некотором $n \geq 2$ формула (19) верна для любых n событий, и, используя (18), докажем, что тогда формула (19) верна и для любых $n+1$ событий:

О случаях применения теоремы умножения

(пока хз)

О неверных пониманиях условной вероятности

(тоже часто её можно криво понимать, опишу тут, как)

3.4.2 Total Probability Formula

В случае несколько исходов, после каждого из которых вероятность дальнейших событий разная, для подсчета вероятности нужно суммировать по всем исходам, в чем нам поможет формула полной вероятности.

(???)

Конструкция формулы полной вероятности в двух словах (?)

Изначальный набор событий $A_1, A_2, \dots$ при условии которых наступает событие B представляет собой разбиение множества событий, то есть разбивает пространство всех событий на непересекающиеся подмножества, полно покрывающих исходное.

Теорема

(Формула полной вероятности) Если $A_1, A_2, \dots$ — не более чем счетное разбиение пространства событий Ω такое, что $P\{A_k\} \neq 0$ ($k = 1, 2, \dots$), то для любого события B справедливо равенство

$$P\{B\} = \sum_{k \geq 1} P\{A_k\} P\{B|A_k\}$$

(?? кст, а для более чем счетного что?)

Доказательство

Для доказательства нужно просто переписать B через пересечение с каждым A_i , а потом просто воспользоваться формулой условной вероятности.

То есть пусть $A_1, A_2, \dots$ - разбиение Ω , то

$$B = B \cap \Omega = B \cap \left(\bigcup_{k \geq 1} A_k \right) = \bigcup_{k \geq 1} (BA_k) \text{ и } BA_k \cap BA_j = \emptyset \text{ при } k \neq j$$

Пользуясь аддитивностью вероятности и теоремой умножения, находим:

$$P\{B\} = \sum_{k=1}^n P\{BA_k\} = \sum_{k=1}^n P\{A_k\} P\{B|A_k\}$$

Теорема доказана.

О случаях применениях полной вероятности

(вроде понимаю, потренироваться нужно тут ещё)

В случае несколько исходов, при каждом из которых вероятность последующего события разная, для подсчета вероятности нужно суммировать по всем исходам, в чем нам поможет формула полной вероятности.

Вероятность хотя бы одного исхода (!?)

Вычисление вероятности появления хотя бы одного из совместных событий $A_1, A_2, \dots, A_n$ можно вычислять как разность между единицей и вероятностью произведения противоположных событий $\bar{A}_1, \bar{A}_2, \dots, \bar{A}_n$:

$$P\left(\sum_{i=1}^n A_i\right) = 1 - P(\bar{A}_1 \bar{A}_2 \dots \bar{A}_n).$$

В частности, если все n событий имеют одинаковую вероятность, равную p , то вероятность появления хотя бы одного из этих событий

$$P(A) = 1 - (1 - p)^n.$$

3.4.3 Bayes Formula.

Если известно $(B|A_k)$, а хотим мы найти $(A_k|B)$, то в этом нам поможет формула Байеса.

(тут посмотрю популярную лекцию, давно в палнах это!!!)

Формула Байеса**Теорема**

(формула Байеса)

Если $A_1, A_2, \dots$ — полная система несовместных событий и событие B таково, что $P(B) \neq 0$, то

$$P\{A_k|B\} = \frac{P\{A_k\} P\{B|A_k\}}{\sum_{k \geq 1} P\{A_k\} P\{B|A_k\}}$$

Доказательство

Достаточно воспользоваться определением условной вероятности, теоремой умножения и формулой полной вероятности:

$$P\{A_k|B\} = \frac{P\{A_k B\}}{P\{B\}} = \frac{P\{A_k\} P\{B|A_k\}}{\sum_{k \geq 1} P\{A_k\} P\{B|A_k\}}$$

Теорема доказана.

О понимании формулы Байеса (??)

(??? что там за вероятность прошлого от будущего???? так и не понял пока)

О случаях применений формулы Байеса (!)

(Зачем она??? хз. самое важное!!)

О Байесовском методе

(он вообще связан с ф-ой Байеса?)

Примеры условной вероятности

Приведем типичные иллюстративные примеры условной вероятности.
(буду решать задачи - приведу)

3.4.4 Independence of Events.

Обсудим подробно независимые события

Конструкция независимых событий (???)**Новая выгрузка нормальной теории**

Как уже было отмечено, говорить о вероятности наступления какого-либо события как о мере возможности наступления этого события можно лишь при выполнении определённого комплекса условий опыта. Так, если к комплексу условий, при которых изучалась вероятность наступления события A , добавить условие наступления события B , получим другое значение вероятности $P\{A|B\}$ - вероятность наступления события A при условии, что событие B произошло (или ² условная вероятность события A при условии B), которая равна по определению

$$P\{A|B\} = \frac{P\{A \cap B\}}{P\{B\}}.$$

Иногда вместо обозначения $\mathbf{P}\{A|B\}$ используют обозначение $\mathbf{P}_B\{A\} = \mathbf{P}\{A|B\}$. Вероятность $\mathbf{P}\{A\}$, в отличие от условной вероятности $\mathbf{P}\{A|B\}$, называется безусловной. Из определения условной вероятности (2.1) следует формула умножения Вероятностей

$$\mathbf{P}\{A \cap B\} = \mathbf{P}\{B\}\mathbf{P}\{A|B\} = \mathbf{P}\{A\}\mathbf{P}\{B|A\},$$

остающаяся справедливой и в случае, когда $\mathbf{P}\{A\} = 0$ или $\mathbf{P}\{B\} = 0$. Говорят, что события A и B являются независимыми, если

$$\mathbf{P}\{A \cap B\} = \mathbf{P}\{A\}\mathbf{P}\{B\}.$$

Формулу (2.3) для независимых событий A и B называют теоремой умножения вероятностей. Очевидно, при $\mathbf{P}\{B\} > 0$ теорема умножения вероятностей (2.3) означает, что условная вероятность события A при условии B совпадает с безусловной вероятностью события A :

$$\mathbf{P}\{A|B\} = \mathbf{P}\{A\}.$$

Формулу умножения вероятностей легко обобщить на случай произвольного конечного числа событий:

$$\mathbf{P}\{A_1 \cap A_2 \cap \dots \cap A_n\} = \mathbf{P}\{A_1\}\mathbf{P}\{A_2|A_1\}\mathbf{P}\{A_3|A_1 \cap A_2\} \dots \mathbf{P}\{A_n|A_1 \cap A_2 \cap \dots \cap A_{n-1}\}.$$

События $A_1, A_2, \dots, A_n$ называются независимыми в совокупности, если для любого их подмножества $A_{i_1}, A_{i_2}, \dots, A_{i_k}$ вероятность одновременного наступления событий $A_{i_1}, A_{i_2}, \dots, A_{i_k}$ равна произведению безусловных вероятностей этих событий:

$$\mathbf{P}\{A_{i_1} \cap A_{i_2} \cap \dots \cap A_{i_k}\} = \mathbf{P}\{A_{i_1}\}\mathbf{P}\{A_{i_2}\} \dots \mathbf{P}\{A_{i_k}\} \quad (k = 1, 2, \dots, n).$$

Если условие (2.6) выполняется только при $k = 2$, то события $A_1, A_2, \dots, A_n$ называются попарно независимыми.² Следует обратить внимание на следующие факты: - из условия несовместности не следует условие независимости; - из условия независимости не следует условие несовместности; - из условия независимости в совокупности следует попарная независимость, но из попарной независимости не следует независимость в совокупности.

Если события $A_1, A_2, \dots, A_n$ независимы в совокупности, то из обобщённой теоремы сложения вероятностей следует, что

$$\begin{aligned} \mathbf{P}\left\{\bigcup_{i=1}^n A_i\right\} &= \mathbf{P}\{A_1 \cup A_2 \cup \dots \cup A_n\} = \\ &= 1 - (1 - \mathbf{P}\{A_1\})(1 - \mathbf{P}\{A_2\}) \dots (1 - \mathbf{P}\{A_n\}) = 1 - \prod_{i=1}^n (1 - \mathbf{P}\{A_i\}). \end{aligned}$$

Если события $H_1, H_2, \dots, H_n$ образуют полную группу, то для вычисления вероятности произвольного события A можно использовать формулу полной Вероятности:

$$\mathbf{P}\{A\} = \mathbf{P}\{A|H_1\}\mathbf{P}\{H_1\} + \mathbf{P}\{A|H_2\}\mathbf{P}\{H_2\} + \dots + \mathbf{P}\{A|H_n\}\mathbf{P}\{H_n\} = \sum_{i=1}^n \mathbf{P}\{A|H_i\}\mathbf{P}\{H_i\},$$

в соответствии с которой вероятность наступления события A может быть представлена как сумма произведений условных вероятностей события A при условии наступления событий H_i на безусловные вероятности этих событий H_i . Поскольку среди событий $H_1, H_2, \dots, H_n$, образующих полную группу, в результате опыта должно наступить одно и только одно, эти события H_i называют гипотезами ($i = 1, 2, \dots, n$).

Формула полной вероятности (2.8) остаётся справедливой и в случае, если условие, состоящее в том, что события $H_1, H_2, \dots, H_n$ образуют полную группу, заменить более слабым: гипотезы $H_1, H_2, \dots, H_n$ попарно несовместны ($H_i \cap H_j = \emptyset$ при $i \neq j$), а их объединение содержит событие A ($A \subseteq H_1 \cup H_2 \cup \dots \cup H_n$). Из формулы полной вероятности следует формула Байеса:

$$\mathbf{P}\{H_k|A\} = \frac{\mathbf{P}\{A|H_k\} \mathbf{P}\{H_k\}}{\mathbf{P}\{A\}} = \frac{\mathbf{P}\{A|H_k\} \mathbf{P}\{H_k\}}{\sum_{i=1}^n \mathbf{P}\{A|H_i\} \mathbf{P}\{H_i\}} (k = 1, 2, \dots, n).$$

Вероятности $\mathbf{P}\{H_i\}$ гипотез H_i называют априорными вероятностями (вероятностями гипотез H_i до проведения опыта) в отличие от апостериорных вероятностей $\mathbf{P}\{H_i|A\}$ (вероятностей гипотез H_i , уточнённых в результате опыта, исходом которого стало событие A).

Конструкция независимых событий

С информационной точки зрения под независимостью события B от события A естественно понимать то, что знание о наступлении события A не влияет на вероятность наступления события B .

(то есть есть причинная независимость, да?)

Definition 3.16. Пусть $(\Omega, \mathcal{F}, \mathbf{P})$ - вероятностное пространство. События $A, B \in \mathcal{F}$ называются независимыми, если

$$\mathbf{P}(AB) = \mathbf{P}(A)\mathbf{P}(B)$$

Интуитивное представление о независимых событиях

Это два круга, которые пересекаются так, что их мера пересечения равна произведению мер этих кругов.

Симметричность независимости

Независимость событий A и B можно было бы записать с помощью условной вероятности, то есть, например, равенством $P(B|A) = P(B)$, но в то же время можно и равенством $P(A|B) = P(A)$. Покажем, что оба эти равенства эквивалентны в случае, если $P(A) \neq 0$ и $P(B) \neq 0$, то из независимости B от A следует независимость A от B , то есть независимость симметрична.

$$\mathbf{P}(A|B) = \frac{\mathbf{P}(AB)}{\mathbf{P}(B)} = \frac{\mathbf{P}(A)\mathbf{P}(B|A)}{\mathbf{P}(B)} = \mathbf{P}(A)$$

(потом посмотрю. чего не взять бы это за опр-е??)

Независимость в совокупности (??)

(тут детали есть, см. примеры, потом дорабатываюсь)

Definition 3.17. События $A_1, \dots, A_n$ независимы в совокупности, если для любых $k = 2, \dots, n$ и $1 \leq i_1 < i_2 < \dots < i_k \leq n$ выполняются соотношения

$$\mathbf{P}(A_{i_1} \cdot \dots \cdot A_{i_k}) = \mathbf{P}(A_{i_1}) \cdot \dots \cdot \mathbf{P}(A_{i_k})$$

Независимые наборы событий

Можно также расширить определение:

Definition 3.18. *Наборы событий $\mathcal{F}_1, \dots, \mathcal{F}_n$ называются независимыми, если для любого набора $A_1 \in \mathcal{F}_1, \dots, A_n \in \mathcal{F}_n$ события $A_1, \dots, A_n$ независимы в совокупности.*

(?? и зачем оно??)

Свойства независимых событий (??)

(?? откуда вообще это я взял? Горяйнова потом посмотрю!!)

Перечислим простейшие свойства независимых событий.

Lemma 3. *Если события A и B независимы, то независимы также пары событий A и $\bar{B}$, $\bar{A}$ и B , $\bar{A}$ и $\bar{B}$:*

$$\mathbf{P}(A\bar{B}) = \mathbf{P}(A)\mathbf{P}(\bar{B}), \quad \mathbf{P}(\bar{A}B) = \mathbf{P}(\bar{A})\mathbf{P}(B)$$

Доказательство

Утверждение достаточно доказать для пары A и $\bar{B}$. Используя равенство $\mathbf{P}(AB) = \mathbf{P}(A)\mathbf{P}(B)$, получаем

$$\mathbf{P}(A\bar{B}) = \mathbf{P}(A \setminus AB) = \mathbf{P}(A) - \mathbf{P}(AB) = \mathbf{P}(A)(1 - \mathbf{P}(B)) = \mathbf{P}(A)\mathbf{P}(\bar{B})$$

Если $A \in F$ – событие, то события A и $\emptyset$, а также A и Ω независимы

$$\mathbf{P}(A\emptyset) = 0 = \mathbf{P}(A) \cdot 0 = \mathbf{P}(A)\mathbf{P}(\emptyset)$$

$$\mathbf{P}(A\Omega) = \mathbf{P}(A) \cdot 1 = \mathbf{P}(A)\mathbf{P}(\Omega)$$

$\bar{A}$ и $\bar{B}$:

$$\mathbf{P}(A\bar{B}) = \mathbf{P}(A \setminus AB) = \mathbf{P}(A) - \mathbf{P}(AB) = \mathbf{P}(A)(1 - \mathbf{P}(B)) = \mathbf{P}(A)\mathbf{P}(\bar{B})$$

аналогично, $\mathbf{P}(\bar{A}B) = \mathbf{P}(\bar{A})\mathbf{P}(B)$; наконец,

$$\mathbf{P}(\bar{A}\bar{B}) = \mathbf{P}(\bar{A} \setminus \bar{A}B) = \mathbf{P}(\bar{A}) - \mathbf{P}(\bar{A}B) = \mathbf{P}(\bar{A})(1 - \mathbf{P}(B)) = \mathbf{P}(\bar{A})\mathbf{P}(\bar{B})$$

Другое про независимые события**Отличие независимых событий от несовместных событий**

- Для несовместных $\mathbf{P}(A_1 A_2) = \mathbf{P}(\emptyset) = 0$.

- Для независимых $\mathbf{P}(A_1 \wedge A_2) = \mathbf{P}(A_1) * \mathbf{P}(A_2)$

Несовместные события могут быть независимыми, если $|\mathbf{P}(A_1) * \mathbf{P}(A_2)| = 0$, то есть только когда одно из событий – невозможное, т.е. вероятность его наступления равна нулю.

title

О неверных пониманиях независимых событий (??)

(тут соберу их, вроде такое часто может появиться. пока хз, вроде всё и понятно...)

Примеры независимых событий

Пример попарно независимых, но не независимых в совокупности событий (??!!!)

Пусть $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, где все элементарные события ω_i равновероятны. Определим $A = \{\omega_1, \omega_2\}$, $B = \{\omega_1, \omega_3\}$, $C = \{\omega_1, \omega_4\}$. Эти события попарно независимы, но не выполняется равенство

$$P(ABC) = P(A)P(B)P(C)$$

Пример независимых в совокупности, но не попарно независимых событий (??!!!)

Подбрасываются две игральные кости. Рассмотрим три события:

$A = \{\text{на первой кости выпала "единица", "двойка" или "пятерка"}\}$,

$B = \{\text{на первой кости выпала "четверка", "пятерка" или "шестерка"}\}$,

$C = \{\text{сумма выпавших очков на двух игровых костях равна девяти}\}$.

В этом случае выполняется равенство

$$P(ABC) = P(A)P(B)P(C)$$

но нет даже попарной независимости.

Пример независимых событий на плоскости

Пусть пространство элементарных событий $\Omega = \{(x, y) : x, y \in [0, 1]\}$ - единичный квадрат, $\mathcal{F}$ - σ -алгебра измеримых по Лебегу подмножеств Ω , P - мера Лебега на Ω (обобщающая геометрическое понятие площади). Если

$$A = \{(x, y) \in \Omega : a_1 \leq x \leq a_2\}, \quad B = \{(x, y) \in \Omega : b_1 \leq y \leq b_2\}$$

То $AB = \{(x, y) \in \Omega : a_1 \leq x \leq a_2, b_1 \leq y \leq b_2\}$, и

$$P\{A\} = \frac{a_2 - a_1}{|\Omega|} = a_2 - a_1, \quad P\{B\} = \frac{b_2 - b_1}{|\Omega|} = b_2 - b_1$$

$$P\{AB\} = \frac{(a_2 - a_1)(b_2 - b_1)}{|\Omega|} = (a_2 - a_1)(b_2 - b_1) = P\{A\}P\{B\}$$

т.е. по определению события A и B независимы.

(?? потом досмотрю)

3.5 Independent Test Sequences

Приведем теорию и примеры типичных последовательностей испытаний.

3.5.1 Bernoulli Scheme

Будем рассматривать схемы, в которых в каждом испытании вероятность наступления определенного события равны p , обозначим вероятность его не наступления $q = 1 - p$.

Схема Бернулли в двух словах

Схема Бернулли это задача о независимых испытаниях, в каждом из которых вероятность наступления события A постоянна и равна p . Ищется вероятность того, что в среди n испытаний событие A наступит ровно m раз?

Будем смотреть на происходящее, как на последовательность белых и черных кружков, где белый - означает, что событие A не наступило, черный - что наступило. Тогда вероятность одной такой последовательности равна $p^m q^{n-m}$. Всего последовательностей существует C_n^m штук. Поэтому вероятность иметь среди n испытаний ровно m событий A дается так называемой формулой Бернулли или биномиальным законом:

$$P_n(m) = C_n^m p^m q^{n-m}$$

Название такое из-за того, что $P_n(m)$ равна коэффициенту при x^m в разложении бинома $(q + px)^n$.

Хотя бы какое-то число раз нужное событие наступит (?)

Заметим, что среди n испытаний событие A точно наступит какое-то число раз от нуля до n . Поэтому сумма вероятностей исходов с разными количествами появления A должна равняться вероятности достоверного события:

$$\sum_{m=0}^n P_n(m) = 1$$

Это следует из формулы бинома Ньютона: $\sum_{m=0}^n P_n(m) = (p + q) = 1$

Вероятность $P_n(m)$ с увеличением m растет, потом убывает. в случае, если $np - q$ целое, максимальное значение она имеет для двух m : $m_1 = np - q, m_2 = np + q$. В случае, если не целое, то максимальное при наименьшем целом, близким к m_1 .

О применениях схемы Бернулли (!?!?)

(потом пропишу, где нужна, очень много задач на нее.)

Полиномиальная схема

В случае k событий $A_1, A_2, A_3, \dots, A_k$ с вероятностью каждого p_i , в схеме Бернулли вероятность, что A_1 встретится m_1 раз, что A_2 встретится m_2 раз, ..., A_k — m_k равна.

$$P_n(m_1, m_2, \dots, m_k) = \frac{n!}{m_1! m_2! \dots m_k!} p_1^{m_1} p_2^{m_2} \dots p_k^{m_k}$$

Аналогично, ее называют полиномиальным распределением

примеры последовательностей независимых испытаний

Посмотрим на интересные (и только интересные!) примеры.

3.5.2 Typical Asymptotics of a Bernoulli Scheme

(тут только самые полезные и не сложные теоремы!!! основное - в главе дальше. короче говоря, сюда только то, что уже использовал и буду переносить!!!)

Есть множество различных частных случаев схем Бернулли, которые встречаются в жизни. Например, случай Пуассона,

Теорема Пуассона

Простейший предел дается формулой Пуассона.

Если $n \rightarrow \infty$ и $p \rightarrow 0$ так, что $np \rightarrow \lambda, 0 < \lambda < \infty$, то при всех $k = 0, 1, 2, \dots$ выполняется:

$$\mathbf{P}(B_n(k)) = C_n^k p^k q^{n-k} \rightarrow \frac{\lambda^k}{k!} e^{-\lambda}$$

Доказательство

Преобразовав формулу Бернулли

$$\begin{aligned} \frac{n!}{k!(n-k)!} p^k q^{n-k} &= \frac{1}{k!} \frac{(n-1) \cdot \dots \cdot (n-k+1)}{n^{k-1}} n^k p^k (1-p)^{n-k} \\ &= \frac{(np)^k}{k!} \left(1 - \frac{1}{n}\right) \cdot \dots \cdot \left(1 - \frac{k-1}{n}\right) (1-p)^{-k} (1-p)^{\frac{1}{p}np}, \end{aligned}$$

видим, что в пределе у нас стоит экспонента:

$$\mathbf{P}(B_n(k)) = \lim_{p \rightarrow 0} (1-p)^{1/p} = e^{-1}$$

Заметим, что в условиях теоремы Пуассона имеет место оценка

$$\left| P(B_n(k)) - \frac{\lambda^k}{k!} e^{-\lambda} \right| \leq \frac{\lambda^2}{n}, \quad \lambda = np$$

т. е. результат применим, когда np^2 мало.

О случаях применений теоремы Пуассона

(пока хз)

3.5.3 Geometric Distribution.

Обозначим через k_1 номер того испытания в процессе Бернулли, когда впервые выпадает единица. Каковы статистические свойства случайной величины k_1 ? В силу статистической независимости испытаний, функция распределения вероятностей $P(k_1)$ равна произведению вероятностей получить нули в каждой из первых $k_1 - 1$ попытках и вероятности выпадения единицы k_1 -й попытке, то есть

$$P(k_1) = p(1-p)^{k_1-1}$$

Данное распределение называется геометрическим. Математическое ожидание числа попыток, необходимых, чтобы встретить первую единицу в последовательности независимых испытаний Бернулли, равно

$$\langle k_1 \rangle = \sum_{k_1=1}^{\infty} k_1 p (1-p)^{k_1-1} = p \frac{\partial}{\partial q} \sum_{k_1=1}^{\infty} q^{k_1} = p \frac{\partial}{\partial q} \frac{q}{1-q} = \frac{1}{p}$$

Дисперсия номера попытки, в которой впервые выпадает единица равна

$$\begin{aligned} \sigma_{k_1}^2 &= \sum_{k_1=1}^{\infty} k_1^2 p (1-p)^{k_1-1} - \frac{1}{p^2} = p \frac{\partial}{\partial q} \sum_{k_1=1}^{\infty} k_1 q^{k_1} - \frac{1}{p^2} = \\ &= p \frac{\partial}{\partial q} \left[q \frac{\partial}{\partial q} \sum_{k_1=1}^{\infty} q^{k_1} \right] - \frac{1}{p^2} = p \frac{\partial}{\partial q} \left[q \frac{\partial}{\partial q} \frac{q}{1-q} \right] - \frac{1}{p^2} = \frac{1-p}{p^2} \end{aligned}$$

Упражнение: Симметричную монетку подбрасывают снова и снова до тех пор пока она не выпадет два раза подряд одной и той же стороной (неважно какой именно). Вычислите математическое ожидание необходимого числа бросаний k .

3.5.4 Poisson Distribution.

Теория

Найдем функцию распределения числа событий n , произошедших за время наблюдения t . Для этого разделим ось времени на малые интервалы длины $\Delta t \ll \lambda^{-1}, t$. Тогда число пуассоновских событий равно числу единиц в серии из $N = t/\Delta t$ независимых испытаний Бернулли, где вероятность получить единицу в очередном исходе составляет $p = \lambda\Delta t$, причем $\Delta t \rightarrow +0$. Воспользовавшись результатом предыдущего раздела, можно записать

$$\begin{aligned} p(n) &= \lim_{\Delta t \rightarrow +0} B\left(n, \frac{t}{\Delta t}, \lambda\Delta t\right) = \lim_{\Delta t \rightarrow +0} \frac{\left(\frac{t}{\Delta t}\right)!}{n! \left(\frac{t}{\Delta t} - n\right)!} (\lambda\Delta t)^n (1 - \lambda\Delta t)^{\frac{t}{\Delta t} - n} = \\ &= \lim_{\Delta t \rightarrow +0} \frac{1}{n!} \left(\frac{t}{\Delta t}\right)^n (\lambda\Delta t)^n (1 - \lambda\Delta t)^{\frac{\lambda t}{\Delta t}} = \frac{(\lambda t)^n}{n!} e^{-\lambda t} \end{aligned}$$

Данное распределение называется распределением Пуассона, а сам факт стремления биномиального распределения к распределению Пуассона в пределе $p \rightarrow +0, N \rightarrow +\infty, pN = \text{const}$ известен как предельная теорема Пуассона, или закон редких событий, или закон малых чисел. Можно сказать, что пуассоновский процесс представляет собой континуальный предел процесса Бернулли. Среднее число событий, произошедших за время t равно

$$\mu_n = \sum_{n=0}^{+\infty} n \frac{(\lambda t)^n}{n!} e^{-\lambda t} = e^{-\lambda t} \sum_{m=0}^{+\infty} \frac{(\lambda t)^{m+1}}{m!} = \lambda t$$

где нами была выполнена замена индекса суммирования $m = n - 1$. Аналогично находим дисперсию числа событий, случившихся в промежутке времени длительностью t

$$\begin{aligned} \sigma_n^2 &= \sum_{n=0}^{+\infty} n^2 \frac{(\lambda t)^n}{n!} e^{-\lambda t} - (\lambda t)^2 = e^{-\lambda t} \sum_{m=0}^{+\infty} (m+1) \frac{(\lambda t)^{m+1}}{m!} - (\lambda t)^2 = \\ &= e^{-\lambda t} \lambda t \sum_{m=1}^{+\infty} \frac{(\lambda t)^m}{(m-1)!} + e^{-\lambda t} \lambda t \sum_{m=0}^{+\infty} \frac{(\lambda t)^m}{m!} - (\lambda t)^2 = \lambda t \end{aligned}$$

Суперпозиция потоков Пуассона

Важным свойством пуассоновского потока является его стабильность относительно сложения. Оказывается, суперпозиция двух независимых пуассоновских потоков с рейтами λ_1 и λ_2 дает пуассоновский поток с рейтом $\lambda_1 + \lambda_2$.

Для доказательства сложим количества событий, n_1 и n_2 , произошедших за один и тот же интервал времени t в двух пуассоновских потоках с рейтами λ_1 и λ_2 , соответственно. Характеристическая функция вероятностного распределения случайной величины $n = n_1 + n_2$ равна

$$\Phi(k) \equiv \langle e^{ik(n_1+n_2)} \rangle = \langle e^{ikn_1} \rangle \langle e^{ikn_2} \rangle = \varphi_1(k) \varphi_2(k),$$

где $\phi_1(k)$ и $\phi_2(k)$ - характеристические функции распределений Пуассона случайных

величин n_1 и n_2 , то есть

$$\begin{aligned}\varphi_1(k) &\equiv \langle e^{ikn_1} \rangle = \exp(-\lambda_1 t) \sum_{n_1=0}^{+\infty} e^{ikn_1} \frac{(\lambda_1 t)^{n_1}}{n_1!} = \\ &= \exp(\lambda_1 t (e^{ik} - 1)), \\ \varphi_2(k) &\equiv \langle e^{ikn_2} \rangle = \exp(\lambda_2 t (e^{ik} - 1)).\end{aligned}$$

Таким образом, находим

$$\Phi(k) = \varphi_1(k)\varphi_2(k) = \exp((\lambda_1 + \lambda_2)t(e^{ik} - 1)).$$

откуда следует что случайная величина n имеет пуассоновское распределение с рейтом $\lambda_1 + \lambda_2$.

Деление потока Пуассона

Аналогично, пуассоновский поток с рейтом λ можно разделить на два независимых пуассоновских подпотока с рейтами λ_1 и $\lambda_2 = \lambda - \lambda_1$. Деление исходного процесса производится посредством подбрасывания монетки с вероятностями орла и решки равными λ_1/λ и $1 - \lambda_1/\lambda$, случайный исход которого определяет, к какому из двух производных процессов будет принадлежать очередное пуассоновское событие. Проще всего указанное свойство можно доказать, показав их справедливость для процесса Бернулли с малой вероятностью единицы, и перейдя затем к континуальному пределу.

Упражнение: Посетители прибывают в торговый комплекс в случайные моменты времени и независимо друг от друга со средней частотой λ . В каждом отдельном случае вероятность того, что очередной посетитель будет женского пола равна q . Найдите математическое ожидание и дисперсию числа женщин, прибывших в торговый комплекс за промежуток времени t .

Неоднородный поток Пуассона

Неоднородным потоком Пуассона называется поток событий, в котором вероятность $\lambda(t)dt$ наблюдать событие в промежутке времени $[t, t + dt]$ как и в однородном случае не зависит предшествующей истории наблюдений, но зависит от времени t . Такой процесс можно рассматривать как предельный случай последовательности независимых испытаний Бернулли, в которой вероятность выпадения единицы в очередной попытке зависит от номера этой попытки. К примеру, события распада радиоактивных атомов в образце, состоящим в начальный момент времени из большого их числа $N_0 \gg 1$, представляют собой неоднородный пуассоновский процесс с рейтом $\lambda(t) = N_0 \lambda e^{-\lambda t}$.

Распределение числа событий

Пусть n - это число событий, произошедших за время t , отсчитываемое от произвольно выбранного момента начала наблюдений. Найдём функцию распределения вероятностей $p(n; t)$. В случае $n = 0$ получаем

$$P(0; t) = \Pr[qT_1 > t] = \int_t^{+\infty} dT_1 \frac{T_1}{\langle T_1 \rangle} P(T_1) \int_{t/T}^1 dq = \int_t^{+\infty} dT_1 \frac{T_1}{\langle T_1 \rangle} P(T_1) \left(1 - \frac{t}{T}\right),$$

где q - это отношение, в котором точка начала наблюдений делит межсобытийный интервал, на который она попала. Если $n > 0$, то искомое распределение можно

записать как

$$P(n; t) = \Pr [(qT_1 + T_2 + \dots + T_n < t) \cap (T_{n+1} > t - qT_1 - T_2 - \dots - T_n)] = \quad (22)$$

$$= \int_0^t d\tau \rho(\tau) \int_{t-\tau}^{+\infty} dT_{n+1} P(T_{n+1}),$$

где $\rho(\tau)$ - плотность вероятности случайной величины $qT_1 + T_2 + \dots + T_n$.

Решение в Лаплас-представлении Из соотношения (23) с учетом теоремы о свертке следует простая связь между Лаплас-образами функциями $p(n; t)$, $\rho(\tau)$ и $P(T)$, а именно

$$\tilde{p}(n; s) = \frac{1 - \tilde{P}(s)}{s} \tilde{\rho}(s),$$

где

$$\tilde{\rho}(s) = \langle e^{-s\tau} \rangle = \langle e^{-s(qT_1 + T_2 + \dots + T_n)} \rangle = \left\langle e^{-sqT_1} \prod_{i=2}^n e^{-sT_i} \right\rangle = \langle e^{-sqT_1} \rangle \prod_{i=2}^n \langle e^{-sT_i} \rangle \quad (25)$$

$$= \tilde{P}^{n-1}(s) \int_0^{+\infty} dT_1 \frac{T_1}{\langle T_1 \rangle} P(T_1) \int_0^1 dq e^{-sqT_1} = \tilde{P}^{n-1}(s) \int_0^{+\infty} dT_1 P(T_1) \frac{1 - e^{-sT_1}}{\langle T_1 \rangle s} = \quad (26)$$

$$= \frac{\tilde{P}^{n-1}(s)(1 - \tilde{P}(s))}{\langle T \rangle s}.$$

Подставляя этот результат в Ур. (24), находим

$$\tilde{p}(n; s) = \frac{(1 - \tilde{P}(s))^2 \tilde{P}^{n-1}(s)}{\langle T \rangle s^2}$$

Упражнение: Проверьте, что в случае экспоненциального распределения интервала времени между последовательными событиями обратное преобразование Лапласа от выражения (28) дает распределение Пуассона.

Теорема о свертке (напоминание) Пусть

$$\varphi(t) = \int_0^t d\tau f(\tau) F(t - \tau)$$

Умнодим обе части этого равенства на e^{-st} и проинтегрируем по t от 0 до $+\infty$

$$\int_0^{+\infty} dt e^{-st} \varphi(t) = \int_0^{+\infty} dt e^{-st} \int_0^t d\tau f(\tau) F(t - \tau)$$

$$\tilde{\varphi}(s) = \int_0^{+\infty} dt \int_0^t d\tau e^{-st} f(\tau) F(t - \tau).$$

Замена переменных $x = t - \tau$, $y = \tau$ позволяет записать

$$\tilde{\varphi}(s) = \int_0^{+\infty} dx \int_0^{+\infty} dy e^{-s(x+y)} f(y) F(x),$$

$$\tilde{\varphi}(s) = \int_0^{+\infty} dx e^{-sx} F(x) \int_0^{+\infty} dy e^{-sy} f(y),$$

$$\tilde{\varphi}(s) = \tilde{F}(s) \tilde{f}(s)$$

3.5.5 Exponential

Теория по Белану

Рассмотрим теперь статистику интервала времени T между парой последовательных пуассоновских событий. Повторяя рассуждения аналогичные упомянутым выше, можно сказать, что T равно произведению Δt на разность номеров соседних единиц в серии из $N = T/\Delta t$ независимых испытаний Бернулли, где вероятность получить единицу в очередном исходе составляет $p = \lambda\Delta t$, причем $\Delta t \rightarrow +0$. Используя геометрическое распределение (4), запишем

$$P(T) = \lim_{\Delta t \rightarrow +0} \frac{1}{\Delta t} \lambda \Delta t (1 - \lambda \Delta t)^{\frac{T}{\Delta t} - 1} = \lambda e^{-\lambda T}$$

Получили экспоненциальное распределение. Вычислим математическое ожидание и дисперсию случайной величины T

$$\mu_T = \int_0^{+\infty} dT \lambda e^{-\lambda T} T = \lambda^{-1}$$

3.5.6 Moivre-laplace Theorems

Локальная теорема Муавра

Теорема

(Локальная теорема Муавра)

В схеме Бернулли, при вероятности испытания A $P(A) = p$, вероятность $P_n(m)$ событий A при $n \rightarrow \infty$ равномерно стремится к:

$$\sqrt{npq} P_n(m) : \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2} \rightarrow 1.$$

где $x \equiv x_{mn} \equiv \frac{m-np}{\sqrt{npq}}$ находится в конечном интервале, $q = (1 - p)$.

Доказательство

(потом пропишу, пока не актуально, гнеденко стр 79)

Доказывается это с помощью формулы Стирлинга:

$$s! = \sqrt{2\pi s} s^s e^{-s} e^{\theta_s}$$

$$|\theta_s| \leq \frac{1}{12s}$$

$$m = np + x\sqrt{npq}$$

$$P_n(m) = \frac{n!}{m!(n-m)!} p^m q^{n-m} = \frac{1}{\sqrt{2\pi}} \sqrt{\frac{n}{m(n-m)}} e^{-\theta} \left(\frac{n}{m} p\right)^m \left(\frac{n}{n-m} q\right)^{n-m}$$

$$\theta = \theta_n - \theta_m - \theta_{n-m} < \frac{1}{12} \left(\frac{1}{n} + \frac{1}{m} + \frac{1}{n-m} \right)$$

$$\ln A_n = \ln \left(\frac{n}{m} p \right)^m \left(\frac{n}{n-m} q \right)^{n-m} = -(np + x\sqrt{npq}) \ln \left(1 + x\sqrt{\frac{q}{np}} \right) - (nq - x\sqrt{npq}) \ln \left(1 - x\sqrt{\frac{p}{nq}} \right) \quad (3.2)$$

$$\begin{aligned} \ln \left(1 + x\sqrt{\frac{q}{np}} \right) &= x\sqrt{\frac{q}{np}} - \frac{1}{2} \frac{qx^2}{np} + O\left(\frac{1}{n^{3/2}}\right) \\ \ln \left(1 - x\sqrt{\frac{p}{nq}} \right) &= -x\sqrt{\frac{p}{nq}} - \frac{1}{2} \frac{px^2}{nq} + O\left(\frac{1}{n^{3/2}}\right) \end{aligned}$$

$$\ln A_n = -\frac{x^2}{2} + O\left(\frac{1}{\sqrt{n}}\right)$$

$$A_n : e^{-\frac{x^2}{2}} \rightarrow 1 \quad (n \rightarrow \infty)$$

$$\sqrt{npq} \cdot \sqrt{\frac{n}{m(n-m)}} \rightarrow 1 \text{ равномерно}$$

(хз)

О случаях применимости локальной теоремы Муавра (??)

(не знаю, интересно)

Пример применения локальной теоремы Муавра: распределение Бернулли с $n = 10000, m = 40, p = 0,005$

пусть

$$n = 10000, m = 40, p = 0,005$$

По теореме Муавра имеем:

$$P_n(m) \sim \frac{1}{\sqrt{2\pi npq}} e^{-\frac{1}{2} \left(\frac{m-np}{\sqrt{npq}} \right)^2}$$

примерно $\sqrt{npq} = \sqrt{10000 \cdot 0,005 \cdot 0,995} = \sqrt{49,75} \sim 7,05$; $\frac{m-np}{\sqrt{npq}} \sim -1,42$

Осталось вычислить функцию

$$\varphi(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$$

Значение ее ищется в таблице.

В итоге

$$P_n(m) \sim \frac{0,1456}{7,05} = 0,0206$$

(потом еще потренируюсь с таблицами)

Теорема Муавра-Лапласа в двух словах**Теорема**

(Теорема Муавра-Лапласа)

вероятность этого события равна p , то равномерно относительно a и b ($-\infty \leq a \leq b \leq +\infty$)

$$P \left\{ a \leq \frac{\mu - np}{\sqrt{npq}} < b \right\} \rightarrow \frac{1}{\sqrt{2\pi}} \int_a^b e^{-\frac{z^2}{2}} dz$$

Доказательство теоремы Муавра-Лапласа

(скорее всего это в части ниже будет, раз оно долгое. если короткое - вставлю в пар-ф выше)

Доказательство

(см гнеденко)

и тут его три страницы, пока что не буду писать его

О случаях применимости теоремы Муавра-Лапласа (??)

(не знаю, интересно)

Обобщение теоремы Муавра на случай нескольких событий

Таким же способом можно доказать обобщение этой теоремы на случай нескольких испытаний $A_1, A_2, \dots, A_k$.

Теорема

(потом запишу, пока что она не нужна)

Вероятность того, что каждое событий встретится m_1 раз удовлетворяет

$$\sqrt{n^{k-1}} P_n(m_1, m_2, \dots, m_k) : \frac{e^{-\frac{1}{2} \sum_{i=1}^k q_i x_i^2}}{(2\pi)^{\frac{k-1}{2}} \sqrt{p_1 p_2 \dots p_k}} \rightarrow 1$$

О случая применений локальной теоремы Муавра

(хз)

О случая применимости теоремы Муавра-Лапласа

(с опытом впишу)

Пример применения теоремы Муавра-Лапласа

(говорят, применений на самом деле много)

Оценим вероятность неравенства $\left| \frac{\mu}{n} - p \right| < \epsilon$:

$$\mathbf{P} \left\{ \left| \frac{\mu}{n} - p \right| < \epsilon \right\} = \mathbf{P} \left\{ -\epsilon \sqrt{\frac{n}{pq}} < \frac{\mu - np}{\sqrt{npq}} < \epsilon \sqrt{\frac{n}{pq}} \right\}$$

и в силу теоремы Муавра-Лапласа

$$\lim_{n \rightarrow \infty} \mathbf{P} \left\{ \left| \frac{\mu}{n} - p \right| < \epsilon \right\} = \frac{1}{\sqrt{2\pi}} \int e^{-\frac{z^2}{2}} dz = 1$$

Что собственно и очевидно.

Это называется законом больших чисел.

(где-то мы это использовать начнем позже...)

Чему равна вероятность того, что частота наступления события A , у которого вероятность в каждом испытании равна p , за n испытаний отклонится от p не более чем на α ?

$$\begin{aligned} P \left\{ \left| \frac{\mu}{n} - p \right| \leq \alpha \right\} &= P \left\{ -\alpha \sqrt{\frac{n}{pq}} \leq \frac{\mu - np}{\sqrt{npq}} \leq \alpha \sqrt{\frac{n}{pq}} \right\} \sim \\ &\sim \frac{1}{\sqrt{2\pi}} \int_{-\alpha \sqrt{\frac{n}{pq}}}^{\alpha \sqrt{\frac{n}{pq}}} e^{-\frac{x^2}{2}} dx = \frac{2}{\sqrt{2\pi}} \int_0^{\alpha \sqrt{\frac{n}{pq}}} e^{-\frac{x^2}{2}} dx \end{aligned}$$

(численно можно досчитать)

еще пример у Гнеденко

(????)

3.6 Random Variables in a Nutshell

Обсудим самую используемые свойства и конструкции случайной величины.

3.6.1 Typical Random Variable Problems

Поиск распределения случайной величины из двух, распределение которых известно (!?!?!?!?!?)

(максимально типичные задачи, пропишу, как решаются они!!!!)

3.6.2 Random Variable Design

Обсудим конструкцию случайной величины.

Определение случайной величины

Основные формулы случайной величины

(тут всё, что и использую чаще всего, скоро оттренирую, впишу. мб метод через индикаторы четко вставляю)

Вкратце о случайной величине и её использовании

Случайной величиной называется функция, у которой прообраз любого борелевского множества принадлежит сигма алгебре исходного множества.

что еще?

(тут дальше пробегаюсь обзором, что вводятся дальше такие-то величины, так-то она используется. пересказ всего в двух словах. полезная такая заготовка, именно это я и думаю, когда представляю весь этот раздел.)

О том, почему такая замысловатая конструкция случайной величины (???)

(а я хз на самом деле, так не понял, что-то тут на самом деле ведь заумно. чего не сказать проще??? очень актуальный вопрос)

Определение и свойства случайной величины

Definition 3.19. Случайная величина $\xi = \xi(\omega)$ в $(\Omega, \mathcal{F}, \mathbf{P})$ на произвольном множестве называется функция $\xi = \xi(\omega)$, принимающая значения в множестве S с σ -алгеброй подмножеств Σ , называется любая $(\mathcal{F}, \Sigma)$ -измеримая функция $\xi : \Omega \rightarrow S$, т.е. такая функция, что для любого подмножества $A \subseteq S, A \in \Sigma$, выполняется условие

$$\xi^{-1}(A) = \{\omega \in \Omega : \xi(\omega) \in A\} \in \mathcal{F}$$

(???? я так и не понял, почему такое определение???)

Если $\xi_1(\omega), \dots, \xi_n(\omega)$ - случайные величины, заданные на одном и том же конечном или счетном вероятностном пространстве $(\Omega, \mathcal{F}, \mathbf{P})$, а $g(X_1, \dots, X_n)$ - функция, заданная на прямом произведении множеств значений этих случайных величин, то $\eta(\omega) = g(\xi_1(\omega), \dots, \xi_n(\omega))$ - тоже случайная величина, определенная на $(\Omega, \mathcal{F}, \mathbf{P})$.

Если $(\Omega, \mathcal{F}, \mathbf{P})$ - произвольное вероятностное пространство, то для того, чтобы суперпозиция $g(\xi_1, \dots, \xi_n)$ была случайной величиной, нужно налагать на функцию $g : \Omega \rightarrow S$ те или иные условия, обеспечивающие измеримость суперпозиции $g(\xi_1, \dots, \xi_n)$.

(?? потом подумаю про то, что выше, но скорее это не нужно)

О других определениях случайной величины

(укажем тут, что квантовая случайная величина есть и что-нибудь еще, что дополняет теорию выше)

Индикаторы событий**Пример: индикатор событий**

(мб про индикатор отдельный раздел сделаю потом)
Простейшая случайная величина - индикатор событий.
Индикатор для события $A \in \mathcal{F}$ определяется как:

$$I_A = I(A) = I_A(\omega) = \begin{cases} 1 & \text{при } \omega \in A \\ 0 & \text{при } \omega \notin A \end{cases}$$

(что про него важное известно?)

О применении индикаторов (!!!)

(почему ими часто пользуются в задачах??? я хз, пока не понимаю это)

Свойства индикаторов

Свойства индикаторов следующие:

- 1) $I(\emptyset) \equiv 0, I(\Omega) \equiv 1, I(A^c) = 1 - I(A)$
- 2) Для любых событий $A_1, \dots, A_n$

$$I(A_1 \cap \dots \cap A_n) = I(A_1) \dots I(A_n)$$

3) Если события $A_1, \dots, A_n$ попарно несовместны, то

$$I(A_1 \cup \dots \cup A_n) = \sum_{k=1}^n I(A_k)$$

4) Если $A \cap B \neq \emptyset$, то $I(A \cup B) \neq I(A) + I(B)$.

5) В общем случае

$$I\left(\bigcup_{k=1}^n A_k\right) = \sum_{s=1}^n (-1)^{s-1} \sum_{1 \leq k_1 < \dots < k_s \leq n} I(A_{k_1} \dots A_{k_s})$$

Доказательство

(пока не смотрел пруф, потом впишу)

Докажем свойства выше. (? их же доказываем?)

Их можно вывести с помощью уже перечисленных свойств и равенства

$$\bigcup_{k=1}^n A_k = \bigcap_{k=1}^n \overline{A_k}$$

Действительно:

$$\begin{aligned} I\left(\bigcup_{k=1}^n A_k\right) &= 1 - I\left(\bigcap_{k=1}^n \overline{A_k}\right) = 1 - I\left(\bigcap_{k=1}^n \overline{A_k}\right) = \\ &= 1 - \prod_{k=1}^n I(\overline{A_k}) = 1 - \prod_{k=1}^n (1 - I(A_k)) = \\ &= \sum_{s=1}^n (-1)^{s-1} \sum_{1 \leq i_1 < i_2 < \dots < i_s \leq n} I(A_{i_1}) \dots I(A_{i_s}) \end{aligned}$$

Заменяя произведения индикаторов индикаторами пересечений событий, получаем:

$$\begin{aligned} I\left(\bigcup_{k=1}^n A_k\right) &= \sum_{k=1}^n I(A_k) - \sum_{1 \leq k_1 < k_2 \leq n} I(A_{k_1} A_{k_2}) + \\ &+ \sum_{1 \leq k_1 < k_2 < k_3 \leq n} I(A_{k_1} A_{k_2} A_{k_3}) - \dots + (-1)^{n-1} I(A_1 \dots A_n) = \\ &= \sum_{s=1}^n (-1)^{s-1} \sum_{1 \leq k_1 < \dots < k_s \leq n} I(A_{k_1} \dots A_{k_s}) \end{aligned}$$

Если I_A - индикатор события A , то

$$\mathbf{M}I_A = \mathbf{P}\{A\}$$

Действительно:

$$\mathbf{M}I_A = \sum_{\omega \in \Omega} I_A(\omega) \mathbf{P}(\omega) = \sum_{\omega \in A} \mathbf{P}(\omega) = \mathbf{P}\{A\}$$

Порождение алгебры случайной величиной

(зубков про начало про распределения)

Пусть на вероятностном пространстве $(\Omega, \mathcal{F}, \mathbf{P})$ задана случайная величина ξ и множество ее значений $S = \{x_1, x_2, \dots\}$ конечно или счетно. Тогда события $A_j = \{\omega : \xi(\omega) = x_j\}$, $j = 1, 2, \dots$, образуют конечное или счетное разбиение Ω . Это разбиение порождает алгебру событий (подмножеств Ω)

$$A(\xi) = \{A = \{\omega \in \Omega : \xi(\omega) \in B\} : B \subseteq S\}$$

поскольку для любого $B \subseteq S$

$$\begin{aligned} \{\omega \in \Omega : \xi(\omega) \in B\} &= \left\{ \omega \in \Omega : \xi(\omega) \in \bigcup_{j: x_j \in B} \{x_j\} \right\} = \\ &= \bigcup_{j: x_j \in B} \{\omega \in \Omega : \xi(\omega) = x_j\} = \bigcup_{j: x_j \in B} A_j \end{aligned}$$

Так как события A_j не пересекаются, то

$$\xi = \xi(\omega) = \sum_{k=1}^n I_{A_k}(\omega) x_k$$

$$A(\xi) = \{A = \{\omega \in \Omega : \xi(\omega) \in B\}, B \in \Sigma\}$$

(?? и что с этого???)

Об ошибках в понимании случайной величины (!!!)

(потом соберем их)

Независимость случайных величин

(16 зубков)

Определение независимости случайных величин

Definition 3.20. Случайные величины $\xi_1, \dots, \xi_n$, заданные на одном и том же вероятностном пространстве и принимающие значения в измеримых пространствах $(S_1, \Sigma_1), \dots, (S_n, \Sigma_n)$ соответственно, называются независимыми в совокупности, если независимы порожденные ими σ -алгебры $\mathcal{F}_{\xi_1}, \dots, \mathcal{F}_{\xi_n}$ или, что то же самое, если для любых событий $B_1 \in \Sigma_1, \dots, B_n \in \Sigma_n$

$$\mathbf{P} \{\xi_1 \in B_1, \dots, \xi_n \in B_n\} = \mathbf{P} \{\xi_1 \in B_1\} \dots \mathbf{P} \{\xi_n \in B_n\}$$

Тут еще пара теорем от Зубкова

(потом вставляю, пока раз не вставил, значит, скипануть можно)

3.6.3 Random Variable Distribution Function

вкратце о функции распределения

Полезны формулы:

$$F_\xi(x+0) - F_\xi(x-0) = \mathbf{P}\{\xi = x\} \quad \forall x$$

? что еще???

О применениях функции распределения

(пока четко не вижу)

Функция распределения используется для получения плотности случайной величины.

Определение функции распределения и ее свойства

Случайная величина ξ со значениями в множестве S порождает меру на подмножествах $B \subseteq S, B \in \Sigma$ по формуле

$$\mathbf{P}_\xi\{B\} = \mathbf{P}\{\xi(\omega) \in B\}, \quad B \in \Sigma$$

Эта мера называется законом распределения (или просто распределением) случайной величины ξ

Закон распределения является мерой на пространстве значений случайной величины и не связан жестко с каким-либо вероятностным пространством.

Случайные величины с одним и тем же законом распределения могут быть определены на разных вероятностных пространствах, и на одном вероятностном пространстве можно определить разные случайные величины с одним и тем же законом распределения. Так как в теории вероятностей, как правило, изучаются только свойства распределений случайных величин, то случайные величины с одинаковыми распределениями обычно считаются эквивалентными

Если случайная величина ξ принимает действительные значения, то закон ее распределения удобно описывать функцией распределения

$$F_\gamma(x) = \mathbf{P}\{\xi \leq x\}, \quad -\infty < x < \infty$$

Свойства функции распределения

(их помню, потом посмотрю док-ва)

Теорема

Функция распределения $F(x)$ обладает следующими свойствами:

1. $F(x)$ не убывает,
2. $\lim_{x \rightarrow -\infty} F(x) = 0, \quad \lim_{x \rightarrow \infty} F(x) = 1$
3. $F(x)$ непрерывна справа $\mathbf{P}\{\xi < x\} = F(x - 0) = \lim_{y \uparrow x} F(y)$

Доказательство

(???? не смотрел, ибо не актуально)

Посмотрим на каждый пункт в отдельности.

1. Если $x < y$, то $\{\xi \leq y\} \supseteq \{\xi \leq x\}$ и

$$F(y) - F(x) = \mathbf{P}\{x < \xi \leq y\} \geq 0$$

2. Из равенства $\Omega = \bigcup_{n \geq 1} (C_n \setminus C_{n-1})$, где $C_n = \{\omega : -n < \xi \leq n\}, C_0 = \emptyset$, получаем, что

$$\begin{aligned} 1 = \mathbf{P}\{\Omega\} &= \sum_{n \geq 1} \mathbf{P}\{C_n \setminus C_{n-1}\} = \lim_{N \rightarrow \infty} \sum_{n=1}^N \mathbf{P}\{C_n \setminus C_{n-1}\} = \\ &= \mathbf{P}\{C_N\} = \lim_{N \rightarrow \infty} (F(N) - F(-N)) \end{aligned}$$

Теперь из того, что $F(x) \in [0, 1]$ при всех x , следуют равенства

$$\lim_{x \rightarrow -\infty} F(x) = 0, \quad \lim_{x \rightarrow \infty} F(x) = 1$$

3. Рассмотрим последовательность непересекающихся множеств

$$B_n = \left(\omega : x + \frac{1}{n+1} < \xi(\omega) \leq x + \frac{1}{n} \right), \quad n = 1, 2, \dots$$

Тогда $\{\omega : x < \xi \leq x+1\} = \bigcup_{n \geq 1} B_n$ и

$$\begin{aligned} F(x+1) - F(x) &= \mathbf{P}\{x < \xi \leq x+1\} = \mathbf{P}\left\{\bigcup_{n \geq 1} B_n\right\} = \\ &= \sum_{n \geq 1} \mathbf{P}\{B_n\} = \sum_{n \geq 1} \left(F\left(x + \frac{1}{n}\right) - F\left(x + \frac{1}{n+1}\right) \right) = \\ &= \lim_{N \rightarrow \infty} \sum_{n=1}^N \left(F\left(x + \frac{1}{n}\right) - F\left(x + \frac{1}{n+1}\right) \right) = \lim_{N \rightarrow \infty} \left(F(x+1) - F\left(x + \frac{1}{N+1}\right) \right) \end{aligned}$$

Следовательно,

$$F(x) = \lim_{N \rightarrow \infty} F\left(x + \frac{1}{N+1}\right) = F(x+0) = \lim_{y \downarrow x} F(y)$$

Аналогично, рассматривая непересекающиеся множества

$$A_1 = \{\omega : \xi \leq x-1\}, \quad A_n = \left\{ x - \frac{1}{n-1} < \xi_n \leq x - \frac{1}{n} \right\}, \quad n \geq 2$$

для которых $\{\omega : \xi < x\} = \bigcup_{n \geq 1} A_n$, находим

$$\begin{aligned} \mathbf{P}\{\xi < x\} &= \sum_{n \geq 1} \mathbf{P}\{A_n\} = \\ &= F(x-1) + \sum_{n \geq 2} \left(F\left(x - \frac{1}{n}\right) - F\left(x - \frac{1}{n-1}\right) \right) = \\ &= \lim_{N \rightarrow \infty} \left(F(x-1) + \sum_{n=2}^N \left(F\left(x - \frac{1}{n}\right) - F\left(x - \frac{1}{n-1}\right) \right) \right) = \\ &= \lim_{N \rightarrow \infty} F\left(x - \frac{1}{N}\right) = \lim_{y \uparrow x} F(y) = F(x-0) \end{aligned}$$

Следствие. $F_\xi(x+0) - F_\xi(x-0) = \mathbf{P}\{\xi = x\}$ при любом x .

$F'(x) = p_\xi(x)$ - **связь с плотностью распределения**

плотностью распределения ξ есть производная $F'(x) = p_\xi(x)$, если она существует. (??? и про неё нужно еще что-то сказать??)

3.6.4 Joint Distributions

Рассмотрим распределение нескольких случайных величин.

Вероятность совместного распределения

На вероятностном пространстве $(\Omega, \mathcal{F}, \mathbf{P})$ по набору случайных величин $\xi_1 = \xi_1(\omega), \dots, \xi_k = \xi_k(\omega)$ можно построить новую случайную величину: случайный вектор

$$\Xi = \Xi(\omega) = (\xi_1, \dots, \xi_k), \quad \omega \in \Omega$$

Пусть случайные величины имеют измеримые пространства значений $(S_1, \Sigma_1), \dots, (S_k, \Sigma_k)$ принимающий значения в пространстве $S = S_1 \times \dots \times S_k$.

В пространстве S рассматривается σ -алгебра подмножеств, порожденная прямыми произведениями элементов $\xi_1, \dots, \xi_k$ пространстве $(\Omega, \mathcal{F}, \mathbf{P})$ задан случайный вектор

$$\Xi = \Xi(\omega) = (\xi_1, \dots, \xi_k), \quad \omega \in \Omega$$

со значениями в пространстве $S = S_1 \times \dots \times S_k$, То по нему однозначно определяются образующие его случайные компоненты

$$\xi_1 = \xi_1(\omega), \dots, \xi_k = \xi_k(\omega)$$

а по распределению Ξ в пространстве $S = S_1 \times \dots \times S_k$ - распределения его компонент:

$$\mathbf{P} \{ \xi_j \in B_j \} = \mathbf{P} \{ \Xi \in S_1 \times \dots \times S_{j-1} \times B_j \times S_{j+1} \times \dots \times S_k \}$$

Компонент случайного вектора называются его маргинальными распределениями.

О случаях применения совместного распределения (!!!)

(!! что-то вроде я уже напарывался на то, что задача не решается, потому что не владею этой темой!! соберу её скоро мб!!)

Совместное распределение для двух случайных величин

В простейшем случае на одном вероятностном пространстве $(\Omega, \mathcal{F}, \mathbf{P})$ есть две случайные величины $\xi = \xi(\omega)$ и $\eta = \eta(\omega)$. Пусть $x_1, \dots, x_m$ И $y_1, \dots, y_n$ - все возможные значения случайных величин ξ и η соответственно:

$$\begin{aligned} \mathbf{P} \{ \xi = x_r \} &= p_r, \quad r = 1, \dots, m, \quad p_1 + \dots + p_m = 1 \\ \mathbf{P} \{ \eta = y_s \} &= q_s, \quad s = 1, \dots, n, \quad q_1 + \dots + q_n = 1 \end{aligned}$$

случайных величин ξ и η Определяется вероятностями

$$v_{rs} = \mathbf{P} \{ \xi = x_r, \quad \eta = y_s \}, \quad r \in \{1, \dots, m\}, \quad s \in \{1, \dots, n\}$$

Пределием $\{v_{rs}\}$ и маргинальными распределениями $\{p_r\}$ и $\{q_s\}$ описывается соотношениями

$$\begin{aligned} p_r &= \mathbf{P} \{ \xi = x_r \} = \mathbf{P} \{ (\xi, \eta) \in \{x_r\} \times \{y_1, \dots, y_n\} \} = \sum_{s=1}^n v_{rs} \\ q_s &= \mathbf{P} \{ \eta = y_s \} = \mathbf{P} \{ (\xi, \eta) \in \{x_1, \dots, x_m\} \times \{y_s\} \} = \sum_{r=1}^m v_{rs} \end{aligned}$$

(тут хотя бы одну типичную задачу вставляю, ибо на формулу выше много задач есть!!!)

Функция распределения совместных величин

(из горяйнова, пока не очень зашло)

Совместное распределение случайных величин также можно описать функцией распределения. Пусть ξ и η - две случайные величины, определенные на одном вероятностном пространстве $(\Omega, \mathcal{A}, \mathbf{P})$. Тогда их совместная функция распределения есть:

$$F_{\xi, \eta}(x, y) = \mathbf{P}(\{\omega \in \Omega : \xi(\omega) \leq x, \eta(\omega) \leq y\})$$

$(x, y) \in \mathbb{R}^2$. Отметим свойства $F_{\xi, \eta}$ как функции двух переменных.

1*. $F_{\xi, \eta}(x, y)$ не убывает по каждой переменной;

2*. $F_{\xi, \eta}(x, y)$ непрерывна справа по каждой переменной и

$$F_{\xi, \eta}(-\infty, y) = F_{\xi, \eta}(x, -\infty) = 0, \quad F_{\xi, \eta}(+\infty, +\infty) = 1$$

3*. Для всех $x, y \in \mathbb{R}, \Delta x \geq 0, \Delta y \geq 0$ выполняется неравенство

$$F_{\xi, \eta}(x + \Delta x, y + \Delta y) - F_{\xi, \eta}(x, y + \Delta y) - F_{\xi, \eta}(x + \Delta x, y) + F_{\xi, \eta}(x, y) \geq 0$$

Свойства 1* – 3* являются характеристическими для совместной функции распределения.

Кроме того, совместная функция распределения содержит в себе индивидуальные (одномерные) функции распределения:

$$F_{\xi}(x) = \lim_{y \rightarrow \infty} F_{\xi, \eta}(x, y), \quad F_{\eta}(y) = \lim_{x \rightarrow \infty} F_{\xi, \eta}(x, y)$$

Определение параметров для совместного распределения (!?!)

(очень часто именно в этом и задача, найти ковариацию, дисперсию их. как решать - все еще хз.)

Совместное распределение для более чем двух случайных величин**3.6.5 Mathematical Expectation.****Конструкция математического ожидания****Основные формулы математического ожидания**

Наиболее важно знать следующие формулы математического ожидания:

Определим простые случайные величины, как величины, принимающие лишь конечное число разных значений.

Для простой случайной величины $\xi = \xi(\omega)$, принимающей значения из конечного множества $\{x_1, \dots, x_n\}$ с вероятностями $p_1, \dots, p_n$, математическое ожидание было определено равенством

$$\mathbf{M}\xi = \sum_{k=1}^n x_k p_k = \sum_{k=1}^n x_k \mathbf{P}\{\xi = x_k\}$$

Если распределение случайной величины ξ имеет плотность $p_{\xi}(x)$ $\int_{\mathbb{R}} |x| p_{\xi}(x) dx < \infty$, то

$$\mathbf{M}\xi = \int_{\mathbb{R}} x p_{\xi}(x) dx$$

Определение

Пусть $(\Omega, \mathcal{F}, \mathbf{P})$ - вероятностное пространство, и пространство элементарных $\mathbf{P}(\omega)$ на элементах ω пространства элементарных событий Ω . Пусть, далее, $\xi = \xi(\omega)$ случайная величина, определенная на $(\Omega, \mathcal{F}, \mathbf{P})$ и принимающая действительные значения.

Definition 3.21. *математическим ожиданием называется величина*

$$\mathbf{M}\xi = \sum_{\omega \in \Omega} \xi(\omega) \mathbf{P}(\omega) = S_+(\xi) + S_-(\xi)$$

если обе величины

$$S_+(\xi) = \sum_{\omega \in \Omega: \xi(\omega) \geq 0} \xi(\omega) \mathbf{P}(\omega), \quad S_-(\xi) = \sum_{\omega \in \Omega: \xi(\omega) < 0} \xi(\omega) \mathbf{P}(\omega)$$

конечны; если же хотя бы одна из них бесконечна, то говорят, что случайная величина ξ не имеет математического ожидания.

Основные свойства математического ожидания (?)

(излишнее перенесу потом в главу ниже)

Рассмотрим основные свойства математического ожидания

Конечная аддитивность математического ожидания: Если каждая из случайных величин $\xi_1, \dots, \xi_n$ имеет конечное математическое ожидание, то

$$\mathbf{M}(\xi_1 + \dots + \xi_n) = \mathbf{M}\xi_1 + \dots + \mathbf{M}\xi_n$$

Доказательство

Рассмотрим сначала сумму двух случайных величин. Ввиду абсолютной сходимости рядов $\sum_{\omega \in \Omega} \xi(\omega) \mathbf{P}(\omega)$ и $\sum_{\omega \in \Omega} \eta(\omega) \mathbf{P}(\omega)$ их члены можно переставлять, не изменяя значения суммы, и поэтому

$$\begin{aligned} \mathbf{M}\xi + \mathbf{M}\eta &= \sum_{\omega \in \Omega} \xi(\omega) \mathbf{P}(\omega) + \sum_{\omega \in \Omega} \eta(\omega) \mathbf{P}(\omega) = \\ &= \sum_{\omega \in \Omega} (\xi(\omega) + \eta(\omega)) \mathbf{P}(\omega) = \mathbf{M}(\xi + \eta) \end{aligned}$$

В последнем переходе использовано определение математического ожидания. Общий случай следует из этого частного индукцией по числу слагаемых:

$$\begin{aligned} \mathbf{M}(\xi_1 + \dots + \xi_n) &= \mathbf{M}((\xi_1 + \dots + \xi_{n-1}) + \xi_n) = \\ &= (\mathbf{M}\xi_1 + \dots + \mathbf{M}\xi_{n-1}) + \mathbf{M}\xi_n \end{aligned}$$

Условие конечности математических ожиданий слагаемых существенно. Например, если $\gamma = \gamma(\omega)$ - случайная величина, не имеющая математического ожидания, а $\xi(\omega) = -\gamma(\omega)$, то $\xi(\omega) + \gamma(\omega) = 0$ при любом $\omega \in \Omega$, и поэтому $\mathbf{M}(\xi + \gamma) = 0$, однако $\mathbf{M}(\xi + \gamma) \neq \mathbf{M}\xi + \mathbf{M}\gamma$, так как правая часть не определена.

Счетная аддитивность: если случайные величины $\xi_1, \xi_2, \dots$ удовлетворяют условию $\sum_{n=1}^{\infty} \mathbf{M}|\xi_n| < \infty$, то

$$\mathbf{M} \left\{ \sum_{n=1}^{\infty} \xi_n \right\} = \sum_{n=1}^{\infty} \mathbf{M}\xi_n$$

Действительно, правую часть (24) можно представить в виде двойного ряда, который по условию (23) сходится абсолютно:

$$\sum_{n=1}^{\infty} M\xi_n = \sum_{n=1}^{\infty} \sum_{\omega \in \Omega} \xi_n(\omega) P(\omega), \quad \sum_{n=1}^{\infty} \sum_{\omega \in \Omega} |\xi_n(\omega) P(\omega)| = \sum_{n=1}^{\infty} M|\xi_n| < \infty$$

а левая часть (24) получается из этого двойного ряда изменением порядка суммирования:

$$M \left\{ \sum_{n=1}^{\infty} \xi_n \right\} = \sum_{\omega \in \Omega} \left(\sum_{n=1}^{\infty} \xi_n(\omega) \right) P(\omega)$$

для абсолютно сходящегося ряда такая операция не изменяет его сумму.

Линейность: если $b, c \in R$ - константы, а ξ - случайная величина, то

$$M(c\xi + b) = cM\xi + b$$

Действительно,

$$M(c\xi + b) = \sum_{\omega \in \Omega} (c\xi(\omega) + b)P(\omega) = c \sum_{\omega \in \Omega} \xi(\omega)P(\omega) + b \sum_{\omega \in \Omega} P(\omega) = cM\xi + b$$

Теоремы о предельном переходе математического ожидания (!!!!)

(тут потом немного комментариев добавлю, чисто как обзорный пар-ф. подробнее - в части ниже. и нужно потренироваться, чисто анализ!!!)

Основные теоремы о предельном переходе под знаком математического ожидания есть следующие.

Теорема

[О монотонной сходимости] Пусть $\xi_1, \xi_2, \dots$ - неотрицательные случайных величин и $\xi_n \nearrow \xi$. Тогда $M\xi_n \nearrow M\xi$ при $n \rightarrow \infty$

Здесь для удобства формулировки теоремы допускается случай $M\xi = +\infty$. Тогда и $M\xi_n \rightarrow \infty$ при $n \rightarrow \infty$. Это допущение предполагается и в последующих формулировках теорем.

(тут комментарий про идею док-ва, пока хз)

Теорема

[Лемма Фату] Пусть $\xi_1, \xi_2, \dots$ — неотрицательные случайные величины. Тогда

$$M \liminf_{n \rightarrow \infty} \xi_n \leq \liminf_{n \rightarrow \infty} M\xi_n$$

(тут комментарий про идею док-ва и применения, пока хз)

Теорема

[Лебега о мажорируемой сходимости] Пусть $\eta, \xi_1, \xi_2, \dots$ — случайные величины, для которых выполняются следующие условия: $|\xi_n| \leq \eta$ для всех n , $\xi_n \rightarrow \xi$ при $n \rightarrow \infty$ и $M\eta < \infty$. Тогда

$$\lim_{n \rightarrow \infty} M\xi_n = M\xi, \quad \lim_{n \rightarrow \infty} M|\xi_n - \xi| = 0$$

(тут комментарий про идею док-ва и применения, пока хз)

Мультипликативное свойство математических ожиданий для независимых случайных величин с не более чем счетными значениями

(вроде, всё просто, да? зубков 17)

Теорема

Если случайные величины $\xi_1, \dots, \xi_n$ определены на одном вероятностном пространстве $(\Omega, \mathcal{F}, \mathbf{P})$, принимают значения из конечного или счетного множества действительных чисел, имеют конечные математические ожидания и независимы, то

$$\mathbf{M}(\xi_1 \dots \xi_n) = \mathbf{M}\xi_1 \cdot \dots \cdot \mathbf{M}\xi_n$$

Доказательство проведем индукцией по числу сомножителей.

(??? тут укажу идею док-ва, пока не так уловил)

Основным случаем является база индукции: $n = 2$. Пусть ξ и η - независимые случайные величины; будем обозначать через $x_1, x_2, \dots$ все возможные значения ξ , через $y_1, y_2, \dots$ - все возможные значения η . Случайные величины ξ и η порождают разбиения

$$\begin{aligned} A_j &= \{\omega \in \Omega : \xi(\omega) = x_j\}, \quad j \geq 1 \\ B_k &= \{\omega \in \Omega : \eta(\omega) = y_k\}, \quad k \geq 1 \end{aligned}$$

пространства элементарных событий Ω ; тогда

$$\begin{aligned} \xi &= \sum_{j \geq 1} x_j I(A_j), \quad \mathbf{M}\xi = \sum_{j \geq 1} x_j \mathbf{P}\{A_j\} \\ \eta &= \sum_{k \geq 1} y_k I(B_k), \quad \mathbf{M}\eta = \sum_{k \geq 1} y_k \mathbf{P}\{B_k\} \end{aligned}$$

и

$$\begin{aligned} \xi\eta &= \left(\sum_{j \geq 1} x_j I(A_j) \right) \cdot \left(\sum_{k \geq 1} y_k I(B_k) \right) = \\ &= \sum_{j \geq 1} \sum_{k \geq 1} x_j y_k I(A_j) I(B_k) = \sum_{j \geq 1} \sum_{k \geq 1} x_j y_k I(A_j B_k) \end{aligned}$$

Остается воспользоваться возможностью почленного умножения абсолютно сходящихся рядов для $\mathbf{M}\xi$ и $\mathbf{M}\eta$, независимостью ξ и η и определением математического ожидания:

$$\begin{aligned} \mathbf{M}\xi\mathbf{M}\eta &= \left(\sum_{j \geq 1} x_j \mathbf{P}\{A_j\} \right) \left(\sum_{k \geq 1} y_k \mathbf{P}\{B_k\} \right) = \\ &= \sum_{j \geq 1} \sum_{k \geq 1} x_j y_k \mathbf{P}\{A_j\} \mathbf{P}\{B_k\} = \sum_{j \geq 1} \sum_{k \geq 1} x_j y_k \mathbf{P}\{A_j B_k\} = \mathbf{M}\xi\eta \end{aligned}$$

Если утверждение теоремы верно для $n - 1$, то по уже доказанному

$$\mathbf{M}(\xi_1, \dots, \xi_n) = \mathbf{M}((\xi_1 \dots \xi_{n-1}) \xi_n) = \mathbf{M}(\xi_1 \dots \xi_{n-1}) \mathbf{M}\xi_n = \mathbf{M}\xi_1 \cdot \dots \cdot \mathbf{M}\xi_n$$

Теорема доказана.

(тут комментарии и дополнения укажу. зачем теорема нужна, всё такое. пока вроде понятно, ну и ок.)

Математическое ожидание случайной величины с произвольным распределением

(зубков)

Определим простые случайные величины, как величины, принимающие лишь конечное число разных значений.

Для простой случайной величины $\xi = \xi(\omega)$, принимающей значения из конечного множества $\{x_1, \dots, x_n\}$ с вероятностями $p_1, \dots, p_n$, математическое ожидание было определено равенством

$$M\xi = \sum_{k=1}^n x_k p_k = \sum_{k=1}^n x_k \mathbf{P}\{\xi = x_k\}$$

При распространении понятия математического ожидания на произвольные случайные величины, принимающие действительные значения, следует его определить сначала для неотрицательных случайных величин, а потом - для произвольных по следующей схеме:

(?? чего проще не получается?? потом подумая.)

а) для неотрицательной случайной величины $\xi \geq 0$ строится монотонно возрастающая последовательность простых случайных величин $\xi_n(\omega)$

$$\xi_n(\omega) \uparrow \xi(\omega) \text{ при } n \uparrow \infty \text{ для всех } \omega \in \Omega$$

и полагают

$$M\xi = \lim_{n \rightarrow \infty} M\xi_n$$

б) произвольная случайная величина ξ представляется в виде разности двух неотрицательных случайных величин: $\xi = \xi_+ - \xi_-$, где $\xi_+ = \max(\xi, 0)$, $\xi_- = \max(-\xi, 0)$, и полагают $M\xi = M\xi_+ - M\xi_-$

Корректность такого определения математического ожидания обосновывается следующим утверждением.

Теорема

Если $\xi = \xi(\omega)$ - неотрицательная случайная величина, а $\xi_n = \xi_n(\omega)$ - $\eta_n = \eta_n(\omega)$ - неубывающие последовательности простых случайных величин, сходящиеся к $\xi(\omega)$ при каждом $\omega \in \Omega$, то

$$\lim_{n \rightarrow \infty} M\xi_n = \lim_{n \rightarrow \infty} M\eta_n$$

Доказательство

очевидная лемма.

Пусть $\eta, \xi_1, \xi_2, \dots$ - простые случайные величины, $\xi_n \uparrow \xi \geq \eta, n \uparrow \infty$.

Тогда $\lim_n M\xi_n \geq M\eta$

Доказательство леммы.

Фиксируем произвольно $\varepsilon > 0$. События $A_n = \{\omega : \xi_n(\omega) \geq \eta(\omega) - \varepsilon\}$ монотонно возрастают по n и в силу условия леммы $\overline{A_n} \downarrow \emptyset$ значит, $\mathbf{P}\{\overline{A_n}\} \rightarrow 0, \mathbf{P}\{A_n\} \rightarrow 1, n \rightarrow \infty$.

Далее

$$\xi_n \geq \xi_n I_{A_n} \geq (\eta - \varepsilon) I_{A_n} = \eta - \varepsilon I_{A_n} - \eta I_{\overline{A_n}}$$

и по свойствам математических ожиданий простых случайных величин

$$M\xi_n \geq M(\eta - \varepsilon I_{A_n} - \eta I_{\overline{A_n}}) = M\eta - \varepsilon \mathbf{P}\{A_n\} - \eta \mathbf{P}\{\overline{A_n}\}$$

где $z < \infty$ выбрано так, что $\mathbf{P}\{\eta \leq z\} = 1$ (η - простая случайная величина). Переходя к пределу по $n \rightarrow \infty$, получаем:

$$\lim_{n \rightarrow \infty} \mathbf{M}\xi_n \geq \mathbf{M}\eta - \varepsilon \text{ при любом } \varepsilon > 0$$

что эквивалентно утверждению леммы. Перейдем к доказательству теоремы. Для каждого $m < \infty$ имеем по условию $\xi_n \uparrow \xi \geq \eta_m, n \uparrow \infty$, и в силу леммы

$$\lim_n \mathbf{M}\xi_n \geq \mathbf{M}\eta_m \text{ при любом } m = 1, 2, \dots$$

переходя к пределу по $m \rightarrow \infty$, получаем:

$$\lim_{n \rightarrow \infty} \mathbf{M}\xi_n \geq \lim_{m \rightarrow \infty} \mathbf{M}\eta_m$$

что эквивалентно утверждению леммы.

Перейдем к доказательству теоремы.

Для каждого $m < \infty$ имеем по условию $\xi_n \uparrow \xi \geq \eta_m, n \uparrow \infty$, и в силу леммы

$$\lim_n \mathbf{M}\xi_n \geq \mathbf{M}\eta_m \text{ при любом } m = 1, 2, \dots$$

переходя к пределу по $m \rightarrow \infty$, получаем:

$$\lim_{n \rightarrow \infty} \mathbf{M}\xi_n \geq \lim_{m \rightarrow \infty} \mathbf{M}\eta_m$$

В силу симметрии $\lim_{m \rightarrow \infty} \mathbf{M}\eta_m \geq \lim_{n \rightarrow \infty} \mathbf{M}\xi_n$. Следовательно

$$\lim_n \mathbf{M}\xi_n = \lim_n \mathbf{M}\eta_n$$

При переходах, описанных в пп. а) и б), сохраняются все установленные нами ранее свойства математических ожиданий и формулы для их вычислений.

Доказательства этого проводятся по одной и той же схеме.

(?? я так и не понял, чего оно такое большое?)

Мультипликативность математического ожидания в для счетнозначных случайных величин

Покажем, что для независимых случайных величин определенное с помощью предельного перехода математическое ожидание обладает свойством мультипликативности.

Теорема.

Если ξ и η - независимые случайные величины с конечными математическими ожиданиями, то

$$\mathbf{M}\xi\eta = \mathbf{M}\xi\mathbf{M}\eta$$

Доказательство.

а) Для простых независимых случайных величин оно уже доказано.

б) Пусть ξ и η - независимые неотрицательные случайные величины. Введем монотонно неубывающие последовательности простых случайных величин

$$\xi_n = \min \{n, [2^n \xi] / 2^n\} \uparrow \xi, \quad \eta_n = \min \{n, [2^n \eta] / 2^n\} \uparrow \eta$$

При каждом n простые случайные величины ξ_n и η_n независимы как функции от независимых случайных величин ξ и η . Последовательность $\xi_n \eta_n$ состоит из простых случайных величин, монотонна и сходится к случайной величине $\xi\eta$. Из соотношений

$$\mathbf{M}\xi_n \rightarrow \mathbf{M}\xi, \quad \mathbf{M}\eta_n \rightarrow \mathbf{M}\eta, \quad \mathbf{M}\xi_n \eta_n \rightarrow \mathbf{M}\xi\eta, \quad n \rightarrow \infty$$

и из того, что $\mathbf{M}\xi_n\eta_n = \mathbf{M}\xi_n\mathbf{M}\eta_n$, следует, что

$$\mathbf{M}\xi\eta = \lim \mathbf{M}\xi_n\eta_n = \lim \mathbf{M}\xi_n\mathbf{M}\eta_n = \lim \mathbf{M}\xi_n \lim \mathbf{M}\eta_n = \mathbf{M}\xi\mathbf{M}\eta$$

в) Пусть ξ и η - независимые случайные величины, принимающие как положительные, так и отрицательные значения. Тогда

$$\xi = \xi^+ - \xi^-, \quad \eta = \eta^+ - \eta^-, \quad x^+ = \max\{0, x\}, \quad x^- = \max\{0, -x\}$$

Случайные векторы (ξ^+, ξ^-) , (η^+, η^-) независимы как функции от независимых случайных величин ξ и η и имеют неотрицательные компоненты. Поэтому

$$\begin{aligned} \mathbf{M}\xi\eta &= \mathbf{M}(\xi^+ - \xi^-)(\eta^+ - \eta^-) = \\ &= \mathbf{M}\xi^+\eta^+ - \mathbf{M}\xi^+\eta^- - \mathbf{M}\xi^-\eta^+ + \mathbf{M}\xi^-\eta^- = \\ &= \mathbf{M}\xi^+\mathbf{M}\eta^+ - \mathbf{M}\xi^+\mathbf{M}\eta^- - \mathbf{M}\xi^-\mathbf{M}\eta^+ + \mathbf{M}\xi^-\mathbf{M}\eta^- = \\ &= (\mathbf{M}\xi^+ - \mathbf{M}\xi^-)(\mathbf{M}\eta^+ - \mathbf{M}\eta^-) = \mathbf{M}\xi\mathbf{M}\eta \end{aligned}$$

Об ошибках в понимании математического ожидания

(хз)

Формулы для вычисления математического ожидания

(зубков)

(потом перепису, я пока не пойму, очевидные же вещи, если это так, то вынесу их отдельно в параграф до этого)

Формулы для математических ожиданий случайных величин с дискретными распределениями выведены раньше.

(да?)

Теория

Для неотрицательной случайной величины ξ с действительными значениями и произвольным распределением математическое ожидание определяется по общей схеме с помощью неубывающей последовательности простых случайных величин, сходящуюся к ξ при $n \rightarrow \infty$,

(????)

например, с помощью последовательности $\xi_n(\omega) = \min \left\{ n, \frac{[2^n \xi(\omega)]}{2^n} \right\}$

$$\mathbf{M}\xi = \lim_{n \rightarrow \infty} \mathbf{M}\xi_n = \lim_{n \rightarrow \infty} \sum_{k=1}^{n2^n} \frac{k-1}{2^n} \mathbf{P} \left\{ \omega : \frac{k-1}{2^n} < \xi(\omega) \leq \frac{k}{2^n} \right\}$$

Этот предел является интегралом Лебега по пространству элементарных событий Ω и обозначается $\int_{\Omega} \xi(\omega) d\mathbf{P}\{\omega\}$

(??????)

Неотрицательная случайная величина определяет меру $P_{\xi}(B) = \mathbf{P}\{\xi \in B\}$ на полуоси $[0, \infty)$ и функцию распределения $F_{\xi}(x)$. В этих терминах формулу для $\mathbf{M}\xi$ можно

представить в виде

$$\begin{aligned} \mathbf{M}\xi &= \lim_{n \rightarrow \infty} \mathbf{M}\xi_n = \lim_{n \rightarrow \infty} \sum_{k=1}^{n2^n} \frac{k-1}{2} \mathbf{P} \left\{ \left(\frac{k-1}{2^n} < \xi(\omega) \leq \frac{k}{2^n} \right) \right\} = \\ &= \lim_{n \rightarrow \infty} \sum_{k=1}^{n2^n} \frac{k-1}{2} \left(F_\xi \left(\frac{k}{2^n} \right) - F_\xi \left(\frac{k-1}{2^n} \right) \right) = \int_0^\infty x dF(x) \end{aligned}$$

т.е. в виде интеграла Лебега-Стилтьеса.

Случайную величину ξ , принимающую значения обоих знаков, можно представить в виде разности двух неотрицательных случайных величин ξ^+ и ξ^- , и, применяя к каждой из них последнюю формулу и объединяя результаты, получим следующее выражение:

$$\mathbf{M}\xi = \int_R x dF(x)$$

Из него следует формула для вычисления математического ожидания случайной величины, имеющей абсолютно непрерывное распределение.

Математическое ожидание для величины с плотностью распределения

(!! думаю, это важно, в какой-то задаче такая же конструкция использовалась!!)

Утверждение 1.

Если распределение случайной величины ξ имеет плотность $p_\xi(x)$ и $\int_{\mathbb{R}} |x| p_\xi(x) dx < \infty$,

то

$$\mathbf{M}\xi = \int_{\mathbb{R}} x p_\xi(x) dx$$

Доказательство.

При условиях утверждения

$$F_\xi(x) = \mathbf{P}\{\xi \leq x\} = F_\xi(0) + \int_0^x p_\xi(u) du$$

и поэтому по правилам интегрального исчисления

$$\mathbf{M}\xi = \int_{\mathbb{R}} x dF_\xi(x) = \int_{\mathbb{R}} x d \left(F_\xi(0) + \int_0^x p_\xi(u) du \right) = \int_{\mathbb{R}} x p_\xi(x) dx$$

Немного сложнее (фактически повторяя процедуру определения математического ожидания) проводится доказательство формулы для математического ожидания функции от случайной величины.

Lemma 4 (Математическое ожидание от функции от случайной величины). Если распределение случайной величины ξ имеет плотность $p_\xi(x)$, функция $g(x)$ кусочно-непрерывна, а также $\int_{\mathbb{R}} |g(x)| p_\xi(x) dx < \infty$, то

$$\mathbf{M}g(\xi) = \int_{\mathbb{R}} g(x) p_\xi(x) dx$$

Доказательство.

Как и при определении математического ожидания, достаточно рассмотреть случай, когда $g(x) \geq 0$ при всех x .

Прежде всего заметим, что формула (42) верна, если $g(x)$ - ступенчатая функция.

Действительно, в этом случае существует такое разбиение действительной оси на не более чем счетное множество интервалов $\Delta_k = (a_k, b_k)$, $k = 1, 2, \dots$, что $g(x) = g_k$ для всех $x \in \Delta_k$.

Множество концов этих интервалов относительно непрерывного распределения имеет нулевую меру и поэтому не влияет на значения математического ожидания и интеграла.

Отсюда и из свойств счетной аддитивности математического ожидания и интеграла следует, что

$$\begin{aligned} \mathbf{M}g(\xi) &= \sum_{k \geq 1} \mathbf{M}g(\xi) \mathbb{I}\{\xi \in \Delta_k\} = \\ &= \sum_{k \geq 1} g_k \mathbf{P}\{\xi \in \Delta_k\} = \sum_{k \geq 1} g_k \int_{\Delta_k} p_\xi(x) dx = \int_{\mathbb{R}} g(x) p_\xi(x) dx \end{aligned}$$

Если функция $g(x)$ кусочно-непрерывна, то существует разбиение действительной оси на не более чем счетное множество интервалов, на каждом из которых функция $g(x)$ непрерывна. Аналогично предыдущему отсюда следует, что достаточно доказать равенство

$$\mathbf{M}g(\xi) \mathbb{I}\{\xi \in (a, b)\} = \int_a^b g(x) p_\xi(x) dx$$

для случая, когда $g(x)$ непрерывна на конечном интервале (a, b) $\int_a^b |g(x)| p_\xi(x) dx < \infty$ Для доказательства, во-первых, построим ступенчатую функцию $g^*(x)$, $x \in (a, b)$ удовлетворяющую условию $g(x) \leq g^*(x) \leq g(x) + 1$ при всех $x \in (a, b)$ (это возможно в силу непрерывности $g(x)$). Тогда случайная величина $g^*(\xi)$ имеет дискретное распределение и $\mathbf{M}g^*(\xi) \leq \int_a^b (g(x) + 1) p_\xi(x) dx < \infty$ Во-вторых, построим последовательность ступенчатых функций $g_n(x)$, сходящуюся к $g(x)$ при всех $x \in (a, b)$. Тогда последовательность случайных величин $g_n(\xi)$ сходится к случайной величине $g(\xi)$, которая не больше случайной величины $g^*(\xi)$ с конечным математическим ожиданием. По теореме о мажорируемой сходимости

$$\lim_{n \rightarrow \infty} \mathbf{M}g_n(\xi) = \mathbf{M}g(\xi)$$

но, как было показано выше, $\mathbf{M}g_n(\xi) = \int_a^b g_n(x) p_\xi(x) dx$, а эти интегралы при $n \rightarrow \infty$ сходятся к $\int_a^b g(x) p_\xi(x) dx$ по построению функций $g_n(x)$. Утверждение доказано.

3.6.6 Variance

Обсудим наиболее используемую теорию дисперсии.

Конструкция дисперсии

Теория

Центральный момент второго порядка всегда неотрицателен и имеет собственное название - дисперсия-и два обозначения: $D\xi$ и $\sigma^2(\xi)$, поскольку является важной и удобной характеристикой распределения случайной величины:

$$D\xi = \sigma^2(\xi) = M(\xi - M\xi)^2$$

Так же, как математическое ожидание является средним значением случайной величины, дисперсия является средним значением квадрата отклонения случайной величины от ее математического ожидания. Величина $\sigma(\xi) = \sqrt{D\xi}$ называется среднеквадратическим отклонением случайной величины ξ и характеризует степень разброса значений случайной величины ξ около $M\xi$. Свойства дисперсии: 1. Выражение через степенные моменты:

$$D\xi = M\xi^2 - (M\xi)^2$$

О применениях дисперсии

(пропишу четко)

3.6.7 Power Moments

кратко о моментах

что тут будет и как их используем?

????

Степенные моменты случайных величин

(14 зубков)

Пусть ξ - случайная величина, принимающая значения из множества $\{x_1, x_2, \dots\}$ а $g(x)$ - Числовая функция. Тогда $\eta = g(\xi)$ - тоже случайная величина,

$$\eta = \sum_{k \geq 1} g(x_k) I \{\xi = x_k\}$$

$$Mg(\xi) = M\eta = \sum_{k \geq 1} g(x_k) P \{\xi = x_k\}$$

Полагая $g(x) = x^n$, $g(x) = |x|^n$, $g(x) = (x - M\xi)^n$, мы получаем величины, каждая из которых имеет свое название:

момент n-го порядка:

$$M\xi^n = \sum_{k \geq 1} x_k^n P \{\xi = x_k\}$$

абсолютный момент n-го порядка:

$$M|\xi|^n = \sum_{k \geq 1} |x_k|^n P \{\xi = x_k\}$$

центральный момент n-го порядка:

$$M(\xi - M\xi)^n = \sum_{k \geq 1} (x_k - M\xi)^n P \{\xi = x_k\}$$

Способы определения момента для данного распределения (??)

(?? есть на это задачи, пока хз, как решать.)

О применениях k-х моментов (???)

3.6.8 Covariation

(18, зубков)

Основные формулы ковариации

(с опытом пропишу)

определение и основные свойства

Определение.

Пусть ξ и η - случайные величины с действительными значениями и конечными математическими ожиданиями.

Ковариацией случайных величин ξ и η называется величина

$$\text{cov}(\xi, \eta) = \mathbf{M}[(\xi - \mathbf{M}\xi)(\eta - \mathbf{M}\eta)]$$

Изучим свойства.

Очевидно, $\text{cov}(\xi, \eta) = \text{cov}(\eta, \xi)$.

1. В частном случае ковариация это дисперсия.

Если случайная величина ξ имеет математическое ожидание (когда?), то $\text{cov}(\xi, \xi) = \mathbf{D}\xi$.

Утверждение следует из определений ковариации и дисперсии:

$$\text{cov}(\xi, \xi) = \mathbf{M}(\xi - \mathbf{M}\xi)(\xi - \mathbf{M}\xi) = \mathbf{M}(\xi - \mathbf{M}\xi)^2 = \mathbf{D}\xi$$

2. Для любых случайных величин с конечными математическими ожиданиями

$$|\text{cov}(\xi, \eta)| \leq \sqrt{\mathbf{D}\xi \mathbf{D}\eta}$$

Определение.

Отношение

$$\frac{\text{cov}(\xi, \eta)}{\sqrt{\mathbf{D}\xi \mathbf{D}\eta}} = \text{corr}(\xi, \eta)$$

называют корреляцией случайных величин ξ и η

Из только что доказанного неравенства следует, что корреляция (если она определена, т.е. если $\mathbf{D}\xi$ и $\mathbf{D}\eta$ конечны) по модулю не превосходит 1 и что она имеет такой же знак, как и ковариация.

3. Если случайные величины ξ и η независимы и имеют конечные математические ожидания, то $\text{cov}(\xi, \eta) = 0$. Существуют зависимые случайные величины, ковариация которых равна 0. Доказательство. Первое утверждение следует из мультипликативности математического ожидания для произведения независимых случайных величин:

$$\begin{aligned} \text{cov}(\xi, \eta) &= \mathbf{M}\{(\xi - \mathbf{M}\xi)(\eta - \mathbf{M}\eta)\} = \mathbf{M}(\xi - \mathbf{M}\xi)\mathbf{M}(\eta - \mathbf{M}\eta) = \\ &= (\mathbf{M}\xi - \mathbf{M}\xi)(\mathbf{M}\eta - \mathbf{M}\eta) = 0 \end{aligned}$$

Для доказательства второго утверждения достаточно привести пример двух случайных величин, которые не являются независимыми, но имеют нулевую ковариацию.

Рассмотрим пару случайных величин (ξ, η) , распределение которой состоит из 4 атомов веса $1/4$, расположенных в концах единичных векторов координатных осей:

$$\begin{aligned} \mathbf{P}\{(\xi, \eta) = (0, 1)\} &= \mathbf{P}\{(\xi, \eta) = (0, 1)\} = \\ &= \mathbf{P}\{(\xi, \eta) = (-1, 0)\} = \mathbf{P}\{(\xi, \eta) = (0, -1)\} = \\ &= \mathbf{P}\{(\xi, \eta) = (-1, 0)\} = \mathbf{P}\{(\xi, \eta) = (0, -1)\} = \frac{1}{4} \end{aligned}$$

Тогда

$$\begin{aligned} \mathbf{P}\{\xi = 1\} &= \mathbf{P}\{\eta = 1\} = \mathbf{P}\{\xi = -1\} = \mathbf{P}\{\eta = -1\} = \frac{1}{4} \\ \mathbf{P}\{\xi = 0\} &= \mathbf{P}\{\eta = 0\} = \frac{1}{2} \end{aligned}$$

и

$$0 = \mathbf{P}\{\xi = 0, \eta = 0\} \neq \mathbf{P}\{\xi = 0\}\mathbf{P}\{\eta = 0\} = \frac{1}{4}$$

Таким образом, если случайные величины ξ и η независимы и имеют конечные математические ожидания, то их ковариация и корреляция равны 0. С другой стороны, если случайные величины ξ и η линейно связаны: $\eta = a\xi + b$ - то $\mathbf{M}\eta = a\mathbf{M}\xi + b$, $\mathbf{D}\eta = a^2\mathbf{D}\xi$ и

$$\begin{aligned} \text{cov}(\xi, \eta) &= \mathbf{M}(\xi - \mathbf{M}\xi)(\eta - \mathbf{M}\eta) = \\ &= \mathbf{M}(\xi - \mathbf{M}\xi)(a\xi + b - (a\mathbf{M}\xi + b)) = a\mathbf{M}(\xi - \mathbf{M}\xi)^2 = a\mathbf{D}\xi \\ \text{corr}(\xi, \eta) &= \frac{\text{cov}(\xi, \eta)}{\sqrt{\mathbf{D}\xi\mathbf{D}\eta}} = \frac{a\mathbf{D}\xi}{\sqrt{\mathbf{D}\xi a^2\mathbf{D}\xi}} = \text{sign}(a) \end{aligned}$$

т.е. ковариация и корреляция имеют тот же знак, что коэффициент a , и при этом корреляция линейно связанных случайных величин с невырожденными распределениями равна либо 1, либо -1. Таким образом, ковариация и корреляция в какой-то мере характеризуют степень зависимости между случайными величинами.

Случайные величины, ковариация которых равна 0, называют некоррелированными. Свойства независимости и некоррелированности одно из другого не следуют. Независимые случайные величины не являются некоррелированными, если они не имеют математического ожидания.

формула для дисперсии суммы случайных величин

4. Теперь докажем формулу для дисперсии суммы случайных величин. Теорема. Для любых случайных величин $\xi_1, \dots, \xi_n$

$$\mathbf{D}(\xi_1 + \dots + \xi_n) = \sum_{k=1}^n \mathbf{D}\xi_k + 2 \sum_{1 \leq k < m \leq n} \text{cov}(\xi_k, \xi_m)$$

Если случайные величины $\xi_1, \dots, \xi_n$ попарно независимы (или хотя бы попарно некоррелированы), то

$$\mathbf{D}(\xi_1 + \dots + \xi_n) = \mathbf{D}\xi_1 + \dots + \mathbf{D}\xi_n$$

Если пары случайных величин $(\xi_1, \eta_1), \dots, (\xi_n, \eta_n)$ попарно независимы, то

$$\text{cov}(\xi_1 + \dots + \xi_n, \eta_1 + \dots + \eta_n) = \text{cov}(\xi_1, \eta_1) + \dots + \text{cov}(\xi_n, \eta_n)$$

Доказательство. Используя определение дисперсии и аддитивность математического ожидания, находим:

$$\begin{aligned} \mathbf{D}(\xi_1 + \dots + \xi_n) &= \mathbf{M}((\xi_1 + \dots + \xi_n) - \mathbf{M}(\xi_1 + \dots + \xi_n))^2 = \\ &= \mathbf{M}\left(\sum_{k=1}^n \xi_k - \sum_{k=1}^n \mathbf{M}\xi_k\right)^2 = \mathbf{M}\left(\sum_{k=1}^n (\xi_k - \mathbf{M}\xi_k)\right)^2 = \\ &= \mathbf{M}\left\{\sum_{k=1}^n (\xi_k - \mathbf{M}\xi_k)^2 + 2 \sum_{1 \leq k < m \leq n} (\xi_k - \mathbf{M}\xi_k)(\xi_m - \mathbf{M}\xi_m)\right\} = \\ &= \sum_{k=1}^n \mathbf{D}\xi_k + 2 \sum_{1 \leq k < m \leq n} \text{cov}(\xi_k, \xi_m) \end{aligned}$$

Тем самым доказано первое утверждение. Для доказательства второго заметим, что если $\xi_1, \dots, \xi_n$ попарно независимы, то все слагаемые второй суммы равны 0 как ковариации независимых случайных величин.

Из попарной независимости пар $(\xi_1, \eta_1), \dots, (\xi_n, \eta_n)$ следует, что случайные величины ξ_i и η_j независимы при любых $i \neq j$ и поэтому $\text{cov}(\xi_i, \eta_j) = 0$ при $i \neq j$. Следовательно,

$$\begin{aligned} \text{cov}(\xi_1 + \dots + \xi_n, \eta_1 + \dots + \eta_n) &= \\ &= \mathbf{M}\left(\sum_{i=1}^n \xi_i - \mathbf{M}\sum_{i=1}^n \xi_i\right)\left(\sum_{j=1}^n \eta_j - \mathbf{M}\sum_{j=1}^n \eta_j\right) = \\ &= \mathbf{M}\sum_{i=1}^n (\xi_i - \mathbf{M}\xi_i) \sum_{j=1}^n (\eta_j - \mathbf{M}\eta_j) = \\ &= \sum_{i,j=1}^n \mathbf{M}(\xi_i - \mathbf{M}\xi_i)(\eta_j - \mathbf{M}\eta_j) = \\ &= \sum_{i=1}^n \mathbf{M}(\xi_i - \mathbf{M}\xi_i)(\eta_i - \mathbf{M}\eta_i) = \sum_{i=1}^n \text{cov}(\xi_i, \eta_i) \end{aligned}$$

Теорема доказана.

Определение

Если $\xi = (\xi_1, \dots, \xi_d)$ - набор случайных величин, определенных на одном вероятностном пространстве (или, что то же самое, случайный вектор), то матрица

$$B = B(\xi) = \|\text{cov}(\xi_i, \xi_j)\| = \begin{vmatrix} \mathbf{D}\xi_1 & \text{cov}(\xi_1, \xi_2) & \dots & \text{cov}(\xi_1, \xi_d) \\ \text{cov}(\xi_2, \xi_1) & \mathbf{D}\xi_2 & \dots & \text{cov}(\xi_2, \xi_d) \\ \dots & \dots & \dots & \dots \\ \text{cov}(\xi_d, \xi_1) & \text{cov}(\xi_d, \xi_2) & \dots & \mathbf{D}\xi_d \end{vmatrix}$$

называется его матрицей ковариаций или ковариационной матрицей.

Утверждение.

Матрица ковариаций $B(\xi) = \|\text{cov}(\xi_i, \xi_j)\|$ случайного вектора $\xi = (\xi_1, \dots, \xi_d)$ симметрична и неотрицательно определена.

Более того, матрица ковариаций $B(\xi)$ вырождена ($\det B = 0$) тогда и только тогда, когда существует максимальный вектор $x \in R^d \setminus \{0\}$ и число $c \in R$, что $\mathbf{P}\left\{\sum_{i=1}^d x_i \xi_i = c\right\} = 1$

Доказательство. Симметричность матрицы ковариаций следует из того, что $\text{cov}(\xi, \eta) = \text{cov}(\eta, \xi)$. Неотрицательная определенность матрицы $B = \|\text{cov}(\xi_i, \xi_j)\|$ означает, что соответствующая ей квадратичная форма неотрицательна:

$$Q_B(x) = xBx^T \geq 0$$

для любого вектора $x = (x_1, \dots, x_d) \in R^d$ той же размерности, что и порядок матрицы B . Покажем, что $Q_B(x)$ неотрицательно определена:

$$\begin{aligned} Q_B(x) &= (x_1, \dots, x_d) \|\text{cov}(\xi_1, \xi_j)\| (x_1, \dots, x_d)^T = \\ &= \sum_{i,j=1}^d x_i \text{cov}(\xi_i, \xi_j) x_j = \sum_{i,j=1}^d x_i x_j \mathbf{M}(\xi_i - \mathbf{M}\xi_i)(\xi_j - \mathbf{M}\xi_j) = \\ &= \mathbf{M} \sum_{i,j=1}^d x_i x_j (\xi_i - \mathbf{M}\xi_i)(\xi_j - \mathbf{M}\xi_j) = \\ &= \mathbf{M} \left(\sum_{i=1}^d x_i (\xi_i - \mathbf{M}\xi_i) \right)^2 \geq 0 \end{aligned}$$

при любом векторе $x = (x_1, \dots, x_d) \in R^d$. Наконец, вырожденность матрицы B означает, что соответствующая ей квадратичная форма Q_B обращается в нуль хотя бы на одном ненулевом векторе x . Но по свойствам математического ожидания равенство в неравенстве

$$Q_B(x) = \mathbf{M} \left(\sum_{i=1}^d x_i (\xi_i - \mathbf{M}\xi_i) \right)^2 \geq 0$$

возможно только в случае, когда

$$\mathbf{P} \left\{ \sum_{i=1}^d x_i (\xi_i - \mathbf{M}\xi_i) = 0 \right\} = 1$$

Т. е.

$$\mathbf{P} \left\{ \sum_{i=1}^d x_i \xi_i = \sum_{i=1}^d x_i \mathbf{M}\xi_i = c \right\} = 1$$

Утверждение полностью доказано.

6.

В п.4 было доказано, что дисперсии и ковариации сумм независимых случайных величин равны соответственно суммам дисперсий и ковариаций слагаемых. Аналогичное утверждение верно для ковариационных матриц. $(\xi_{n1}, \dots, \xi_{nd})$ независимы, $B(\xi_1), \dots, B(\xi_n)$ — их ковариационные матрицы $u\xi = \xi_1 + \dots + \xi_n = (\Xi_1, \dots, \Xi_d)$, то

$$B(\Xi) = B(\xi_1) + \dots + B(\xi_n)$$

Доказательство. Из независимости векторов $\xi_1, \dots, \xi_n$ следует, что при любых $k, m \in \{1, \dots, d\}$ их компоненты ξ_{ik} и ξ_{jm} независимы, если $i \neq j$. Для доказательства теоремы рассмотрим произвольный элемент $\text{coy}(\Xi_k, \Xi_m)$ ковариационной матрицы вектора

$$\Xi = (\Xi_1, \dots, \Xi_d) = \left(\sum_{i=1}^n \xi_{i1}, \dots, \sum_{i=1}^n \xi_{id} \right)$$

Этот элемент имеет вид

$$\text{cov}(\Xi_k, \Xi_m) = \text{cov}(\xi_{1k} + \dots + \xi_{nk}, \xi_{1m} + \dots + \xi_{nm})$$

причем пары $(\xi_{1k}, \xi_{1m}), \dots, (\xi_{nk}, \xi_{nm})$ независимы. По теореме о ковариации суммы независимых пар случайных величин

$$\text{cov}(\Xi_k, \Xi_m) = \text{cov}(\xi_{1k}, \xi_{1m}) + \dots + \text{cov}(\xi_{nk}, \xi_{nm})$$

Таким образом, каждый элемент ковариационной матрицы вектора Ξ равен сумме соответствующих элементов ковариационных матриц векторов $\xi_1, \dots, \xi_n$. Тем самым Теорема доказана.

7. При умножении случайной величины на константу c ее дисперсия умножается на c^2 . Обобщение этого свойства на ковариационные матрицы выглядит следующим образом.

Утверждение 2.

Если $\bar{\gamma} = (\gamma_1, \dots, \gamma_d) \in \mathbb{R}^d$ - случайный вектор, $B(\gamma)$ - его матрица ковариаций, $c \in \mathbb{R}$, $\bar{a} \in \mathbb{R}^d$, $C = \|c_{ij}\| - (d \times d)$ - матрица, то

$$B(c\gamma + a) = c^2 B \quad \text{и} \quad B(C\gamma + a) = CB(\gamma)C^T$$

Доказательство.

Первое равенство следует из того, что при любых i, j

$$\text{cov}(c\gamma_i + a_i, c\gamma_j + a_j) = \mathbf{M}(c\gamma_i - \mathbf{M}c\gamma_i, c\gamma_j - \mathbf{M}c\gamma_j) = c^2 \text{cov}(\gamma_i, \gamma_j)$$

Для доказательства второго равенства заметим, что $(C\gamma)_i = \sum_{j=1}^d c_{ij}\gamma_j$, поэтому в силу билинейности оператора ковариации

$$\begin{aligned} & \text{cov}((C\gamma + a)_r, (C\gamma + a)_s) = \\ &= \mathbf{M} \left(\sum_{j=1}^d c_{rj}\gamma_j - \mathbf{M} \sum_{j=1}^d c_{rj}\gamma_j \right) \left(\sum_{m=1}^d c_{sm}\gamma_m - \mathbf{M} \sum_{m=1}^d c_{sm}\gamma_m \right) = \\ &= \sum_{j,m=1}^d \mathbf{M}(c_{rj}\gamma_j - \mathbf{M}c_{rj}\gamma_j)(c_{sm}\gamma_m - \mathbf{M}c_{sm}\gamma_m) = \\ &= \sum_{j,m=1}^d c_{rj} \text{cov}(\gamma_j, \gamma_m) c_{sm} = (CB(\gamma)C^T)_{rs} \end{aligned}$$

что и требовалось доказать.

Способы определения ковариации (??)

(в 1ю продублирую)

Если задана таблица условных распределений, то считаем мат ожидания, а потом вставляем в формулу

3.6.9 Typical Inequalities in Probability Theory

Приведем наиболее используемые неравенства.

(тут тренироваться нужно много!!!)

Неравенство Маркова

.

Конструкция

Для любой случайной величины и любого $x > 0$

$$\mathbf{P}\{|\xi| \geq x\} \leq \frac{\mathbf{M}|\xi|}{x}$$

О случаях применений (??)**Доказательство**

Доказательство.

Рассмотрим разбиение Ω на события $\{|\xi| \geq x\}$ и $\{|\xi| < x\}$. Тогда

$$\begin{aligned} \mathbf{M}|\xi| &= \mathbf{M}|\xi|(I\{|\xi| \geq x\} + I\{|\xi| < x\}) = \\ &= \mathbf{M}|\xi|I\{|\xi| \geq x\} + \mathbf{M}|\xi|I\{|\xi| < x\} \geq \\ &\geq \mathbf{M}xI\{|\xi| \geq x\} + \mathbf{M}0 \cdot I\{|\xi| < x\} = x\mathbf{P}\{|\xi| \geq x\} \end{aligned}$$

Разумеется, неравенство Маркова содержательно только при $x > \mathbf{M}|\xi|$.

Неравенства Чебышёва

.

Конструкция

Для любой случайной величины ξ и для любого $x > 0$

$$\mathbf{P}\{|\xi - \mathbf{M}\xi| \geq x\} \leq \frac{\mathbf{D}\xi}{x^2}, \quad \mathbf{P}\{|\xi - \mathbf{M}\xi| < x\} \geq 1 - \frac{\mathbf{D}\xi}{x^2}$$

О случаях применений (??)

Доказательство

Доказательство.

Первое неравенство - следствие неравенства Маркова, примененного к квадрату центрированной случайной величины: $\eta = (\xi - \mathbf{M}\xi)^2$.

Второе неравенство эквивалентно первому, так как

$$\mathbf{P}\{|\xi - \mathbf{M}\xi| < x\} = 1 - \mathbf{P}\{|\xi - \mathbf{M}\xi| \geq x\}$$

Содержательный смысл неравенства Чебышёва в том, что основная масса распределения случайной величины ξ с конечной дисперсией $\mathbf{D}\xi$ сосредоточена в интервале, длина которого имеет тот же порядок, что $\sqrt{\mathbf{D}\xi}$.

В частности, если $\mathbf{D}\xi$ мало, то почти все распределение ξ сосредоточено около $\mathbf{M}\xi$.

Неравенство Коши-Буняковского

.

Конструкция

$$\mathbf{M}|\xi\eta| \leq \sqrt{\mathbf{M}\xi^2\mathbf{M}\eta^2}$$

О случаях применений (??)

Доказательство

(и там у зубкова еще много неравенств, но пока что пофиг на них, мне до этого темы доучивать нужно)

Закон больших чисел в двух словах

????

(выпишу сюда суть ее из подробной теории вероятностей)

Суть

(потом её напишу)

3.6.10 About Other Structures

(потом мб сделаю указания на другие конструкции, которые не так актуальны, как то, что выше)

4 Introduction to Probability by Blitzstein, Hwang

4.0.1 Preface

This book provides a modern introduction to probability and develops a foundation for understanding statistics, randomness, and uncertainty. A variety of applications and examples are explored, from basic coin-tossing and the study of coincidences to Google PageRank and Markov chain Monte Carlo. As probability is often considered to be a counterintuitive subject, many intuitive explanations, diagrams, and practice problems are given. Each chapter ends with a section showing how to explore the ideas of that chapter in R , a free software environment for statistical calculations and simulations.

Lecture videos from Stat 110 at Harvard, the course which gave rise to this book, are freely available at <http://stat110.net>. Additional supplementary materials such as R code, animations, and solutions to the exercises marked with (S), are also available at this site.

Calculus is a prerequisite for this book; there is no statistics prerequisite. The main mathematical challenge lies not in performing technical calculus derivations, but in translating between abstract concepts and concrete examples. Some major themes and features are listed below.

1. Stories. Throughout this book, definitions, theorems, and proofs are presented through stories: real-world interpretations that preserve mathematical precision and generality. We explore probability distributions using the generative stories that make them widely used in statistical modeling. When possible, we refrain from tedious derivations and instead aim to give interpretations and intuitions for why key results are true. Our experience is that this approach promotes long-term retention of the material by providing insight instead of demanding rote memorization.
2. Pictures. Since pictures are thousand-word stories, we supplement definitions with illustrations so that key concepts are associated with memorable diagrams. In many fields, the difference between a novice and an expert has been described as follows: the novice struggles to memorize a large number of seemingly disconnected facts and formulas, whereas the expert sees a unified structure in which a few principles and ideas connect these facts coherently. To help students see the structure of probability, we emphasize the connections between ideas (both verbally and visually), and at the end of most chapters we present recurring, ever-expanding maps of concepts and distributions.
3. Dual teaching of concepts and strategies. Our intent is that in reading this book, students will learn not only the concepts of probability, but also a set of problem-solving strategies that are widely applicable outside of probability. In the worked examples, we explain each step of the solution but also comment on how we knew to take the approach we did. Often we present multiple solutions to the same problem.

We explicitly identify and name important strategies such as symmetry and pattern recognition, and we proactively dispel common misunderstandings, which are marked with the (bio-hazard) symbol.

4. Practice problems. The book contains about 600 exercises of varying difficulty. The exercises are intended to reinforce understanding of the material and strengthen problem-solving skills instead of requiring repetitive calculations. Some are strategic practice problems, grouped by theme to facilitate practice of a particular topic, while others are mixed practice, in which several earlier topics may need to be synthesized. About 250 exercises have detailed online solutions for practice and self-study.

5. Simulation, Monte Carlo, and R. Many probability problems are too difficult to solve exactly, and in any case it is important to be able to check one's answer. We introduce techniques for exploring probability via simulation, and show that often a few lines of R code suffice to create a simulation for a seemingly complicated problem.

6. Focus on real-world relevance and statistical thinking. Examples and exercises in this book have a clear real-world motivation, with a particular focus on building a strong foundation for further study of statistical inference and modeling. We preview important statistical ideas such as sampling, simulation, Bayesian inference, and Markov chain Monte Carlo; other application areas include genetics, medicine, computer science, and information theory. Our choice of examples and exercises is intended to highlight the power, applicability, and beauty of probabilistic thinking.

The second edition benefits from hundreds of comments, questions, and reviews from students who took courses using the book, faculty who taught with the book, and readers using the book for self-study. We have added many new examples, exercises, and explanations based on our experience teaching with the book and the feedback we have received.

New supplementary materials have also been added at <http://stat110.net>, including animations and interactive visualizations that were created in connection with the edX online version of Stat 110. These are intended to help make probability feel more intuitive, visual, and tangible.

4.1 Probability and Counting

Luck. Coincidence. Randomness. Uncertainty. Risk. Doubt. Fortune. Chance. You've probably heard these words countless times, but chances are that they were used in a vague, casual way. Unfortunately, despite its ubiquity in science and everyday life, probability can be deeply counterintuitive. If we rely on intuitions of doubtful validity, we run a serious risk of making inaccurate predictions or overconfident decisions. The goal of this book is to introduce probability as a logical framework for quantifying uncertainty and randomness in a principled way. We'll also aim to strengthen intuition, both when our initial guesses coincide with logical reasoning and when we're not so lucky.

4.1.1 1.1 Why Study Probability?

Mathematics is the logic of certainty; probability is the logic of uncertainty. Probability is extremely useful in a wide variety of fields, since it provides tools for understanding and explaining variation, separating signal from noise, and modeling complex phenomena. To give just a small sample from a continually growing list of applications:

1. Statistics: Probability is the foundation and language for statistics, enabling many powerful methods for using data to learn about the world.

2. Physics: Einstein famously said "God does not play dice with the universe", but current understanding of quantum physics heavily involves probability at the most fundamental level of nature. Statistical mechanics is another major branch of physics that is built on probability.
3. Biology: Genetics is deeply intertwined with probability, both in the inheritance of genes and in modeling random mutations.
4. Computer science: Randomized algorithms make random choices while they are run, and in many important applications they are simpler and more efficient than any currently known deterministic alternatives. Probability also plays an essential role in studying the performance of algorithms, and in machine learning and artificial intelligence.
5. Meteorology: Weather forecasts are (or should be) computed and expressed in terms of probability.
6. Gambling: Many of the earliest investigations of probability were aimed at answering questions about gambling and games of chance.
7. Finance: At the risk of redundancy with the previous example, it should be pointed out that probability is central in quantitative finance. Modeling stock prices over time and determining "fair" prices for financial instruments are based heavily on probability.
8. Political science: In recent years, political science has become more and more quantitative and statistical, with applications such as analyzing surveys of public opinion, assessing gerrymandering, and predicting elections.
9. Medicine: The development of randomized clinical trials, in which patients are randomly assigned to receive treatment or placebo, has transformed medical research in recent years. As the biostatistician David Harrington remarked, "Some have conjectured that it could be the most significant advance in scientific medicine in the twentieth century....In one of the delightful ironies of modern science, the randomized trial 'adjusts' for both observed and unobserved heterogeneity in a controlled experiment by introducing chance variation into the study design." [16]
10. Life: Life is uncertain, and probability is the logic of uncertainty. While it isn't practical to carry out a formal probability calculation for every decision made in life, thinking hard about probability can help us avert some common fallacies, shed light on coincidences, and make better predictions.

Probability provides procedures for principled problem-solving, but it can also produce pitfalls and paradoxes. For example, we'll see in this chapter that even Gottfried Wilhelm von Leibniz and Sir Isaac Newton, the two people who independently discovered calculus in the 17th century, were not immune to basic errors in probability. Throughout this book, we will use the following strategies to help avoid potential pitfalls.

1. Simulation: A beautiful aspect of probability is that it is often possible to study problems via simulation. Rather than endlessly debating an answer with someone who disagrees with you, you can run a simulation and see empirically who is right. Each chapter in this book ends with a section that gives examples of how to do calculations and simulations in R, a free statistical computing environment.

2. Biohazards: Studying common mistakes is important for gaining a stronger understanding of what is and is not valid reasoning in probability. In this book, common mistakes are called biohazards and are denoted by (since making such mistakes can be hazardous to one's health!).
3. Sanity checks: After solving a problem one way, we will often try to solve the same problem in a different way or to examine whether our answer makes sense in simple and extreme cases.

4.1.2 1.2 Sample Spaces and Pebble World

The mathematical framework for probability is built around sets. Imagine that an experiment is performed, resulting in one out of a set of possible outcomes. Before the experiment is performed, it is unknown which outcome will be the result; after, the result "crystallizes" into the actual outcome.

Definition 1.2.1 (Sample space and event). The sample space S of an experiment is the set of all possible outcomes of the experiment. An event A is a subset of the sample space S , and we say that A occurred if the actual outcome is in A .




FIGURE 1.1

A sample space as Pebble World, with two events A and B spotlighted.

The sample space of an experiment can be finite, countably infinite, or uncountably infinite (see Section A.1.5 of the math appendix for an explanation of countable and uncountable sets). When the sample space is finite, we can visualize it as Pebble World, as shown in Figure 1.1. Each pebble represents an outcome, and an event is a set of pebbles.

Performing the experiment amounts to randomly selecting one pebble. If all the pebbles are of the same mass, all the pebbles are equally likely to be chosen. This special case is the topic of the next two sections. In Section 1.6, we give a general definition of probability that allows the pebbles to differ in mass.

Set theory is very useful in probability, since it provides a rich language for expressing and working with events; Section A. 1 of the math appendix provides a review of set theory. Set operations, especially unions, intersections, and complements, make it easy to build new events in terms of already-defined events. These concepts also let us express an event in more than one way; often, one expression for an event is much easier to work with than another expression for the same event.

For example, let S be the sample space of an experiment and let $A, B \subseteq S$ be events. Then the union $A \cup B$ is the event that occurs if and only if at least one of A, B occurs, the intersection $A \cap B$ is the event that occurs if and only if both A and B occur, and the complement A^c is the event that occurs if and only if A does not occur. We also have De Morgan's laws:

$$(A \cup B)^c = A^c \cap B^c \text{ and } (A \cap B)^c = A^c \cup B^c,$$

since saying that it is not the case that at least one of A and B occur is the same as saying that A does not occur and B does not occur, and saying that it is not the case that both occur is the same as saying that at least one does not occur. Analogous results hold for unions and intersections of more than two events.

In the example shown in Figure 1.1, A is a set of 5 pebbles, B is a set of 4 pebbles, $A \cup B$ consists of the 8 pebbles in A or B (including the pebble that is in both), $A \cap B$ consists of the pebble that is in both A and B , and A^c consists of the 4 pebbles that are not in A .

The notion of sample space is very general and abstract, so it is important to have some concrete examples in mind.

Example 1.2.2 (Coin flips). A coin is flipped 10 times. Writing Heads as H and Tails as T , a possible outcome (pebble) is $HHHTHHTTHT$, and the sample space is the set of all possible strings of length 10 of H 's and T 's. We can (and will) encode H as 1 and T as 0, so that an outcome is a sequence $(s_1, \dots, s_{10})$ with $s_j \in \{0, 1\}$, and the sample space is the set of all such sequences. Now let's look at some events:

1. Let A_1 be the event that the first flip is Heads. As a set,

$$A_1 = \{(1, s_2, \dots, s_{10}) : s_j \in \{0, 1\} \text{ for } 2 \leq j \leq 10\}.$$

This is a subset of the sample space, so it is indeed an event; saying that A_1 occurs is the same thing as saying that the first flip is Heads. Similarly, let A_j be the event that the j th flip is Heads for $j = 2, 3, \dots, 10$.

2. Let B be the event that at least one flip was Heads. As a set,

$$B = \bigcup_{j=1}^{10} A_j$$

3. Let C be the event that all the flips were Heads. As a set,

$$C = \bigcap_{j=1}^{10} A_j.$$

4. Let D be the event that there were at least two consecutive Heads. As a set,

$$D = \bigcup_{j=1}^9 (A_j \cap A_{j+1}).$$

Example 1.2.3 (Pick a card, any card). Pick a card from a standard deck of 52 cards. The sample space S is the set of all 52 cards (so there are 52 pebbles, one for each card). Consider the following four events:

- A : card is an ace.

- B : card has a black suit.
- D : card is a diamond.
- H : card is a heart.

As a set, H consists of 13 cards:

{Ace of Hearts, Two of Hearts, . . . , King of Hearts }.

We can create various other events in terms of the events A, B, D , and H . Unions, intersections, and complements are especially useful for this. For example:

- $A \cap H$ is the event that the card is the Ace of Hearts.
- $A \cap B$ is the event {Ace of Spades, Ace of Clubs}.
- $A \cup D \cup H$ is the event that the card is red or an ace.
- $(A \cup B)^c = A^c \cap B^c$ is the event that the card is a red non-ace.

Also, note that $(D \cup H)^c = D^c \cap H^c = B$, so B can be expressed in terms of D and H . On the other hand, the event that the card is a spade can't be written in terms of A, B, D, H since none of them are fine-grained enough to be able to distinguish between spades and clubs.

There are many other events that could be defined using this sample space. In fact, the counting methods introduced later in this chapter show that there are $2^{52} \approx 4.5 \times 10^{15}$ events in this problem, even though there are only 52 pebbles.

What if the card drawn were a joker? That would indicate that we had the wrong sample space; we are assuming that the outcome of the experiment is guaranteed to be an element of S .

As the preceding examples demonstrate, events can be described in English or in set notation. Sometimes the English description is easier to interpret while the set notation is easier to manipulate. Let S be a sample space and s_{actual} be the actual outcome of the experiment (the pebble that ends up getting chosen when the experiment is performed). A mini-dictionary for converting between English and sets is given on the next page.

English	Sets
Events and occurrences	
sample space	S
s is a possible outcome	$s \in S$
A is an event	$A \subseteq S$
A occurred	$s_{\text{actual}} \in A$
something must happen	$s_{\text{actual}} \in S$
New events from old events	
A or B (inclusive)	$A \cup B$
A and B not A	$A \cap B^c$
A or B , but not both at least one of $A_1, \dots, A_n$	$(A \cap B^c) \cup (A^c \cap B)$
	$A_1 \cup \dots \cup A_n$
all of $A_1, \dots, A_n$	$A_1 \cap \dots \cap A_n$
Relationships between events	
A implies B	$A \subseteq B$
A and B are mutually exclusive	$A \cap B = \emptyset$
$A_1, \dots, A_n$ are a partition of S	$A_1 \cup \dots \cup A_n = S, A_i \cap A_j = \emptyset$ for $i \neq j$

4.1.3 1.3 Naive Definition of Probability

Historically, the earliest definition of the probability of an event was to count the number of ways the event could happen and divide by the total number of possible outcomes for the experiment. We call this the naive definition since it is restrictive and relies on strong assumptions; nevertheless, it is important to understand, and useful when not misused.

Definition 1.3.1 (Naive definition of probability). Let A be an event for an experiment with a finite sample space S . The naive probability of A is

$$P_{\text{naive}}(A) = \frac{|A|}{|S|} = \frac{\text{number of outcomes favorable to } A}{\text{total number of outcomes in } S}.$$

(We use $|A|$ to denote the size of A ; see Section A.1.5 of the math appendix.)

In terms of Pebble World, the naive definition just says that the probability of A is the fraction of pebbles that are in A . For example, in Figure 1.1 it says

$$P_{\text{naive}}(A) = \frac{5}{9}, P_{\text{naive}}(B) = \frac{4}{9}, P_{\text{naive}}(A \cup B) = \frac{8}{9}, P_{\text{naive}}(A \cap B) = \frac{1}{9}.$$

For the complements of the events just considered,

$$P_{\text{naive}}(A^c) = \frac{4}{9}, P_{\text{naive}}(B^c) = \frac{5}{9}, P_{\text{naive}}((A \cup B)^c) = \frac{1}{9}, P_{\text{naive}}((A \cap B)^c) = \frac{8}{9}.$$

In general,

$$P_{\text{naive}}(A^c) = \frac{|A^c|}{|S|} = \frac{|S| - |A|}{|S|} = 1 - \frac{|A|}{|S|} = 1 - P_{\text{naive}}(A).$$

In Section 1.6, we will see that this result about complements always holds for probability, even when we go beyond the naive definition. A good strategy when trying to find the probability of an event is to start by thinking about whether it will be easier to find the probability of the event or the probability of its complement. De Morgan's laws are especially useful in this context, since it may be easier to work with an intersection than a union, or vice versa.

The naive definition is very restrictive in that it requires S to be finite, with equal mass for each pebble. It has often been misapplied by people who assume equally likely outcomes without justification and make arguments to the effect of "either it will happen or it won't, and we don't know which, so it's 50 – 50". In addition to sometimes giving absurd probabilities, this type of reasoning isn't even internally consistent. For example, it would say that the probability of life on Mars is 1/2 ("either there is or there isn't life there"), but it would also say that the probability of intelligent life on Mars is 1/2, and it is clear intuitively - and by the properties of probability developed in Section 1.6 - that the latter should have strictly lower probability than the former. But there are several important types of problems where the naive definition is applicable:

- when there is symmetry in the problem that makes outcomes equally likely. It is common to assume that a coin has a 50% chance of landing Heads when tossed, due to the physical symmetry of the coin.¹ For a standard, well-shuffled deck of cards, it is reasonable to assume that all orders are equally likely. There aren't certain overeager cards that especially like to be near the top of the deck; any particular location in the deck is equally likely to house any of the 52 cards.
- when the outcomes are equally likely by design. For example, consider conducting a survey of n people in a population of N people. A common goal is to obtain a simple random sample, which means that the n people are chosen randomly with all subsets of

size n being equally likely. If successful, this ensures that the naive definition is applicable, but in practice this may be hard to accomplish because of various complications, such as not having a complete, accurate list of contact information for everyone in the population.

- when the naive definition serves as a useful null model. In this setting, we assume that the naive definition applies just to see what predictions it would yield, and then we can compare observed data with predicted values to assess whether the hypothesis of equally likely outcomes is tenable.

4.1.4 1.4 How to Count

Calculating the naive probability of an event A involves counting the number of pebbles in A and the number of pebbles in the sample space S . Often the sets we need to count are extremely large. This section introduces some fundamental methods for counting; further methods can be found in books on combinatorics, the branch of mathematics that studies counting.

1.4.1 Multiplication Rule

In some problems, we can directly count the number of possibilities using a basic but versatile principle called the multiplication rule. We'll see that the multiplication rule leads naturally to counting rules for sampling with replacement and sampling without replacement, two scenarios that often arise in probability and statistics.

Theorem 1.4.1 (Multiplication rule). Consider a compound experiment consisting of two sub-experiments, Experiment A and Experiment B. Suppose that Experiment A has a possible outcomes, and for each of those outcomes Experiment B has b possible outcomes. Then the compound experiment has ab possible outcomes.

To see why the multiplication rule is true, imagine a tree diagram as in Figure 1.2. Let the tree branch a ways according to the possibilities for Experiment A, and for each of those branches create b further branches for Experiment B. Overall, there are $\underbrace{b + b + \cdots + b}_a = ab$ possibilities.

1.4.2. It is often easier to think about the experiments as being in chronological order, but there is no requirement in the multiplication rule that Experiment A has to be performed before Experiment B.

Example 1.4.3 (Runners). Suppose that 10 people are running a race. Assume that ties are not possible and that all 10 will complete the race, so there will be welldefined first place, second place, and third place winners. How many possibilities are there for the first, second, and third place winners?

Solution: There are 10 possibilities for who gets first place, then once that is fixed there are 9 possibilities for who gets second place, and once these are both fixed

⁰¹ See Diaconis, Holmes, and Montgomery [7] for a physical argument that the chance of a tossed coin coming up the way it started is about 0.51 (close to but slightly more than $1/2$), and Gelman and Nolan [11] for an explanation of why the probability of Heads is close to $1/2$ even for a coin that is manufactured to have different weights on the two sides (for standard coin-tossing; allowing the coin to spin is a different matter).




FIGURE 1.2

Tree diagram illustrating the multiplication rule. If Experiment A has 3 possible outcomes, for each of which Experiment B has 4 possible outcomes, then overall there are $3 \cdot 4 = 12$ possible outcomes.

there are 8 possibilities for third place. So by the multiplication rule, there are $10 \cdot 9 \cdot 8 = 720$ possibilities.

We did not have to consider the first place winner first. We could just as well have said that there are 10 possibilities for who got third place, then once that is fixed there are 9 possibilities for second place, and once those are both fixed there are 8 possibilities for first place. Or imagine that there are 3 platforms, which the first, second, and third place runners will stand on after the race. The platforms are gold, silver, and bronze, allocated to the first, second, and third place runners, respectively. Again there are $10 \cdot 9 \cdot 8 = 720$ possibilities for how the platforms will be occupied after the race, and there is no reason that the platforms must be considered in the order (gold, silver, bronze). $\square$

Example 1.4.4 (Chessboard). How many squares are there in an 8×8 chessboard, as in Figure 1.3? Even the name " 8×8 chessboard" makes this easy: there are $8 \cdot 8 = 64$ squares

on the board. The grid structure makes this clear, but we can also think of this as an example of the multiplication rule: to specify a square, we can specify which row and which column it is in. There are 8 choices of row, for each of which there are 8 choices of column.

Furthermore, we can see without doing any calculations that half the squares are white and half are black. Imagine rotating the chessboard 90 degrees clockwise. Then all the positions that had a white square now contain a black square, and vice versa, so the number of white squares must equal the number of black squares. We




FIGURE 1.3

An 8×8 chessboard (left) and a crossword puzzle grid (right). The chessboard has $8 \cdot 8 = 64$ squares, whereas counting the number of white squares in the crossword puzzle grid requires more work.

can also count the number of white squares using the multiplication rule: in each of the 8 rows there are 4 white squares, giving a total of $8 \cdot 4 = 32$ white squares.

In contrast, it would require more effort to count the number of white squares in the crossword puzzle grid shown in Figure 1.3. The multiplication rule does not apply, since different rows sometimes have different numbers of white squares. $\square$

Example 1.4.5 (Ice cream cones). Suppose you are buying an ice cream cone. You can choose whether to have a cake cone or a waffle cone, and whether to have chocolate, vanilla, or strawberry as your flavor. This decision process can be visualized with a tree diagram, as in Figure 1.4.




FIGURE 1.4




FIGURE 1.4

Tree diagram for choosing an ice cream cone. Regardless of whether the type of cone or the flavor is chosen first, there are $2 \cdot 3 = 3 \cdot 2 = 6$ possibilities.

By the multiplication rule, there are $2 \cdot 3 = 6$ possibilities. This is a very simple example,

but is worth thinking through in detail as a foundation for thinking about and visualizing more complicated examples. Soon we will encounter examples where drawing the tree in a legible size would take up more space than exists in the known universe, yet where conceptually we can still think in terms of the ice cream example. Some things to note:

1. It doesn't matter whether you choose the type of cone first ("I'd like a waffle cone with chocolate ice cream") or the flavor first ("I'd like chocolate ice cream on a waffle cone"). Either way, there are $2 \cdot 3 = 3 \cdot 2 = 6$ possibilities.
2. It doesn't matter whether the same flavors are available on a cake cone as on a waffle cone. What matters is that there are exactly 3 flavor choices for each cone choice. If for some strange reason it were forbidden to have chocolate ice cream on a waffle cone, with no substitute flavor available (aside from vanilla and strawberry), there would be $3 + 2 = 5$ possibilities and the multiplication rule wouldn't apply. In larger examples, such complications could make counting the number of possibilities vastly more difficult.

Now suppose you buy two ice cream cones on a certain day, one in the afternoon and the other in the evening. Write, for example, (cakeC, waffleV) to mean a cake cone with chocolate in the afternoon, followed by a waffle cone with vanilla in the evening. By the multiplication rule, there are $6^2 = 36$ possibilities in your delicious compound experiment.

But what if you're only interested in what kinds of ice cream cones you had that day, not the order in which you had them, so you don't want to distinguish, for example, between (cakeC, waffleV) and (waffleV, cakeC)? Are there now $36/2 = 18$ possibilities? No, since possibilities like (cakeC, cakeC) were already only listed once each. There are $6 \cdot 5 = 30$ ordered possibilities (x, y) with $x \neq y$, which turn into 15 possibilities if we treat (x, y) as equivalent to (y, x) , plus 6 possibilities of the form (x, x) , giving a total of 21 possibilities. Note that if the 36 original ordered pairs (x, y) are equally likely, then the 21 possibilities here are not equally likely.

Example 1.4.6 (Subsets). A set with n elements has 2^n subsets, including the empty set $\emptyset$ and the set itself. This follows from the multiplication rule since for each element, we can choose whether to include it or exclude it. For example, the set $\{1, 2, 3\}$ has the 8 subsets $\emptyset, \{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}$. This result explains why in Example 1.2.3 there are 2^{52} events that can be defined.

We can use the multiplication rule to arrive at formulas for sampling with and without replacement. Many experiments in probability and statistics can be interpreted in one of these two contexts, so it is appealing that both formulas follow directly from the same basic counting principle.

Theorem 1.4.7 (Sampling with replacement). Consider n objects and making k choices from them, one at a time with replacement (i.e., choosing a certain object does not preclude it from being chosen again). Then there are n^k possible outcomes (where order matters, in the sense that, e.g., choosing object 3 and then object 7 is counted as a different outcome than choosing object 7 and then object 3.)

For example, imagine a jar with n balls, labeled from 1 to n . We sample balls one at a time with replacement, meaning that each time a ball is chosen, it is returned to the jar. Each sampled ball is a sub-experiment with n possible outcomes, and there are k sub-experiments. Thus, by the multiplication rule there are n^k ways to obtain a sample of size k .

Theorem 1.4.8 (Sampling without replacement). Consider n objects and making k choices from them, one at a time without replacement (i.e., choosing a certain object precludes it from being chosen again). Then there are $n(n-1) \cdots (n-k+1)$ possible outcomes for $1 \leq k \leq n$, and 0 possibilities for $k > n$ (where order matters). By convention, $n(n-1) \cdots (n-k+1) = n$ for $k = 1$.

This result also follows directly from the multiplication rule: each sampled ball is again a sub-experiment, and the number of possible outcomes decreases by 1 each time. Note that for sampling k out of n objects without replacement, we need $k \leq n$, whereas in sampling with replacement the objects are inexhaustible.

Example 1.4.9 (Permutations and factorials). A permutation of $1, 2, \dots, n$ is an arrangement of them in some order, e.g., $3, 5, 1, 2, 4$ is a permutation of $1, 2, 3, 4, 5$. By Theorem 1.4.8 with $k = n$, there are $n!$ permutations of $1, 2, \dots, n$. For example, there are $n!$ ways in which n people can line up for ice cream. (Recall that $n!$ is $n(n-1)(n-2) \cdots 1$ for any positive integer n , and $0! = 1$.)

Theorems 1.4.7 and 1.4.8 are theorems about counting, but when the naive definition applies, we can use them to calculate probabilities. This brings us to our next example, a famous problem in probability called the birthday problem. The solution incorporates both sampling with replacement and sampling without replacement.

Example 1.4.10 (Birthday problem). There are k people in a room. Assume each person's birthday is equally likely to be any of the 365 days of the year (we exclude February 29), and that people's birthdays are independent (we will define independence formally later, but intuitively it means that knowing some people's birthdays gives us no information about other people's birthdays; this would not hold if, e.g., we knew that two of the people were twins). What is the probability that at least one pair of people in the group have the same birthday?

Solution:

There are 365^k ways to assign birthdays to the people in the room, since we can imagine the 365 days of the year being sampled k times, with replacement. By assumption, all of these possibilities are equally likely, so the naive definition of probability applies.

Used directly, the naive definition says we just need to count the number of ways to assign birthdays to k people such that there are two people who share a birthday.

But this counting problem is hard, since it could be Emma and Steve who share a birthday, or Steve and Naomi, or all three of them, or the three of them could share a birthday while two others in the group share a different birthday, or various other possibilities.

Instead, let's count the complement: the number of ways to assign birthdays to k people such that no two people share a birthday. This amounts to sampling the 365 days of the year without replacement, so the number of possibilities is $365 \cdot 364 \cdot 363 \cdots (365 - k + 1)$ for $k \leq 365$. Therefore the probability of no birthday matches in a group of k people is

$$P(\text{ no birthday match }) = \frac{365 \cdot 364 \cdots (365 - k + 1)}{365^k},$$

and the probability of at least one birthday match is

$$P(\text{ at least 1 birthday match }) = 1 - \frac{365 \cdot 364 \cdots (365 - k + 1)}{365^k}.$$

Figure 1.5 plots the probability of at least one birthday match as a function of k . The first value of k for which the probability of a match exceeds 0.5 is $k = 23$. Thus, in a group of 23 people, there is a better than 50% chance that there is at least one birthday match. At $k = 57$, the probability of a match already exceeds 99%.




FIGURE 1.5 Probability that in a room of k people, at least two were born on the same day. This probability first exceeds 0.5 when $k = 23$.

Of course, for $k = 366$ we are guaranteed to have a match, but it's surprising that even with a much smaller number of people it's overwhelmingly likely that there is a birthday match. For a quick intuition into why it should not be so surprising, note that with 23 people there are $\binom{23}{2} = 253$ pairs of people, any of which could be a birthday match.

Problems 26 and 27 show that the birthday problem is much more than a fun party game, and much more than a way to build intuition about coincidences; there are also important applications in statistics and computer science. Problem 62 explores the more general setting in which the probability is not necessarily $1/365$ for each day. It turns out that in the non-equal probability case, having at least one match becomes even more likely.

1.4.11 (Labeling objects). Drawing a sample from a population is a very fundamental concept in statistics. It is important to think of the objects or people in the population as named or labeled. For example, if there are n balls in a jar, we can imagine that they have labels from 1 to n , even if the balls look the same to the human eye. In the birthday problem, we can give each person an ID (identification) number, rather than thinking of the people as indistinguishable particles or a faceless mob.

A related example is an instructive blunder made by Leibniz in a seemingly simple problem (see Gorroochurn [14] for discussion of this and a variety of other probability problems from a historical perspective).

Example 1.4.12 (Leibniz's mistake). If we roll two fair dice, which is more likely: a sum of 11 or a sum of 12?

Solution:

Label the dice A and B, and consider each die to be a sub-experiment. By the multiplication rule, there are 36 possible outcomes for ordered pairs of the form (value of A, value of B), and they are equally likely by symmetry. Of these, (5, 6) and (6, 5) are favorable to a sum of 11, while only (6, 6) is favorable to a sum of 12. Therefore a sum of 11 is twice as likely as a sum of 12; the probability is $1/18$ for the former, and $1/36$ for the latter.

However, Leibniz wrongly argued that a sum of 11 and a sum of 12 are equally likely, claiming that each of these sums can be attained in only one way. Here Leibniz was making the mistake of treating the two dice as indistinguishable objects, viewing (5, 6) and (6, 5) as the same outcome.

What are the antidotes to Leibniz's mistake? First, as explained in 1.4.11, we should label the objects in question instead of treating them as indistinguishable. If Leibniz had labeled his dice A and B, or green and orange, or left and right, he would not have made this mistake. Second, before we use counting for probability, we should ask ourselves whether the naive definition applies (see 1.4.23 for another example showing that caution is needed before applying the naive definition).

1.4.2 Adjusting for Overcounting

In many counting problems, it is not easy to directly count each possibility once and only once. If, however, we are able to count each possibility exactly c times for some c , then we can adjust by dividing by c . For example, if we have exactly double-counted each possibility, we can divide by 2 to get the correct count. We call this adjusting for overcounting.

Example 1.4.13 (Committees and teams). Consider a group of four people.

- How many ways are there to choose a two-person committee?
- How many ways are there to break the people into two teams of two?

Solution:

(a) One way to count the possibilities is by listing them out: labeling the people as 1, 2, 3, 4, the possibilities are 12, 13, 14, 23, 24, 34.

Another approach is to use the multiplication rule with an adjustment for overcounting. By the multiplication rule, there are 4 ways to choose the first person on the committee and 3 ways to choose the second person on the committee, but this counts each possibility twice, since picking 1 and 2 to be on the committee is the same as picking 2 and 1 to be on the committee. Since we have overcounted by a factor of 2, the number of possibilities is $(4 \cdot 3)/2 = 6$.

(b) Here are 3 ways to see that there are 3 ways to form the teams. Labeling the people as 1, 2, 3, 4, we can directly list out the possibilities:

1	2	3	4
1	2	3	4
1	2	3	4

\hline 12 & 34

\hline
13/24, and
\hline 14 & 23

\hline
\end{tabular} . Listing out all possibilities would quickly become tedious or infeasible with more people though. Another approach is to note that it suffices to specify person 1's teammate (and then the other team is determined). A third way is to use (a) to see that there are 6 ways to choose one team. This overcounts by a factor of 2, since picking 1 and 2 to be a team is equivalent to picking 3 and 4 to be a team. So again the answer is $6/2 = 3$. $\square$

A binomial coefficient counts the number of subsets of a certain size for a set, such as the number of ways to choose a committee of size k from a set of n people. Sets and subsets are by definition unordered, e.g., $\{3, 1, 4\} = \{4, 1, 3\}$, so we are counting the number of ways to choose k objects out of n , without replacement and without distinguishing between the different orders in which they could be chosen.

Definition 1.4.14 (Binomial coefficient). For any nonnegative integers k and n , the binomial coefficient $\binom{n}{k}$, read as " n choose k ", is the number of subsets of size k for a set of size n .

For example, $\binom{4}{2} = 6$, as shown in Example 1.4.13. The binomial coefficient $\binom{n}{k}$ is sometimes called a combination, but we do not use that terminology here since "combination" is such a useful general-purpose word. Algebraically, binomial coefficients can be computed as follows.

Theorem 1.4.15 (Binomial coefficient formula). For $k \leq n$, we have

$$\binom{n}{k} = \frac{n(n-1) \cdots (n-k+1)}{k!} = \frac{n!}{(n-k)!k!}.$$

For $k > n$, we have $\binom{n}{k} = 0$.

Proof. Let A be a set with $|A| = n$. Any subset of A has size at most n , so $\binom{n}{k} = 0$ for $k > n$. Now let $k \leq n$. By Theorem 1.4.8, there are $n(n-1) \cdots (n-k+1)$ ways to make an ordered

choice of k elements without replacement. This overcounts each subset of interest by a factor of $k!$ (since we don't care how these elements are ordered), so we can get the correct count by dividing by $k!$. $\square$

1.4.16. The binomial coefficient $\binom{n}{k}$ is often defined in terms of factorials, but keep in mind that $\binom{n}{k}$ is 0 if $k > n$, even though the factorial of a negative number is undefined. Also, the middle expression in Theorem 1.4.15 is often better for computation than the expression with factorials, since factorials grow extremely fast. For example,

$$\binom{100}{2} = \frac{100 \cdot 99}{2} = 4950$$

can even be done by hand, whereas computing $\binom{100}{2} = 100!/(98! \cdot 2!)$ by first calculating $100!$ and $98!$ would be wasteful and possibly dangerous because of the extremely large numbers involved ($100! \approx 9.33 \times 10^{157}$).

Example 1.4.17 (Club officers). In a club with n people, there are $n(n-1)(n-2)$ ways to choose a president, vice president, and treasurer, and there are

$$\binom{n}{3} = \frac{n(n-1)(n-2)}{3!}$$

ways to choose 3 officers without predetermined titles. $\square$

Example 1.4.18 (Permutations of a word). How many ways are there to permute the letters in the word LALALAAA? To determine a permutation, we just need to choose where the 5 A's go (or, equivalently, just decide where the 3 L's go). So there are

$$\binom{8}{5} = \binom{8}{3} = \frac{8 \cdot 7 \cdot 6}{3!} = 56 \text{ permutations.}$$

How many ways are there to permute the letters in the word STATISTICS? Here are two approaches. We could choose where to put the S's, then where to put the T's (from the remaining positions), then where to put the I's, then where to put the A (and then the C is determined). Alternatively, we can start with $10!$ and then adjust for overcounting, dividing by $3!3!2!$ to account for the fact that the S's can be permuted among themselves in any way, and likewise for the T's and I's. This gives

$$\binom{10}{3} \binom{7}{3} \binom{4}{2} \binom{2}{1} = \frac{10!}{3!3!2!} = 50400 \text{ possibilities.}$$

$\square$

Example 1.4.19 (Binomial theorem). The binomial theorem states that

$$(x+y)^n = \sum_{k=0}^n \binom{n}{k} x^k y^{n-k},$$

for any nonnegative integer n . To prove the binomial theorem, expand out the product

$$\underbrace{(x+y)(x+y)\dots(x+y)}_{n \text{ factors}}.$$

Just as $(a+b)(c+d) = ac + ad + bc + bd$ is the sum of terms where we pick the a or the b from the first factor (but not both) and the c or the d from the second factor (but not both), the terms of $(x+y)^n$ are obtained by picking either the x or the y (but not both) from each factor. There are $\binom{n}{k}$ ways to choose exactly k of the x 's, and each such choice yields the term $x^k y^{n-k}$. The binomial theorem follows.

We can use binomial coefficients to calculate probabilities in many problems for which the naive definition applies.

Example 1.4.20 (Full house in poker). A 5-card hand is dealt from a standard, well-shuffled 52 -card deck. The hand is called a full house in poker if it consists of three cards of some rank and two cards of another rank, e.g., three 7's and two 10's (in any order). What is the probability of a full house?

Solution:

All of the $\binom{52}{5}$ possible hands are equally likely by symmetry, so the naive definition is applicable. To find the number of full house hands, use the multiplication rule (and imagine the tree). There are 13 choices for what rank we have three of; for concreteness, assume we have three 7's and focus on that branch of the tree. There are $\binom{4}{3}$ ways to choose which 7's we have. Then there are 12 choices for what rank we have two of, say 10 's for concreteness, and $\binom{4}{2}$ ways to choose two 10 's. Thus,

$$P(\text{ full house }) = \frac{13\binom{4}{3}12\binom{4}{2}}{\binom{52}{5}} = \frac{3744}{2598960} \approx 0.00144$$

The decimal approximation is more useful when playing poker, but the answer in terms of binomial coefficients is exact and self-annotating (seeing $\binom{52}{5}$ is a much bigger hint about its origin than seeing 2598960).

Example 1.4.21 (Newton-Pepys problem). Isaac Newton was consulted about the following problem by Samuel Pepys, who wanted the information for gambling purposes. Which of the following events has the highest probability?

A: At least one 6 appears when 6 fair dice are rolled.

B : At least two 6 's appear when 12 fair dice are rolled.

C : At least three 6 's appear when 18 fair dice are rolled.

Solution:

The three experiments have 6^6 , 6^{12} , and 6^{18} possible outcomes, respectively, and by symmetry the naive definition applies in all three experiments.

A : Instead of counting the number of ways to obtain at least one 6 , it is easier to count the number of ways to get no 6 's. Getting no 6 's is equivalent to sampling the numbers 1 through 5 with replacement 6 times, so 5^6 outcomes are favorable to A^c (and $6^6 - 5^6$ are favorable to A). Thus

$$P(A) = 1 - \frac{5^6}{6^6} \approx 0.67$$

B: Again we count the outcomes in B^c first. There are 5^{12} ways to get no 6 's in 12 die rolls. There are $\binom{12}{1}5^{11}$ ways to get exactly one 6 : we first choose which die lands 6 , then sample the numbers 1 through 5 with replacement for the other 11 dice. Adding these, we get the number of ways to fail to obtain at least two 6 's. Then

$$P(B) = 1 - \frac{5^{12} + \binom{12}{1}5^{11}}{6^{12}} \approx 0.62.$$

C : We count the outcomes in C^c , i.e., the number of ways to get zero, one, or two 6 's in 18 die rolls. There are 5^{18} ways to get no 6 's, $\binom{18}{1}5^{17}$ ways to get exactly one 6 , and $\binom{18}{2}5^{16}$ ways to get exactly two 6 's (choose which two dice will land 6 , then decide how the other 16 dice will land).

$$P(C) = 1 - \frac{5^{18} + \binom{18}{1}5^{17} + \binom{18}{2}5^{16}}{6^{18}} \approx 0.60.$$

Therefore A has the highest probability.

Newton arrived at the correct answer using similar calculations. Newton also provided Pepys with an intuitive argument for why A was the most likely of the three; however, his intuition was invalid. As explained in Stigler [24], using loaded dice could result in a different ordering of A, B, C , but Newton's intuitive argument did not depend on the dice being fair.

In this book, we care about counting not for its own sake, but because it sometimes helps us to find probabilities. Here is an example of a neat but treacherous counting problem; the solution is elegant, but it is rare that the result can be used with the naive definition of probability.

Example 1.4.22 (Bose-Einstein). How many ways are there to choose k times from a set of n objects with replacement, if order doesn't matter (we only care about how many times each object was chosen, not the order in which they were chosen)?

Solution:

When order does matter, the answer is n^k by the multiplication rule, but this problem is much harder. We will solve it by solving an isomorphic problem (the same problem in a different guise).

Let us find the number of ways to put k indistinguishable particles into n distinguishable boxes. That is, swapping the particles in any way is not considered a separate possibility: all that matters are the counts for how many particles are in each box. This scenario is known as a Bose-Einstein problem, since the physicists Satyendra Nath Bose and Albert Einstein studied related problems about indistinguishable particles in the 1920s, using their ideas to successfully predict the existence of a strange state of matter known as a Bose-Einstein condensate.

Any configuration can be encoded as a sequence of $|$'s and $\bullet$'s in a natural way, as illustrated in Figure 1.6.




FIGURE 1.6

Bose-Einstein encoding: putting $k = 7$ indistinguishable particles into $n = 4$ distinguishable boxes can be expressed as a sequence of $|$'s and $\bullet$'s, where $|$ denotes a wall and $\bullet$ denotes a particle.

To be valid, a sequence must start and end with a $|$, and have exactly $n - 1$ $|$'s and exactly k $\bullet$'s in between the starting and ending $|$'s; conversely, any such sequence is a valid encoding for some configuration of particles in boxes. Imagine that we have written down the starting and ending $|$'s, which represent the outer walls, and in between there are $n + k - 1$ fill-in-the-blank slots in between the outer walls. We need only choose where to put the k $\bullet$'s (since then where the $n + k - 1$ interior $|$'s go is completely determined). So the number of possibilities is $\binom{n+k-1}{k}$.

This counting method is sometimes called the stars and bars argument, where here we have dots in place of stars.

To relate this result back to the original question, we can let each box correspond to one of the n objects and use the particles as "check marks" to tally how many times each object is selected. For example, if a certain box contains exactly 3 particles, that means the object corresponding to that box was chosen exactly 3 times. The particles being indistinguishable corresponds to the fact that we don't care about the order in which the objects are chosen. Thus, the answer to the original question is also $\binom{n+k-1}{k}$.

Another isomorphic problem is to count the number of solutions $(x_1, \dots, x_n)$ to the equation $x_1 + x_2 + \dots + x_n = k$, where the x_i are nonnegative integers. This is equivalent since we can think of x_i as the number of particles in the i th box.

1.4.23. The Bose-Einstein result should not be used in the naive definition of probability except in very special circumstances. For example, consider a survey where a sample of size k is collected by choosing people from a population of size n

one at a time, with replacement and with equal probabilities. Then the n^k ordered samples are equally likely, making the naive definition applicable, but the $\binom{n+k-1}{k}$ unordered samples (where all that matters is how many times each person was sampled) are not equally likely.

As another example, with $n = 365$ days in a year and k people, how many possible unordered birthday lists are there? For example, for $k = 3$, we want to count lists like (May 1, March 31, April 11), where all permutations are considered equivalent. We can't do a simple adjustment for overcounting such as $n^k/3!$ since, e.g., there are 6 permutations of (May 1, March 31, April 11) but only 3 permutations of (March 31, March 31, April 11). By Bose-Einstein, the number of lists is $\binom{n+k-1}{k}$. But the ordered birthday lists are equally likely, not the unordered lists, so the Bose-Einstein value should not be used in calculating birthday probabilities.

4.1.5 1.5 Story Proofs

A story proof is a proof by interpretation. For counting problems, this often means counting the same thing in two different ways, rather than doing tedious algebra. A story proof often avoids messy calculations and goes further than an algebraic proof toward explaining why the result is true. The word "story" has several meanings, some more mathematical than others, but a story proof (in the sense in which we're using the term) is a fully valid mathematical proof. Here are some examples of story proofs, which also serve as further examples of counting.

Example 1.5.1 (Choosing the complement). For any nonnegative integers n and k with $k \leq n$, we have

$$\binom{n}{k} = \binom{n}{n-k}.$$

This is easy to check algebraically (by writing the binomial coefficients in terms of factorials), but a story proof makes the result easier to understand intuitively.

Story proof: Consider choosing a committee of size k in a group of n people. We know that there are $\binom{n}{k}$ possibilities. But another way to choose the committee is to specify which $n - k$ people are not on the committee; specifying who is on the committee determines who is not on the committee, and vice versa. So the two sides are equal, as they are two ways of counting the same thing.

Example 1.5.2 (The team captain). For any positive integers n and k with $k \leq n$,

$$n \binom{n-1}{k-1} = k \binom{n}{k}.$$

This is again easy to check algebraically (using the fact that $m! = m(m-1)!$ for any positive integer m), but a story proof is more insightful.

Story proof: Consider a group of n people, from which a team of k will be chosen, one of whom will be the team captain. To specify a possibility, we could first choose the team captain and then choose the remaining $k-1$ team members; this gives the left-hand side. Equivalently, we could first choose the k team members and then choose one of them to be captain; this gives the right-hand side.

Example 1.5.3 (Vandermonde's identity). A famous relationship between binomial coefficients, called Vandermonde's identity,² says that

$$\binom{m+n}{k} = \sum_{j=0}^k \binom{m}{j} \binom{n}{k-j}.$$

This identity will come up several times in this book. Trying to prove it with a brute force expansion of all the binomial coefficients would be a nightmare. But a story proves the result elegantly and makes it clear why the identity holds.

Story proof: Consider a student organization consisting of m juniors and n seniors, from which a committee of size k will be chosen. There are $\binom{m+n}{k}$ possibilities. If there are j juniors in the committee, then there must be $k-j$ seniors in the committee. The right-hand side of the identity sums up the cases for j .

Example 1.5.4 (Partnerships). Let's use a story proof to show that

$$\frac{(2n)!}{2^n \cdot n!} = (2n-1)(2n-3) \cdots 3 \cdot 1.$$

Story proof: We will show that both sides count the number of ways to break $2n$ people into n partnerships. Take $2n$ people, and give them ID numbers from 1 to $2n$. We can form partnerships by lining up the people in some order and then saying the first two are a pair, the next two are a pair, etc. This overcounts by a factor of $n! \cdot 2^n$ since the order of pairs doesn't matter, nor does the order within each pair. Alternatively, count the number of possibilities by noting that there are $2n-1$ choices for the partner of person 1, then $2n-3$ choices for person 2 (or person 3, if person 2 was already paired to person 1), and so on.

4.1.6 1.6 Non-naive Definition of Probability

We have now seen several methods for counting outcomes in a sample space, allowing us to calculate probabilities if the naive definition applies. But the naive definition can only take us so far, since it requires equally likely outcomes and can't handle

an infinite sample space. To generalize the notion of probability, we'll write down a short wish list of how we want probability to behave (in math, the items on the wish list are called axioms), and then define a probability function to be something that satisfies the properties we want!

Here is the general definition of probability that we'll use for the rest of this book. It requires just two axioms, but from these axioms it is possible to prove a vast array of results about probability.

Definition 1.6.1 (General definition of probability). A probability space consists of a sample space S and a probability function P which takes an event $A \subseteq S$ as input and returns $P(A)$, a real number between 0 and 1, as output. The function P must satisfy the following axioms:

1. $P(\emptyset) = 0, P(S) = 1$.

⁰² Vandermonde's identity is named after the 18th century French mathematician AlexandreThéophile Vandermonde, but it was discovered much earlier, and stated in 1303 by the Chinese mathematician Zhu Shijie.

2. If $A_1, A_2, \dots$ are disjoint events, then

$$P\left(\bigcup_{j=1}^{\infty} A_j\right) = \sum_{j=1}^{\infty} P(A_j).$$

(Saying that these events are disjoint means that they are mutually exclusive: $A_i \cap A_j = \emptyset$ for $i \neq j$.)

In Pebble World, the definition says that probability behaves like mass: the mass of an empty pile of pebbles is 0, the total mass of all the pebbles is 1, and if we have non-overlapping piles of pebbles, we can get their combined mass by adding the masses of the individual piles. Unlike in the naive case, we can now have pebbles of differing masses, and we can also have a countably infinite number of pebbles as long as their total mass is 1.

We can even have uncountable sample spaces, such as having S be an area in the plane. In this case, instead of pebbles, we can visualize mud spread out over a region, where the total mass of the mud is 1.

Any function P (mapping events to numbers in the interval $[0, 1]$) that satisfies the two axioms is considered a valid probability function. However, the axioms don't tell us how probability should be interpreted; different schools of thought exist.

The frequentist view of probability is that it represents a long-run frequency over a large number of repetitions of an experiment: if we say a coin has probability $1/2$ of Heads, that means the coin would land Heads 50% of the time if we tossed it over and over and over.

The Bayesian view of probability is that it represents a degree of belief about the event in question, so we can assign probabilities to hypotheses like "candidate A will win the election" or "the defendant is guilty" even if it isn't possible to repeat the same election or the same crime over and over again.

The Bayesian and frequentist perspectives are complementary, and both will be helpful for developing intuition in later chapters. Regardless of how we choose to interpret probability, we can use the two axioms to derive other properties of probability, and these results will hold for any valid probability function.

Theorem 1.6.2 (Properties of probability). Probability has the following properties, for any events A and B .

1. $P(A^c) = 1 - P(A)$.
2. If $A \subseteq B$, then $P(A) \leq P(B)$.
3. $P(A \cup B) = P(A) + P(B) - P(A \cap B)$.

Proof.


1. Since A and A^c are disjoint and their union is S , the second axiom gives

$$P(S) = P(A \cup A^c) = P(A) + P(A^c),$$

But $P(S) = 1$ by the first axiom. So $P(A) + P(A^c) = 1$.

2. If $A \subseteq B$, then we can write B as the union of A and $B \cap A^c$, where $B \cap A^c$ is the part of

B not also in A . This is illustrated in the Venn diagram below.




Since A and $B \cap A^c$ are disjoint, we can apply the second axiom:

$$P(B) = P(A \cup (B \cap A^c)) = P(A) + P(B \cap A^c).$$

Probability is nonnegative, so $P(B \cap A^c) \geq 0$, proving that $P(B) \geq P(A)$.

3. The intuition for this result can be seen using a Venn diagram like the one below.



The shaded region represents $A \cup B$, but the probability of this region is not the sum $P(A) + P(B)$, because that would count the football-shaped region $A \cap B$ twice. To correct for this, we subtract $P(A \cap B)$. This is a useful intuition, but not a proof.

For a proof using the axioms of probability, we can write $A \cup B$ as the union of the disjoint events A and $B \cap A^c$. Then by the second axiom,

$$P(A \cup B) = P(A \cup (B \cap A^c)) = P(A) + P(B \cap A^c).$$

So it suffices to show that $P(B \cap A^c) = P(B) - P(A \cap B)$. Since $A \cap B$ and $B \cap A^c$ are disjoint and their union is B , another application of the second axiom gives us

$$P(A \cap B) + P(B \cap A^c) = P(B).$$

So $P(B \cap A^c) = P(B) - P(A \cap B)$, as desired. $\square$


The third property is a special case of inclusion-exclusion, a formula for finding the probability of a union of events when the events are not necessarily disjoint. We showed above that for two events A and B ,

$$P(A \cup B) = P(A) + P(B) - P(A \cap B).$$

For three events, inclusion-exclusion says

$$\begin{aligned}
P(A \cup B \cup C) = & P(A) + P(B) + P(C) \\
& - P(A \cap B) - P(A \cap C) - P(B \cap C) \\
& + P(A \cap B \cap C).
\end{aligned}$$

For intuition, consider a triple Venn diagram like the one below.



To get the total area of the shaded region $A \cup B \cup C$, we start by adding the areas of the three circles, $P(A) + P(B) + P(C)$. The three football-shaped regions have each been counted twice, so we then subtract $P(A \cap B) + P(A \cap C) + P(B \cap C)$. Finally, the region in the center has been added three times and subtracted three times, so in order to count it exactly once, we must add it back again. This ensures that each region of the diagram has been counted once and exactly once.

Now we can write inclusion-exclusion for n events.

Theorem 1.6.3 (Inclusion-exclusion). For any events $A_1, \dots, A_n$,

$$\begin{aligned}
P\left(\bigcup_{i=1}^n A_i\right) = & \sum_i P(A_i) - \sum_{i < j} P(A_i \cap A_j) + \sum_{i < j < k} P(A_i \cap A_j \cap A_k) - \dots \\
& + (-1)^{n+1} P(A_1 \cap \dots \cap A_n).
\end{aligned}$$

This formula can be proven by induction using just the axioms, but instead we'll present a shorter proof in Chapter 4 after introducing some additional tools. The rationale behind the alternating addition and subtraction in the general formula is analogous to the special cases we've already considered.

The next example, de Montmort's matching problem, is a famous application of inclusion-exclusion. Pierre Rémond de Montmort was a French mathematician who studied probability in the context of gambling and wrote a treatise [19] devoted to the analysis of various card games. He posed the following problem in 1708, based on a card game called Treize.

Example 1.6.4 (de Montmort's matching problem). Consider a well-shuffled deck of n cards, labeled 1 through n . You flip over the cards one by one, saying the numbers 1 through n as you do so. You win the game if, at some point, the number you say aloud is the same as the number on the card being flipped over (for example, if the 7th card in the deck has the label 7). What is the probability of winning?

Solution:

Let A_i be the event that the i th card in the deck has the number i written on it. We are interested in the probability of the union $A_1 \cup \dots \cup A_n$: as long as at least one of the cards has a number matching its position in the deck, you will win the game. (An ordering for which you

lose is called a derangement, though hopefully no one has ever become deranged due to losing at this game.)

To find the probability of the union, we'll use inclusion-exclusion. First,

$$P(A_i) = \frac{1}{n}$$

for all i . One way to see this is with the naive definition of probability, using the full sample space: there are $n!$ possible orderings of the deck, all equally likely, and $(n-1)!$ of these are favorable to A_i (fix the card numbered i to be in the i th position in the deck, and then the remaining $n-1$ cards can be in any order). Another way to see this is by symmetry: the card numbered i is equally likely to be in any of the n positions in the deck, so it has probability $1/n$ of being in the i th spot. Second,

$$P(A_i \cap A_j) = \frac{(n-2)!}{n!} = \frac{1}{n(n-1)},$$

since we require the cards numbered i and j to be in the i th and j th spots in the deck and allow the remaining $n-2$ cards to be in any order, so $(n-2)!$ out of $n!$ possibilities are favorable to $A_i \cap A_j$. Similarly,

$$P(A_i \cap A_j \cap A_k) = \frac{1}{n(n-1)(n-2)},$$

and the pattern continues for intersections of 4 events, etc.

In the inclusion-exclusion formula, there are n terms involving one event, $\binom{n}{2}$ terms involving two events, $\binom{n}{3}$ terms involving three events, and so forth. By the symmetry of the problem, all n terms of the form $P(A_i)$ are equal, all $\binom{n}{2}$ terms of the form $P(A_i \cap A_j)$ are equal, and the whole expression simplifies considerably:

$$\begin{aligned} P\left(\bigcup_{i=1}^n A_i\right) &= \frac{n}{n} - \frac{\binom{n}{2}}{n(n-1)} + \frac{\binom{n}{3}}{n(n-1)(n-2)} - \cdots + (-1)^{n+1} \cdot \frac{1}{n!} \\ &= 1 - \frac{1}{2!} + \frac{1}{3!} - \cdots + (-1)^{n+1} \cdot \frac{1}{n!} \end{aligned}$$

Comparing this to the Taylor series for $1/e$ (see Section A. 8 of the math appendix),

$$e^{-1} = 1 - \frac{1}{1!} + \frac{1}{2!} - \frac{1}{3!} + \cdots,$$

we see that for large n , the probability of winning the game is extremely close to $1 - 1/e$, or about 0.63. Interestingly, as n grows, the probability of winning approaches $1 - 1/e$ instead of going to 0 or 1. With a lot of cards in the deck, the number of possible locations for matching cards increases while the probability of any particular match decreases, and these two forces offset each other and balance to give a probability of about $1 - 1/e$.

Inclusion-exclusion is a very general formula for the probability of a union of events, but it helps us the most when there is symmetry among the events A_j ; otherwise the sum can be extremely tedious. In general, when symmetry is lacking, we should try to use other tools before turning to inclusion-exclusion as a last resort.

4.1.7 1.7 Recap

Probability allows us to quantify uncertainty and randomness in a principled way. Probabilities arise when we perform an experiment: the set of all possible outcomes of the experiment is called the sample space, and a subset of the sample space is called an event. It is useful

to be able to go back and forth between describing events in English and writing them down mathematically as sets (often using unions, intersections, and complements).

Pebble World can help us visualize sample spaces and events when the sample space is finite. In Pebble World, each outcome is a pebble, and an event is a set of pebbles. If all the pebbles have the same mass (i.e., are equally likely), we can apply the naive definition of probability, which lets us calculate probabilities by counting.

To this end, we discussed several tools for counting. When counting the number of possibilities, we often use the multiplication rule. For example, there are $n!$ permutations of the numbers $1, 2, \dots, n$ and there are 2^n subsets of a set with n elements. If we can't directly use the multiplication rule, we can sometimes count each possibility exactly c times for some c , and then divide by c to get the actual number of possibilities. For example, this strategy is useful for finding an expression for binomial coefficients in terms of factorials.

An important pitfall to avoid is misapplying the naive definition of probability, implicitly or explicitly assuming equally likely outcomes without justification. One technique to help avoid this is to give objects labels, for precision and so that we are not tempted to treat them as indistinguishable.

Moving beyond the naive definition, we define probability to be a function that takes an event and assigns to it a real number between 0 and 1. We require a valid probability function to satisfy two axioms:

1. $P(\emptyset) = 0, P(S) = 1$.
2. If $A_1, A_2, \dots$ are disjoint events, then

$$P\left(\bigcup_{j=1}^{\infty} A_j\right) = \sum_{j=1}^{\infty} P(A_j).$$

Many useful properties can be derived from these axioms. For example,

$$P(A^c) = 1 - P(A)$$

for any event A , and we have the inclusion-exclusion formula

$$P\left(\bigcup_{i=1}^n A_i\right) = \sum_i P(A_i) - \sum_{i < j} P(A_i \cap A_j) + \sum_{i < j < k} P(A_i \cap A_j \cap A_k) - \dots + (-1)^{n+1} P(A_1 \cap \dots \cap A_n)$$

for any events $A_1, \dots, A_n$. For $n = 2$, this is the much nicer-looking result

$$P(A_1 \cup A_2) = P(A_1) + P(A_2) - P(A_1 \cap A_2).$$

Figure 1.7 illustrates how a probability function maps events to numbers between 0 and 1. We'll add many new concepts to this diagram as we continue our journey through the field of probability.




FIGURE 1.7

It is important to distinguish between events and probabilities. The former are sets, while the latter are numbers. Before the experiment is done, we generally don't know whether or not a particular event will occur (happen). So we assign it a probability of happening, using a probability function P . We can use set operations to define new events in terms of old events, and the properties of probabilities to relate the probabilities of the new events to those of the old events.

4.1.8 1.8R

R is a very powerful, popular environment for statistical computing and graphics, freely available for Mac OS X, Windows, and UNIX systems. Knowing how to use R is an extremely useful skill. R and various supporting information can be obtained at <https://www.r-project.org>. RStudio is an excellent alternative interface for R, freely available at <https://www.rstudio.com>.

In the R section at the end of each chapter, we provide R code to let you try out some of the examples from the chapter, especially via simulation. These sections are not intended to be a full introduction to R; many R tutorials are available for free online, and many books on R are available. But the R sections show how to implement various simulations, computations, and visualizations that naturally accompany the material from each chapter. The R code at the end of each chapter is also available at <http://stat110.net>.

Vectors

R is built around vectors, and getting familiar with "vectorized thinking" is very important for using R effectively. To create a vector, we can use the `c` command (which stands for combine or concatenate). For example,

```
v <- c(3, 1, 4, 1, 5, 9)
```

defines $\mathbf{v}$ to be the vector $(3, 1, 4, 1, 5, 9)$. (The left arrow `<-` is typed as `<` followed by `-`. The symbol `=` can be used instead, but the arrow is more suggestive of the fact that the variable on the left is being set equal to the value on the right.) Similarly, `n <- 110` sets n equal to 110; R views n as a vector of length 1.

```
sum(v)
```

adds up the entries of $\mathbf{v}$, `max(v)` gives the largest value, `min(v)` gives the smallest value, and `length(v)` gives the length.

A shortcut for getting the vector $(1, 2, \dots, n)$ is to type `1:n`; more generally, if m and n are integers then `m:n` gives the sequence of integers from m to n (in increasing order if $m \leq n$ and in decreasing order otherwise).

To access the i th entry of a vector $\mathbf{v}$, use `v[i]`. We can also get subvectors easily:

```
v[c(1,3,5)]
```

gives the vector consisting of the 1st, 3rd, and 5th entries of $\mathbf{v}$. It's also possible to get a subvector by specifying what to exclude, using a minus sign:

```
v[-(2:4)]
```

gives the vector obtained by removing the 2nd through 4th entries of $\mathbf{v}$ (the parentheses are needed since $-2 : 4$ would be $(-2, -1, \dots, 4)$).

Many operations in R are interpreted componentwise. For example, in math the cube of a vector doesn't have a standard definition, but in R typing $\mathbf{v}^3$ simply cubes each entry individually. Similarly,

```
1/ ( 1 : 100)^2
```

is a very compact way to get the vector $(1, \frac{1}{2^2}, \frac{1}{3^2}, \dots, \frac{1}{100^2})$.

In math, $\mathbf{v} + \mathbf{w}$ is undefined if $\mathbf{v}$ and $\mathbf{w}$ are vectors of different lengths, but in R the shorter vector gets "recycled"! For example, $\mathbf{v} + 3$ adds 3 to each entry of $\mathbf{v}$.

Factorials and Binomial Coefficients

We can compute $n!$ using `factorial(n)` and $\binom{n}{k}$ using `choose(n,k)`. As we have seen, factorials grow extremely quickly. What is the largest n for which R returns a number for `factorial(n)`? Beyond that point, R will return `Inf` (infinity), with a warning message. But it may still be possible to use `lfactorial(n)` for larger values of n , which computes $\log(n!)$. Similarly, `lchoose(n,k)` computes $\log\binom{n}{k}$.

vSampling and simulation The `sample` command is a useful way of drawing random samples in R. (Technically, they are pseudo-random since there is an underlying deterministic algorithm, but they "look like" random samples for almost all practical purposes.) For example,

```
n <- 10; k <- 5
```

```
sample(n,k)
```

generates an ordered random sample of 5 of the numbers from 1 to 10, without replacement, and with equal probabilities given to each number. To sample with replacement instead, just add in `replace = TRUE`:

```
n <- 10; k <- 5
```

```
sample(n,k,replace=TRUE)
```

To generate a random permutation of $1, 2, \dots, n$ we can use `sample(n,n)`, which because of R's default settings can be abbreviated to `sample(n)`.

We can also use `sample` to draw from a non-numeric vector. For example, `letters` is built into R as the vector consisting of the 26 lowercase letters of the English alphabet, and `sample(letters,7)` will generate a random 7-letter "word" by sampling from the alphabet, without replacement.

The `sample` command also allows us to specify general probabilities for sampling each number. For example,

```
sample(4, 3, replace=TRUE, prob=c(0.1,0.2,0.3,0.4))
```

samples three numbers between 1 and 4, with replacement, and with probabilities given by $(0.1, 0.2, 0.3, 0.4)$. If the sampling is without replacement, then at each stage the probability of any not-yet-chosen number is proportional to its original probability.

Generating many random samples allows us to perform a simulation for a probability problem. The `replicate` command, which is explained below, is a convenient way to do this.

Matching Problem Simulation

Let's show by simulation that the probability of a matching card in Example 1.6.4 is approximately $1 - 1/e$ when the deck is sufficiently large. Using R, we can perform the experiment

a bunch of times and see how many times we encounter at least one matching card:

```
n <- 100
r <- replicate(10^4, sum(sample(n) == (1:n)))
sum(r >= 1) / 10^4
```

In the first line, we choose how many cards are in the deck (here, 100 cards). In the second line, let's work from the inside out:

- `sample(n) == (1 : n)` is a vector of length `n`, the i th element of which equals 1 if the i th card matches its position in the deck and 0 otherwise. That's because for two numbers a and b , the expression `a == b` is TRUE if $a = b$ and FALSE otherwise, and TRUE is encoded as 1 and FALSE is encoded as 0.
- `sum` adds up the elements of the vector, giving us the number of matching cards in this run of the experiment.
- `replicate` does this 10^4 times. We store the results in `r`, a vector of length 10^4 containing the numbers of matched cards from each run of the experiment.

In the last line, we add up the number of times where there was at least one matching card, and we divide by the number of simulations.

To explain what the code is doing within the code rather than in separate documentation, we can add comments using the `#` symbol to mark the start of a comment. Comments are ignored by R but can make the code much easier to understand for the reader (who could be you—even if you will be the only one using your code, it is often hard to remember what everything means and how the code is supposed to work when looking at it a month after writing it). Short comments can be on the same line as the corresponding code; longer comments should be on separate lines. For example, a commented version of the above simulation is:

```
n <- 100 # number of cards
r <- replicate(10^4, sum(sample(n) == (1:n))) # shuffle; count matches
sum(r >= 1) / 10^4 # proportion with a match
```

What did you get when you ran the code? We got 0.63, which is quite close to $1 - 1/e$.

Birthday Problem Calculation and Simulation

The following code uses `prod` (which gives the product of a vector) to calculate the probability of at least one birthday match in a group of 23 people:

```
k <- 23
1 - prod((365 - k + 1) : 365) / 365^k
```

Better yet, R has built-in functions, `pbirthday` and `qbirthday`, for the birthday problem! `pbirthday(k)` returns the probability of at least one match if the room has k people. `qbirthday(p)` returns the number of people needed in order to have probability p of at least one match. For example, `pbirthday(23)` is 0.507 and `qbirthday(0.5)` is 23.

We can also find the probability of having at least one triple birthday match, i.e., three people with the same birthday; all we have to do is add `coincident = 3` to say we're looking for triple matches. For example, `pbirthday(23, coincident = 3)` returns 0.014, so 23 people give us only a 1.4% chance of a triple birthday match. `qbirthday(0.5, coincident = 3)` returns 88, so we'd need 88 people to have at least a 50% chance of at least one triple birthday match.

To simulate the birthday problem, we can use

```
b <- sample(1:365,23,replace=TRUE)
tabulate(b)
```

to generate random birthdays for 23 people and then tabulate the counts of how many people were born on each day (the command `table(b)` creates a prettier table, but is slower). We can run 10^4 repetitions as follows:

```
r <- replicate(10^4, max(tabulate(sample(1:365,23,replace=TRUE))))
sum(r>=2)/10^4
```

If the probabilities of various days are not all equal, the calculation becomes much more difficult, but the simulation can easily be extended since `sample` allows us to specify the probability of each day (by default `sample` assigns equal probabilities, so in the above the probability is $1/365$ for each day).

4.1.9 1.9 Exercises

Exercises marked with (S) have detailed solutions at <http://stat110.net>.

Counting

1. How many ways are there to permute the letters in the word MISSISSIPPI?
2. (a) How many 7 -digit phone numbers are possible, assuming that the first digit can't be a 0 or a 1 ?
(b) Re-solve (a), except now assume also that the phone number is not allowed to start with 911 (since this is reserved for emergency use, and it would not be desirable for the system to wait to see whether more digits were going to be dialed after someone has dialed 911).
3. Fred is planning to go out to dinner each night of a certain week, Monday through Friday, with each dinner being at one of his ten favorite restaurants.
(a) How many possibilities are there for Fred's schedule of dinners for that Monday through Friday, if Fred is not willing to eat at the same restaurant more than once?
(b) How many possibilities are there for Fred's schedule of dinners for that Monday through Friday, if Fred is willing to eat at the same restaurant more than once, but is not willing to eat at the same place twice in a row (or more)?
4. A round-robin tournament is being held with n tennis players; this means that every player will play against every other player exactly once.
(a) How many possible outcomes are there for the tournament (the outcome lists out who won and who lost for each game)?
(b) How many games are played in total?
5. A knock-out tournament is being held with 2^n tennis players. This means that for each round, the winners move on to the next round and the losers are eliminated, until only one person remains. For example, if initially there are $2^4 = 16$ players, then there are 8 games in the first round, then the 8 winners move on to round 2 , then the 4 winners move on to round 3 , then the 2 winners move on to round 4 , the winner of which is declared the winner of the tournament. (There are various systems for determining who plays whom within a round, but these do not matter for this problem.)
(a) How many rounds are there?
(b) Count how many games in total are played, by adding up the numbers of games

- played in each round.
- (c) Count how many games in total are played, this time by directly thinking about it without doing almost any calculation.
- Hint: How many players need to be eliminated?
6. There are 20 people at a chess club on a certain day. They each find opponents and start playing. How many possibilities are there for how they are matched up, assuming that in each game it does matter who has the white pieces (in a chess game, one player has the white pieces and the other player has the black pieces)?
7. Two chess players, A and B, are going to play 7 games. Each game has three possible outcomes: a win for A (which is a loss for B), a draw (tie), and a loss for A (which is a win for B). A win is worth 1 point, a draw is worth 0.5 points, and a loss is worth 0 points.
- (a) How many possible outcomes for the individual games are there, such that overall player A ends up with 3 wins, 2 draws, and 2 losses?
- (b) How many possible outcomes for the individual games are there, such that A ends up with 4 points and B ends up with 3 points?
- (c) Now assume that they are playing a best-of-7 match, where the match will end when either player has 4 points or when 7 games have been played, whichever is first. For example, if after 6 games the score is 4 to 2 in favor of A, then A wins the match and they don't play a 7th game. How many possible outcomes for the individual games are there, such that the match lasts for 7 games and A wins by a score of 4 to 3?
8. (S) (a) How many ways are there to split a dozen people into 3 teams, where one team has 2 people, and the other two teams have 5 people each?
- (b) How many ways are there to split a dozen people into 3 teams, where each team has 4 people?
9. (S) (a) How many paths are there from the point $(0, 0)$ to the point $(110, 111)$ in the plane such that each step either consists of going one unit up or one unit to the right?
- (b) How many paths are there from $(0, 0)$ to $(210, 211)$, where each step consists of going one unit up or one unit to the right, and the path has to go through $(110, 111)$?
10. To fulfill the requirements for a certain degree, a student can choose to take any 7 out of a list of 20 courses, with the constraint that at least 1 of the 7 courses must be a statistics course. Suppose that 5 of the 20 courses are statistics courses.
- (a) How many choices are there for which 7 courses to take?
- (b) Explain intuitively why the answer to (a) is not $\binom{5}{1} \cdot \binom{19}{6}$.
11. Let A and B be sets with $|A| = n$, $|B| = m$.
- (a) How many functions are there from A to B (i.e., functions with domain A , assigning an element of B to each element of A)?
- (b) How many one-to-one functions are there from A to B ? (See Section A.2.1 of the math appendix for information about one-to-one functions.)
12. Four players, named A, B, C, and D, are playing a card game. A standard, well-shuffled deck of cards is dealt to the players (so each player receives a 13-card hand).
- (a) How many possibilities are there for the hand that player A will get? (Within a hand, the order in which cards were received doesn't matter.)
- (b) How many possibilities are there overall for what hands everyone will get, assuming that it matters which player gets which hand, but not the order of cards within a hand?

- (c) Explain intuitively why the answer to Part (b) is not the fourth power of the answer to Part (a).
13. A certain casino uses 10 standard decks of cards mixed together into one big deck, which we will call a superdeck. Thus, the superdeck has $52 \cdot 10 = 520$ cards, with 10 copies of each card. How many different 10 -card hands can be dealt from the superdeck? The order of the cards does not matter, nor does it matter which of the original 10 decks the cards came from. Express your answer as a binomial coefficient.
Hint: Bose-Einstein.
14. You are ordering two pizzas. A pizza can be small, medium, large, or extra large, with any combination of 8 possible toppings (getting no toppings is allowed, as is getting all 8). How many possibilities are there for your two pizzas?

Story Proofs

15. (S) Give a story proof that $\sum_{k=0}^n \binom{n}{k} = 2^n$.
16. (S) Show that for all positive integers n and k with $n \geq k$,

$$\binom{n}{k} + \binom{n}{k-1} = \binom{n+1}{k},$$

doing this in two ways: (a) algebraically and (b) with a story, giving an interpretation for why both sides count the same thing.

Hint for the story proof: Imagine an organization consisting of $n+1$ people, with one of them pre-designated as the president of the organization.

17. Give a story proof that

$$\sum_{k=0}^n \binom{n}{k}^2 = \binom{2n}{n},$$

for all positive integers n .

18. Give a story proof that

$$\sum_{k=1}^n k \binom{n}{k}^2 = n \binom{2n-1}{n-1}$$

for all positive integers n .

Hint: Consider choosing a committee of size n from two groups of size n each, where only one of the two groups has people eligible to become the chair of the committee.

19. Give a story proof that

$$\sum_{k=2}^n \binom{k}{2} \binom{n-k+2}{2} = \binom{n+3}{5},$$

for all integers $n \geq 2$.

Hint: Consider the middle number in a subset of $\{1, 2, \dots, n+3\}$ of size 5.

20. (S) (a) Show using a story proof that

$$\binom{k}{k} + \binom{k+1}{k} + \binom{k+2}{k} + \cdots + \binom{n}{k} = \binom{n+1}{k+1},$$

where n and k are positive integers with $n \geq k$. This is called the hockey stick identity. Hint: Imagine arranging a group of people by age, and then think about the oldest person in a chosen subgroup.

(b) Suppose that a large pack of Haribo gummi bears can have anywhere between 30 and 50 gummi bears. There are 5 delicious flavors: pineapple (clear), raspberry (red), orange (orange), strawberry (green, mysteriously), and lemon (yellow). There are 0 nondelicious flavors. How many possibilities are there for the composition of such a pack of gummi bears? You can leave your answer in terms of a couple binomial coefficients, but not a sum of lots of binomial coefficients.

21. Define $\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\}$ as the number of ways to partition $\{1, 2, \dots, n\}$ into k nonempty subsets, or the number of ways to have n students split up into k groups such that each group has at least one student. For example, $\left\{ \begin{smallmatrix} 4 \\ 2 \end{smallmatrix} \right\} = 7$ because we have the following possibilities.

- $\{1\}, \{2, 3, 4\}$
- $\{2\}, \{1, 3, 4\}$
- $\{3\}, \{1, 2, 4\}$
- $\{4\}, \{1, 2, 3\}$
- $\{1, 2\}, \{3, 4\}$
- $\{1, 3\}, \{2, 4\}$
- $\{1, 4\}, \{2, 3\}$

Prove the following identities:

(a)

$$\left\{ \begin{smallmatrix} n+1 \\ k \end{smallmatrix} \right\} = \left\{ \begin{smallmatrix} n \\ k-1 \end{smallmatrix} \right\} + k \left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\}.$$

Hint: I'm either in a group by myself or I'm not.

(b)

$$\sum_{j=k}^n \binom{n}{j} \left\{ \begin{smallmatrix} j \\ k \end{smallmatrix} \right\} = \left\{ \begin{smallmatrix} n+1 \\ k+1 \end{smallmatrix} \right\}.$$

Hint: First decide how many people are not going to be in my group.

22. The Dutch mathematician R.J. Stroeker remarked:

Every beginning student of number theory surely must have marveled at the miraculous fact that for each natural number n the sum of the first n positive consecutive cubes is a perfect square. [26]

Furthermore, it is the square of the sum of the first n positive integers! That is,

$$1^3 + 2^3 + \dots + n^3 = (1 + 2 + \dots + n)^2.$$

Usually this identity is proven by induction, but that does not give much insight into why the result is true, nor does it help much if we wanted to compute the left-hand side but didn't already know this result. In this problem, you will give a story proof of the identity.

(a) Give a story proof of the identity

$$1 + 2 + \cdots + n = \binom{n+1}{2}$$

Hint: Consider a round-robin tournament (see Exercise 4).

(b) Give a story proof of the identity

$$1^3 + 2^3 + \cdots + n^3 = 6\binom{n+1}{4} + 6\binom{n+1}{3} + \binom{n+1}{2}.$$

It is then just basic algebra (not required for this problem) to check that the square of the right-hand side in (a) is the right-hand side in (b).

Hint: Imagine choosing a number between 1 and n and then choosing 3 numbers between 0 and n smaller than the original number, with replacement. Then consider cases based on how many distinct numbers were chosen.

Naive Definition of Probability

23. Three people get into an empty elevator at the first floor of a building that has 10 floors. Each presses the button for their desired floor (unless one of the others has already pressed that button). Assume that they are equally likely to want to go to floors 2 through 10 (independently of each other). What is the probability that the buttons for 3 consecutive floors are pressed?
24. (S) A certain family has 6 children, consisting of 3 boys and 3 girls. Assuming that all birth orders are equally likely, what is the probability that the 3 eldest children are the 3 girls?
25. (S) A city with 6 districts has 6 robberies in a particular week. Assume the robberies are located randomly, with all possibilities for which robbery occurred where equally likely. What is the probability that some district had more than 1 robbery?
26. A survey is being conducted in a city with 1 million residents. It would be far too expensive to survey all of the residents, so a random sample of size 1000 is chosen (in practice, there are many challenges with sampling, such as obtaining a complete list of everyone in the city, and dealing with people who refuse to participate). The survey is conducted by choosing people one at a time, with replacement and with equal probabilities.
 - (a) Explain how sampling with vs. without replacement here relates to the birthday problem.
 - (b) Find the probability that at least one person will get chosen more than once.
27. A hash table is a commonly used data structure in computer science, allowing for fast information retrieval. For example, suppose we want to store some people's phone numbers. Assume that no two of the people have the same name. For each name x , a hash function h is used, letting $h(x)$ be the location that will be used to store x 's phone number. After such a table has been computed, to look up x 's phone number one just recomputes $h(x)$ and then looks up what is stored in that location.

The hash function h is deterministic, since we don't want to get different results every time we compute $h(x)$. But h is often chosen to be pseudorandom. For this problem, assume that true randomness is used. Let there be k people, with each person's phone number stored in a random location (with equal probabilities for each location, independently of where the other people's numbers are stored), represented by an integer between 1 and n . Find the probability that at least one location has more than one phone number stored there.

28. (S) A college has 10 time slots for its courses, and blithely assigns courses to completely random time slots, independently. The college offers exactly 3 statistics courses. What is the probability that 2 or more of the statistics courses are in the same time slot?
29. (S) For each part, decide whether the blank should be filled in with $=$, $<$, or $>$, and give a clear explanation.
- (a) (probability that the total after rolling 4 fair dice is 21) _____ (probability that the total after rolling 4 fair dice is 22)
- (b) (probability that a random 2-letter word is a palindrome ³) _____ (probability that a random 3-letter word is a palindrome)
30. With definitions as in the previous problem, find the probability that a random n -letter word is a palindrome for $n = 7$ and for $n = 8$.
31. (S) Elk dwell in a certain forest. There are N elk, of which a simple random sample of size n are captured and tagged ("simple random sample" means that all $\binom{N}{n}$ sets of n elk are equally likely). The captured elk are returned to the population, and then a new sample is drawn, this time with size m . This is an important method that is widely used in ecology, known as capture-recapture. What is the probability that exactly k of the m elk in the new sample were previously tagged? (Assume that an elk that was captured before doesn't become more or less likely to be captured again.)
32. Four cards are face down on a table. You are told that two are red and two are black, and you need to guess which two are red and which two are black. You do this by pointing to the two cards you're guessing are red (and then implicitly you're guessing that the other two are black). Assume that all configurations are equally likely, and that you do not have psychic powers. Find the probability that exactly j of your guesses are correct, for $j = 0, 1, 2, 3, 4$.
33. (S) A jar contains r red balls and g green balls, where r and g are fixed positive integers. A ball is drawn from the jar randomly (with all possibilities equally likely), and then a second ball is drawn randomly.
- (a) Explain intuitively why the probability of the second ball being green is the same as the probability of the first ball being green.
- (b) Define notation for the sample space of the problem, and use this to compute the probabilities from (a) and show that they are the same.
- (c) Suppose that there are 16 balls in total, and that the probability that the two balls are the same color is the same as the probability that they are different colors. What are r and g (list all possibilities)?
34. (S) A random 5-card poker hand is dealt from a standard deck of cards. Find the probability of each of the following possibilities (in terms of binomial coefficients).
- (a) A flush (all 5 cards being of the same suit; do not count a royal flush, which is a flush with an ace, king, queen, jack, and 10).
- (b) Two pair (e.g., two 3's, two 7's, and an ace).
35. A random 13-card hand is dealt from a standard deck of cards. What is the probability that the hand contains at least 3 cards of every suit?
36. A group of 30 dice are thrown. What is the probability that 5 of each of the values 1, 2, 3, 4, 5, 6 appear?
37. A deck of cards is shuffled well. The cards are dealt one by one, until the first time an ace

⁰³ A palindrome is an expression such as "A man, a plan, a canal: Panama" that reads the same backwards as forwards (ignoring spaces, capitalization, and punctuation). Assume for this problem that all words of the specified length are equally likely, that there are no spaces or punctuation, and that the alphabet consists of the lowercase letters a, b, ..., z. A word is any string of letters from the alphabet; it does not need to be a word that has a meaning in the English language.

appears.

- (a) Find the probability that no kings, queens, or jacks appear before the first ace.
- (b) Find the probability that exactly one king, exactly one queen, and exactly one jack appear (in any order) before the first ace.

38. Tyrion, Cersei, and ten other people are sitting at a round table, with their seating arrangement having been randomly assigned. What is the probability that Tyrion and Cersei are sitting next to each other? Find this in two ways:

- (a) using a sample space of size $12!$, where an outcome is fully detailed about the seating;
- (b) using a much smaller sample space, which focuses on Tyrion and Cersei.

39. An organization with $2n$ people consists of n married couples. A committee of size k is selected, with all possibilities equally likely. Find the probability that there are exactly j married couples within the committee.

40. There are n balls in a jar, labeled with the numbers $1, 2, \dots, n$. A total of k balls are drawn, one by one with replacement, to obtain a sequence of numbers.

- (a) What is the probability that the sequence obtained is strictly increasing?
- (b) What is the probability that the sequence obtained is increasing (but not necessarily strictly increasing, i.e., there can be repetitions)?

41. Each of n balls is independently placed into one of n boxes, with all boxes equally likely. What is the probability that exactly one box is empty?

42. (S) A *norepeatword* is a sequence of at least one (and possibly all) of the usual 26 letters $a, b, c, \dots, z$, with repetitions not allowed. For example, "course" is a *norepeatword*, but "statistics" is not. Order matters, e.g., "course" is not the same as "source".

A *norepeatword* is chosen randomly, with all *norepeatwords* equally likely. Show that the probability that it uses all 26 letters is very close to $1/e$.

Axioms of Probability

43. Show that for any events A and B ,

$$P(A) + P(B) - 1 \leq P(A \cap B) \leq P(A \cup B) \leq P(A) + P(B).$$

For each of these three inequalities, give a simple criterion for when the inequality is actually an equality (e.g., give a simple condition such that $P(A \cap B) = P(A \cup B)$ if and only if the condition holds).

44. Let A and B be events. The difference $B - A$ is defined to be the set of all elements of B that are not in A . Show that if $A \subseteq B$, then

$$P(B - A) = P(B) - P(A)$$

directly using the axioms of probability.

45. Let A and B be events. The symmetric difference $A \triangle B$ is defined to be the set of all elements that are in A or B but not both. In logic and engineering, this event is also called the *XOR* (exclusive or) of A and B . Show that

$$P(A \triangle B) = P(A) + P(B) - 2P(A \cap B)$$

directly using the axioms of probability.

46. Let $A_1, A_2, \dots, A_n$ be events. Let B_k be the event exactly k of the A_i occur, and C_k be the event that at least k of the A_i occur, for $0 \leq k \leq n$. Find a simple expression for $P(B_k)$ in terms of $P(C_k)$ and $P(C_{k+1})$.

47. Events A and B are independent if $P(A \cap B) = P(A)P(B)$ (independence is explored in detail in the next chapter).

- (a) Give an example of independent events A and B in a finite sample space S (with neither equal to $\emptyset$ or S), and illustrate it with a Pebble World diagram.
- (b) Consider the experiment of picking a random point in the rectangle

$$R = \{(x, y) : 0 < x < 1, 0 < y < 1\}$$

where the probability of the point being in any particular region contained within R is the area of that region. Let A_1 and B_1 be rectangles contained within R , with areas not equal to 0 or 1. Let A be the event that the random point is in A_1 , and B be the event that the random point is in B_1 . Give a geometric description of when it is true that A and B are independent. Also, give an example where they are independent and another example where they are not independent.

- (c) Show that if A and B are independent, then

$$P(A \cup B) = P(A) + P(B) - P(A)P(B) = 1 - P(A^c)P(B^c)$$

48. (S) Arby has a belief system assigning a number $P_{\text{Arby}}(A)$ between 0 and 1 to every event A (for some sample space). This represents Arby's degree of belief about how likely A is to occur. For any event A , Arby is willing to pay a price of $1000 \cdot P_{\text{Arby}}(A)$ dollars to buy a certificate such as the one shown below:

Certificate

The owner of this certificate can redeem it for \$1000 if A occurs. No value if A does not occur, except as required by federal, state, or local law. No expiration date.

Likewise, Arby is willing to sell such a certificate at the same price. Indeed, Arby is willing to buy or sell any number of certificates at this price, as Arby considers it the "fair" price. Arby stubbornly refuses to accept the axioms of probability. In particular, suppose that there are two disjoint events A and B with

$$P_{\text{Arby}}(A \cup B) \neq P_{\text{Arby}}(A) + P_{\text{Arby}}(B).$$

Show how to make Arby go bankrupt, by giving a list of transactions Arby is willing to make that will guarantee that Arby will lose money (you can assume it will be known whether A occurred and whether B occurred the day after any certificates are bought/sold).

Inclusion-exclusion

49. A fair die is rolled n times. What is the probability that at least 1 of the 6 values never appears?
50. (S) A card player is dealt a 13-card hand from a well-shuffled, standard deck of cards. What is the probability that the hand is void in at least one suit ("void in a suit" means having no cards of that suit)?
51. (S) For a group of 7 people, find the probability that all 4 seasons (winter, spring, summer, fall) occur at least once each among their birthdays, assuming that all seasons are equally likely.
52. A certain class has 20 students, and meets on Mondays and Wednesdays in a classroom with exactly 20 seats. In a certain week, everyone in the class attends both days. On both days, the students choose their seats completely randomly (with one student per seat). Find the probability that no one sits in the same seat on both days of that week.

56. For each part, decide whether the blank should be filled in with $=$, $<$, or $>$, and give a clear explanation. In (a) and (b), order doesn't matter.

(a) (number of ways to choose 5 people out of 10) ____ (number of ways to choose 6 people out of 10)

(b) (number of ways to break 10 people into 2 teams of 5) ____ (number of ways to break 10 people into a team of 6 and a team of 4)

(c) (probability that all 3 people in a group of 3 were born on January 1) ____ (probability that in a group of 3 people, 1 was born on each of January 1, 2, and 3)

(d) (probability that Martin wins) 1/2

57. Take a deep breath before attempting this problem. In the book *Innumeracy* [20], John Allen Paulos writes:

Now for better news of a kind of immortal persistence. First, take a deep breath. Assume Shakespeare's account is accurate and Julius Caesar gasped ["Et tu, Brute!"] before breathing his last. What are the chances you just inhaled a molecule which Caesar exhaled in his dying breath?

127

Suppose that the molecules in the atmosphere now are the same as those in the atmosphere when Caesar was alive, and that in the 2000 years or so since Caesar, these molecules have been scattered completely randomly through the atmosphere. Also assume that Caesar's last breath was sampled without replacement but that your breathing is sampled with replacement (without replacement makes more sense but with replacement is easier to work with, and is a good approximation since the number of molecules in the atmosphere is so much larger than the number of molecules in one breath).

Find the probability that at least one molecule in the breath you just took was shared with Caesar's last breath, and give a simple approximation in terms of e .

Hint: As discussed in the math appendix, $(1 + \frac{x}{n})^n \approx e^x$ for n large.

58. A widget inspector inspects 12 widgets and finds that exactly 3 are defective. Unfortunately, the widgets then get all mixed up and the inspector has to find the 3 defective widgets again by testing widgets one by one.

(a) Find the probability that the inspector will now have to test at least 9 widgets.

(b) Find the probability that the inspector will now have to test at least 10 widgets.

59. There are 15 chocolate bars and 10 children. In how many ways can the chocolate bars be distributed to the children, in each of the following scenarios?

(a) The chocolate bars are fungible (interchangeable).

(b) The chocolate bars are fungible, and each child must receive at least one.

Hint: First give each child a chocolate bar, and then decide what to do with the rest.

(c) The chocolate bars are not fungible (it matters which particular bar goes where).

(d) The chocolate bars are not fungible, and each child must receive at least one.

Hint: The strategy suggested in (b) does not apply. Instead, consider randomly giving the chocolate bars to the children, and apply inclusion-exclusion.

60. Given $n \geq 2$ numbers $(a_1, a_2, \dots, a_n)$ with no repetitions, a bootstrap sample is a sequence $(x_1, x_2, \dots, x_n)$ formed from the a_j 's by sampling with replacement with equal probabilities. Bootstrap samples arise in a widely used statistical method known as the bootstrap. For example, if $n = 2$ and $(a_1, a_2) = (3, 1)$, then the possible bootstrap samples are $(3, 3)$, $(3, 1)$, $(1, 3)$, and $(1, 1)$.

(a) How many possible bootstrap samples are there for $(a_1, \dots, a_n)$?

(b) How many possible bootstrap samples are there for $(a_1, \dots, a_n)$, if order does not matter (in the sense that it only matters how many times each a_j was chosen, not the order in which they were chosen)?

(c) One random bootstrap sample is chosen (by sampling from $a_1, \dots, a_n$ with replacement, as described above). Show that not all unordered bootstrap samples (in the sense of (b)) are equally likely. Find an unordered bootstrap sample $\mathbf{b}_1$ that is as likely as possible, and an unordered bootstrap sample $\mathbf{b}_2$ that is as unlikely as possible. Let p_1 be the probability of getting $\mathbf{b}_1$ and p_2 be the probability of getting $\mathbf{b}_2$ (so p_i is the probability of getting the specific unordered bootstrap sample $\mathbf{b}_i$). What is p_1/p_2 ? What is the ratio of the probability of getting an unordered bootstrap sample whose probability is p_1 to the probability of getting an unordered sample whose probability is p_2 ?

61. (S) There are 100 passengers lined up to board an airplane with 100 seats (with each seat assigned to one of the passengers). The first passenger in line crazily decides to sit in a randomly chosen seat (with all seats equally likely). Each subsequent passenger takes their assigned seat if available, and otherwise sits in a random available seat. What is the probability that the last passenger in line gets to sit in their assigned seat? (This is a common interview problem, and a beautiful example of the power of symmetry.)

Hint: Call the seat assigned to the j th passenger in line "seat j " (regardless of whether the airline calls it seat 23 A or whatever). What are the possibilities for which seats are available to the last passenger in line, and what is the probability of each of these possibilities?

62. In the birthday problem, we assumed that all 365 days of the year are equally likely (and excluded February 29). In reality, some days are slightly more likely as birthdays than others. For example, scientists have long struggled to understand why more babies are born 9 months after a holiday. Let $\mathbf{p} = (p_1, p_2, \dots, p_{365})$ be the vector of birthday probabilities, with p_j the probability of being born on the j th day of the year (February 29 is still excluded, with no offense intended to Leap Dayers).

The k th elementary symmetric polynomial in the variables $x_1, \dots, x_n$ is defined by

$$e_k(x_1, \dots, x_n) = \sum_{1 \leq j_1 < j_2 < \dots < j_k \leq n} x_{j_1} \dots x_{j_k}$$

This just says to add up all of the $\binom{n}{k}$ terms we can get by choosing and multiplying k of the variables. For example, $e_1(x_1, x_2, x_3) = x_1 + x_2 + x_3$, $e_2(x_1, x_2, x_3) = x_1x_2 + x_1x_3 + x_2x_3$, and $e_3(x_1, x_2, x_3) = x_1x_2x_3$.

Now let $k \geq 2$ be the number of people.

(a) Find a simple expression for the probability that there is at least one birthday match, in terms of $\mathbf{p}$ and an elementary symmetric polynomial.

(b) Explain intuitively why it makes sense that P (at least one birthday match) is minimized when $p_j = \frac{1}{365}$ for all j , by considering simple and extreme cases.

(c) The famous arithmetic mean-geometric mean inequality says that for $x, y \geq 0$,

$$\frac{x+y}{2} \geq \sqrt{xy}$$

This inequality follows from adding $4xy$ to both sides of $x^2 - 2xy + y^2 = (x - y)^2 \geq 0$. Define $\mathbf{r} = (r_1, \dots, r_{365})$ by $r_1 = r_2 = (p_1 + p_2)/2$, $r_j = p_j$ for $3 \leq j \leq 365$. Using the arithmetic mean-geometric mean bound and the fact, which you should verify, that

$$e_k(x_1, \dots, x_n) = x_1x_2e_{k-2}(x_3, \dots, x_n) + (x_1 + x_2)e_{k-1}(x_3, \dots, x_n) + e_k(x_3, \dots, x_n)$$

show that

$$P(\text{at least one birthday match} \mid \mathbf{p}) \geq P(\text{at least one birthday match} \mid \mathbf{r}),$$

with strict inequality if $\mathbf{p} \neq \mathbf{r}$, where the "given $\mathbf{r}$ " notation means that the birthday probabilities are given by $\mathbf{r}$. Using this, show that the value of $\mathbf{p}$ that minimizes the probability of at least one birthday match is given by $p_j = \frac{1}{365}$ for all j .

4.2 Conditional Probability

We have introduced probability as a language for expressing our degrees of belief or uncertainties about events. Whenever we observe new evidence (i.e., obtain data), we acquire information that may affect our uncertainties. A new observation that is consistent with an existing belief could make us more sure of that belief, while a surprising observation could throw that belief into question. Conditional probability is the concept that addresses this fundamental question: how should we update our beliefs in light of the evidence we observe?

4.2.1 2.1 The Importance of Thinking Conditionally

Conditional probability is essential for scientific, medical, and legal reasoning, since it shows how to incorporate evidence into our understanding of the world in a logical, coherent manner. In fact, a useful perspective is that all probabilities are conditional; whether or not it's written explicitly, there is always background knowledge (or assumptions) built into every probability.

Suppose, for example, that one morning we are interested in the event R that it will rain that day. Let $P(R)$ be our assessment of the probability of rain before looking outside. If we then look outside and see ominous clouds in the sky, then presumably our probability of rain should increase; we denote this new probability by $P(R | C)$ (read as "probability of R given C "), where C is the event of there being ominous clouds. When we go from $P(R)$ to $P(R | C)$, we say that we are "conditioning on C ". As the day progresses, we may obtain more and more information about the weather conditions, and we can continually update our probabilities. If we observe that events $B_1, \dots, B_n$ occurred, then we write our new conditional probability of rain given this evidence as $P(R | B_1, \dots, B_n)$. If eventually we observe that it does start raining, our conditional probability becomes 1.

Furthermore, we will see that conditioning is a very powerful problem-solving strategy, often making it possible to solve a complicated problem by decomposing it into manageable pieces with case-by-case reasoning. Just as in computer science a common strategy is to break a large problem up into bite-size pieces (or even byte-size pieces), in probability a common strategy is to reduce a complicated probability problem to a bunch of simpler conditional probability problems. In particular, we will discuss a strategy known as first-step analysis, which often allows us to obtain recursive solutions to problems where the experiment has multiple stages.

Due to the central importance of conditioning, both as the means by which we update beliefs to reflect evidence and as a problem-solving strategy, we say that

Conditioning is the soul of statistics.

4.2.2 2.2 Definition and Intuition

Definition 2.2.1 (Conditional probability). If A and B are events with $P(B) > 0$, then the conditional probability of A given B , denoted by $P(A | B)$, is defined as

$$P(A | B) = \frac{P(A \cap B)}{P(B)}.$$

Here A is the event whose uncertainty we want to update, and B is the evidence we observe (or want to treat as given). We call $P(A)$ the prior probability of A and $P(A | B)$ the posterior probability of A ("prior" means before updating based on the evidence, and "posterior" means after updating based on the evidence).

It is important to interpret the event appearing after the vertical conditioning bar as the evidence that we have observed or that is being conditioned on: $P(A | B)$ is the probability of A given the evidence B , not the probability of some entity called $A | B$. As discussed in 2.4.1, there is no such event as $A | B$.

For any event A , $P(A | A) = P(A \cap A)/P(A) = 1$. Upon observing that A has occurred, our updated probability for A is 1. If this weren't the case, we would demand a new definition of conditional probability!

Example 2.2.2 (Two cards). A standard deck of cards is shuffled well. Two cards are drawn randomly, one at a time without replacement. Let A be the event that the first card is a heart, and B be the event that the second card is red. Find $P(A | B)$ and $P(B | A)$.

Solution:

By the naive definition of probability and the multiplication rule,

$$P(A \cap B) = \frac{13 \cdot 25}{52 \cdot 51} = \frac{25}{204},$$

since a favorable outcome is determined by choosing any of the 13 hearts and then any of the remaining 25 red cards. Also, $P(A) = 1/4$ since the 4 suits are equally likely, and

$$P(B) = \frac{26 \cdot 51}{52 \cdot 51} = \frac{1}{2}$$

since there are 26 favorable possibilities for the second card, and for each of those, the first card can be any other card (recall from Chapter 1 that chronological order is not needed in the multiplication rule).

A neater way to see that $P(B) = 1/2$ is by symmetry: from a vantage point before having done the experiment, the second card is equally likely to be any card in the deck.

We now have all the pieces needed to apply the definition of conditional probability:

$$\begin{aligned} P(A | B) &= \frac{P(A \cap B)}{P(B)} = \frac{25/204}{1/2} = \frac{25}{102} \\ P(B | A) &= \frac{P(B \cap A)}{P(A)} = \frac{25/204}{1/4} = \frac{25}{51} \end{aligned}$$

This is a simple example, but already there are several things worth noting.

1. It's extremely important to be careful about which events to put on which side of the conditioning bar. In particular, $P(A | B) \neq P(B | A)$. The next section explores how $P(A | B)$ and $P(B | A)$ are related in general. Confusing these two quantities is called the prosecutor's fallacy and is discussed in Section 2.8. If instead we had defined B to be the event that the second card is a heart, then the two conditional probabilities would have been equal.
2. Both $P(A | B)$ and $P(B | A)$ make sense (intuitively and mathematically); the chronological order in which cards were chosen does not dictate which conditional probabilities we can look at. When we calculate conditional probabilities, we are considering what information observing one event provides about another event, not whether one event causes another. For further intuition, imagine that someone spreads out the cards and draws one card with their left hand and another card with their right hand, at the same time. Defining A and B based on the left hand's card and the right hand's card rather than the first card and second card would not change the structure of the problem in any important way.
3. We can also see that $P(B | A) = 25/51$ by a direct interpretation of what conditional probability means: if the first card drawn is a heart, then the remaining cards consist of 25 red cards and 26 black cards (all of which are equally likely to be drawn next), so the conditional probability of getting a red card is $25/(25 + 26) = 25/51$. It is harder to find $P(A | B)$ in this way: if we learn that the second card is red, we might think "that's nice to know, but what we really want to know is whether it's a heart!" The conditional probability results from later sections in this chapter give us methods for getting around this issue.

To shed more light on what conditional probability means, here are two intuitive interpretations.

Intuition 2.2.3 (Pebble World). Consider a finite sample space, with the outcomes visualized as pebbles with total mass 1. Since A is an event, it is a set of pebbles, and likewise for B . Figure 2.1(a) shows an example.




FIGURE 2.1

Pebble World intuition for $P(A | B)$. From left to right: (a) Events A and B are subsets of the sample space. (b) Because we know B occurred, get rid of the outcomes in B^c . (c) In the restricted sample space, renormalize so the total mass is still 1.

Now suppose that we learn that B occurred. In Figure 2.1(b), upon obtaining this information, we get rid of all the pebbles in B^c because they are incompatible with the knowledge that B has occurred. Then $P(A \cap B)$ is the total mass of the pebbles remaining in A . Finally, in Figure 2.1(c), we renormalize, that is, divide all the masses by a constant so that the new total mass of the remaining pebbles is 1. This is achieved by dividing by $P(B)$, the total mass of the pebbles in B . The updated mass of the outcomes corresponding to event A is the conditional probability $P(A | B) = P(A \cap B) / P(B)$.

In this way, our probabilities have been updated in accordance with the observed evidence. Outcomes that contradict the evidence are discarded, and their mass is redistributed among the remaining outcomes, preserving the relative masses of the remaining outcomes. For example, if pebble 2 weighs twice as much as pebble 1 initially, and both are contained in B , then after conditioning on B it is still true that pebble 2 weighs twice as much as pebble 1. But if pebble 2 is not contained in B , then after conditioning on B its mass is updated to 0. $\square$

Intuition 2.2.4 (Frequentist interpretation). Recall that the frequentist interpretation of probability is based on relative frequency over a large number of repeated trials. Imagine repeating our experiment many times, generating a long list of observed outcomes. The conditional probability of A given B can then be thought of in a natural way: it is the fraction of times that A occurs, restricting attention to the trials where B occurs.

In Figure 2.2, our experiment has outcomes which can be written as a string of 0's and 1's; B is the event that the first digit is 1 and A is the event that the second digit is 1. Conditioning on B , we circle all the repetitions where B occurred, and then we look at the fraction of circled repetitions in which event A also occurred.

In symbols, let n_A, n_B, n_{AB} be the number of occurrences of $A, B, A \cap B$ respectively in a large number n of repetitions of the experiment. The frequentist interpretation is that

$$P(A) \approx \frac{n_A}{n}, P(B) \approx \frac{n_B}{n}, P(A \cap B) \approx \frac{n_{AB}}{n}.$$

Then $P(A | B)$ is interpreted as n_{AB} / n_B , which equals $(n_{AB} / n) / (n_B / n)$. This interpretation again translates to $P(A | B) = P(A \cap B) / P(B)$. $\square$




FIGURE 2.2 Frequentist intuition for $P(A \mid B)$. The repetitions where B occurred are circled; among these, the repetitions where A occurred are highlighted in bold. $P(A \mid B)$ is the long-run relative frequency of the repetitions where A occurs, within the subset of repetitions where B occurs.

For practice with applying the definition of conditional probability, let's do some more examples. The next three examples all start with the same basic scenario of a family with two children, but subtleties arise depending on the exact assumptions and the exact information we condition on.

Example 2.2.5 (Two children). Martin Gardner posed the following puzzle in the 1950s, in his column in *Scientific American*.

Mr. Jones has two children. The older child is a girl. What is the probability that both children are girls?

Mr. Smith has two children. At least one of them is a boy. What is the probability that both children are boys?

At first glance this problem seems like it should be a simple application of conditional probability, but for decades there have been controversies about whether or why the two parts of the problem should have different answers, and the extent to which the problem is ambiguous. Gardner gave the answers $1/2$ and $1/3$ to the two parts, respectively, which may seem paradoxical: why should it matter whether we learn the older child's gender, as opposed to just learning one child's gender?

It is important to clarify the assumptions of the problem. Several implicit assumptions are being made to obtain the answers that Gardner gave.

- It assumes that gender is binary, so that each child can be definitively categorized as a boy or a girl. In fact, many people don't neatly fit into either of the categories "male" or "female", and identify themselves as having a non-binary gender.
- It assumes that $P(\text{boy}) = P(\text{girl})$, both for the elder child and for the younger child. In fact, in most countries slightly more boys are born than girls. For example, in the United States it is commonly estimated that 105 boys are born for every 100 girls who are born.
- It assumes that the genders of the two children are independent, i.e., knowing the elder child's gender gives no information about the younger child's gender, and vice versa. This would be unrealistic if, e.g., the children were identical twins.

Under these (admittedly problematic) simplifying assumptions, we can solve the problem as follows.

Solution:

With the assumptions listed above, the definition of conditional probability gives

$$P(\text{both girls} \mid \text{elder is a girl}) = \frac{P(\text{both girls, elder is a girl})}{P(\text{elder is a girl})} = \frac{1/4}{1/2} = 1/2,$$

$$P(\text{both girls at least one girl}) = \frac{P(\text{both girls, at least one girl})}{P(\text{at least one girl})} = \frac{1/4}{3/4} = 1/3.$$

(We solved the second part of the problem in terms of girls rather than boys to make it a bit easier to compare the two parts of the problem.) It may seem counterintuitive that the two results are different, since there is no reason for us to care whether the elder child is a girl as opposed to the younger child. Indeed, by symmetry,

$$P(\text{both girls} \mid \text{younger is a girl}) = P(\text{both girls} \mid \text{elder is a girl}) = 1/2.$$

However, there is no such symmetry between the conditional probabilities $P(\text{both girls} \mid \text{elder is a girl})$ and $P(\text{both girls at least one girl})$. Saying that the elder child is a girl designates a specific child, and then the other child (the younger child) has a 50% chance of being a girl. "At least one" does not refer to a specific child. Conditioning on a specific child being a girl knocks away 2 of the 4 "pebbles" in the sample space $\{GG, GB, BG, BB\}$, where, for example, GB means the elder child is a girl and the younger child is a boy. In contrast, conditioning on at least one child being a girl knocks away only BB .

Example 2.2.6 (Random child is a girl). A family has two children. You randomly run into one of the two, and learn that she is a girl. With assumptions as in the previous example, what is the conditional probability that both are girls? Also assume that you are equally likely to run into either child, and that which one you run into has nothing to do with gender.

Solution:

Intuitively, the answer should be $1/2$: imagine that the child we encountered is in front of us and the other is at home. Both being girls just says that the child who is at home is a girl, which seems to have nothing to do with the fact that the child in front of us is a girl. But let us check this more carefully, using the definition of conditional probability. This is also good practice with writing events in set notation.

Let G_1, G_2 , and G_3 be the events that the elder, younger, and random child is a girl, respectively. By assumption, $P(G_1) = P(G_2) = P(G_3) = 1/2$. By then naive definition, or by independence as explained in Section 2.5, $P(G_1 \cap G_2) = 1/4$. Thus,

$$P(G_1 \cap G_2 \mid G_3) = P(G_1 \cap G_2 \cap G_3) / P(G_3) = (1/4) / (1/2) = 1/2,$$

since $G_1 \cap G_2 \cap G_3 = G_1 \cap G_2$ (if both children are girls, it guarantees that the random child is a girl).

Keep in mind though that in order to arrive at $1/2$, an assumption was needed about how the random child was selected. In statistical language, we say that we collected a random sample; here the sample consists of one of the two children. One of the most important principles in statistics is that it is essential to think carefully about how the sample was collected, not just stare at the raw data without understanding where they came from. To take a simple extreme

case, suppose that a repressive law forbids a boy from leaving the house if he has a sister. Then "the random child is a girl" is equivalent to "at least one of the children is a girl", so the problem reduces to the first part of Example 2.2.5.

Example 2.2.7 (A girl born in winter). A family has two children. Find the probability that both children are girls, given that at least one of the two is a girl who was born in winter. In addition to the assumptions from Example 2.2.5, assume that the four seasons are equally likely and that gender is independent of season. (This means that knowing the gender gives no information about the probabilities of the seasons, and vice versa; see Section 2.5 for much more about independence.)

Solution:

By definition of conditional probability,

$$P(\text{both girls} \mid \text{at least one winter girl}) = \frac{P(\text{both girls, at least one winter girl})}{P(\text{at least one winter girl})}.$$

Since the probability that a specific child is a winter-born girl is $1/8$, the denominator equals

$$P(\text{at least one winter girl}) = 1 - (7/8)^2.$$

To compute the numerator, use the fact that "both girls, at least one winter girl" is the same event as "both girls, at least one winter child"; then use the assumption that gender and season are independent:

$$\begin{aligned} P(\text{both girls, at least one winter girl}) &= P(\text{both girls, at least one winter child}) \\ &= (1/4)P(\text{at least one winter child}) \\ &= (1/4)(1 - P(\text{both are non-winter})) \\ &= (1/4)(1 - (3/4)^2) \end{aligned}$$

Thus,

$$P(\text{both girls} \mid \text{at least one winter girl}) = \frac{(1/4)(1 - (3/4)^2)}{1 - (7/8)^2} = \frac{7/64}{15/64} = 7/15.$$

At first this result seems absurd! In Example 2.2.5, the result was that the conditional probability of both children being girls, given that at least one is a girl, is $1/3$; why should it be any different when we learn that at least one is a winter-born girl? The point is that information about the birth season brings "at least one is a girl" closer to "a specific one is a girl". Conditioning on more and more specific information brings the probability closer and closer to $1/2$.

For example, conditioning on "at least one is a girl who was born on a March 31 at 8:20 pm" comes very close to specifying a child, and learning information about a specific child does not give us information about the other child. The seemingly irrelevant information such as season of birth interpolates between the two parts of Example 2.2.5. Exercise 29 generalizes this example to an arbitrary characteristic that is independent of gender.

4.2.3 2.3 Bayes' Rule and the law of Total Probability

The definition of conditional probability is simple - just a ratio of two probabilities but it has far-reaching consequences. The first consequence is obtained easily by moving the denominator in the definition to the other side of the equation.

Theorem 2.3.1 (Probability of the intersection of two events). For any events A and B with positive probabilities,

$$P(A \cap B) = P(B)P(A | B) = P(A)P(B | A).$$

This follows from taking the definition of $P(A | B)$ and multiplying both sides by $P(B)$, and then taking the definition of $P(B | A)$ and multiplying both sides by $P(A)$. At first sight this theorem may not seem very useful: it is the definition of conditional probability, just written slightly differently, and anyway it seems circular to use $P(A | B)$ to help find $P(A \cap B)$ when $P(A | B)$ was defined in terms of $P(A \cap B)$. But we will see that the theorem is in fact very useful, since it often turns out to be possible to find conditional probabilities without going back to the definition, and in such cases Theorem 2.3.1 can help us more easily find $P(A \cap B)$.

Applying Theorem 2.3.1 repeatedly, we can generalize to the intersection of n events.

Theorem 2.3.2 (Probability of the intersection of n events). For any events $A_1, \dots, A_n$ with $P(A_1, A_2, \dots, A_{n-1}) > 0$,

$$P(A_1, A_2, \dots, A_n) = P(A_1)P(A_2 | A_1)P(A_3 | A_1, A_2) \cdots P(A_n | A_1, \dots, A_{n-1}),$$

The commas denote intersections, e.g., $P(A_3 | A_1, A_2)$ is $P(A_3 | A_1 \cap A_2)$.

In fact, this is $n!$ theorems in one, since we can permute $A_1, \dots, A_n$ however we want without affecting the left-hand side. Often the right-hand side will be much easier to compute for some orderings than for others. For example,

$$P(A_1, A_2, A_3) = P(A_1)P(A_2 | A_1)P(A_3 | A_1, A_2) = P(A_2)P(A_3 | A_2)P(A_1 | A_2, A_3),$$

and there are 4 other expansions of this form too. It often takes practice and thought to be able to know which ordering to use.

We are now ready to introduce the two main theorems of this chapter-Bayes' rule and the law of total probability-which will allow us to compute conditional probabilities in a wide range of problems. Bayes' rule is an extremely famous, extremely useful result that relates $P(A | B)$ to $P(B | A)$.

Theorem 2.3.3 (Bayes' rule).

$$P(A | B) = \frac{P(B | A)P(A)}{P(B)}.$$

This follows immediately from Theorem 2.3.1, which in turn followed immediately from the definition of conditional probability. Yet Bayes' rule has important implications and applications in probability and statistics, since it is so often necessary to find conditional probabilities, and often $P(B | A)$ is much easier to find directly than $P(A | B)$ (or vice versa).

Another way to write Bayes' rule is in terms of odds rather than probability.

Definition 2.3.4 (Odds). The odds of an event A are

$$\text{odds}(A) = P(A)/P(A^c).$$

For example, if $P(A) = 2/3$, we say the odds in favor of A are 2 to 1. (This is sometimes written as 2 : 1, and is sometimes stated as 1 to 2 odds against A ; care is needed since some

sources do not explicitly state whether they are referring to odds in favor or odds against an event.) Of course we can also convert from odds back to probability:

$$P(A) = \text{odds}(A)/(1 + \text{odds}(A)).$$

By taking the Bayes' rule expression for $P(A | B)$ and dividing it by the Bayes' rule expression for $P(A^c | B)$, we arrive at the odds form of Bayes' rule.

Theorem 2.3.5 (Odds form of Bayes' rule). For any events A and B with positive probabilities, the odds of A after conditioning on B are

$$\frac{P(A | B)}{P(A^c | B)} = \frac{P(B | A)}{P(B | A^c)} \frac{P(A)}{P(A^c)}.$$

In words, this says that the posterior odds $P(A | B)/P(A^c | B)$ are equal to the prior odds $P(A)/P(A^c)$ times the factor $P(B | A)/P(B | A^c)$, which is known in statistics as the likelihood ratio. Sometimes it is convenient to work with this form of Bayes' rule to get the posterior odds, and then if desired we can convert from odds back to probability.

The law of total probability (LOTP) relates conditional probability to unconditional probability. It is essential for fulfilling the promise that conditional probability can be used to decompose complicated probability problems into simpler pieces, and it is often used in tandem with Bayes' rule.

Theorem 2.3.6 (Law of total probability). Let $A_1, \dots, A_n$ be a partition of the sample space S (i.e., the A_i are disjoint events and their union is S), with $P(A_i) > 0$ for all i . Then

$$P(B) = \sum_{i=1}^n P(B | A_i) P(A_i).$$

Proof. Since the A_i form a partition of S , we can decompose B as

$$B = (B \cap A_1) \cup (B \cap A_2) \cup \dots \cup (B \cap A_n).$$

This is illustrated in Figure 2.3, where we have chopped B into the smaller pieces $B \cap A_1$ through $B \cap A_n$. By the second axiom of probability, because these pieces are disjoint, we can add their probabilities to get $P(B)$:

$$P(B) = P(B \cap A_1) + P(B \cap A_2) + \dots + P(B \cap A_n).$$

Now we can apply Theorem 2.3.1 to each of the $P(B \cap A_i)$:

$$P(B) = P(B | A_1) P(A_1) + \dots + P(B | A_n) P(A_n).$$

The law of total probability tells us that to get the unconditional probability of B , we can divide the sample space into disjoint slices A_i , find the conditional probability of B within each of the slices, then take a weighted sum of the conditional probabilities, where the weights are the probabilities $P(A_i)$. The choice of how to divide up the sample space is crucial: a well-chosen partition will reduce a complicated problem into simpler pieces, whereas a poorly chosen partition will only exacerbate our problems, requiring us to calculate n difficult probabilities instead of just one!

The next few examples show how we can use Bayes' rule together with the law of total probability to update our beliefs based on observed evidence.




FIGURE 2.3

The $\{A_i\}$

partition the sample space; $P(B)$ is equal to $\sum_i P(B \cap A_i)$.

Example 2.3.7 (Random coin). You have one fair coin, and one biased coin which lands Heads with probability $3/4$. You pick one of the coins at random and flip it three times. It lands Heads all three times. Given this information, what is the probability that the coin you picked is the fair one?

Solution:

Let A be the event that the chosen coin lands Heads three times and let F be the event that we picked the fair coin. We are interested in $P(F | A)$, but it is easier to find $P(A | F)$ and $P(A | F^c)$ since it helps to know which coin we have; this suggests using Bayes' rule and the law of total probability. Doing so, we have

$$\begin{aligned} P(F | A) &= \frac{P(A | F)P(F)}{P(A)} \\ &= \frac{P(A | F)P(F)}{P(A | F)P(F) + P(A | F^c)P(F^c)} \\ &= \frac{(1/2)^3 \cdot 1/2}{(1/2)^3 \cdot 1/2 + (3/4)^3 \cdot 1/2} \\ &\approx 0.23 \end{aligned}$$

Before flipping the coin, we thought we were equally likely to have picked the fair coin as the biased coin: $P(F) = P(F^c) = 1/2$. Upon observing three Heads, however, it becomes more likely that we've chosen the biased coin than the fair coin, so $P(F | A)$ is only about 0.23 .

2.3.8 (Prior vs. posterior). It would not be correct in the calculation in the above example to say after the first step, " $P(A) = 1$ because we know A happened." It is true that $P(A | A) = 1$, but $P(A)$ is the prior probability of A and $P(F)$ is the prior probability of F -both are the probabilities before we observe any data in the experiment. These must not be confused with posterior probabilities conditional on the evidence A .

Example 2.3.9 (Testing for a rare disease). A patient named Fred is tested for a disease called conditionitis, a medical condition that afflicts 1% of the population. The test result is positive, i.e., the test claims that Fred has the disease. Let D be the event that Fred has the disease and T be the event that he tests positive.

Suppose that the test is "95% accurate"; there are different measures of the accuracy of a test, but in this problem it is assumed to mean that $P(T \mid D) = 0.95$ and $P(T^c \mid D^c) = 0.95$. The quantity $P(T \mid D)$ is known as the sensitivity or true positive rate of the test, and $P(T^c \mid D^c)$ is known as the specificity or true negative rate.

Find the conditional probability that Fred has conditionitis, given the evidence provided by the test result.

Solution:

Applying Bayes' rule and the law of total probability, we have

$$\begin{aligned} P(D \mid T) &= \frac{P(T \mid D)P(D)}{P(T)} \\ &= \frac{P(T \mid D)P(D)}{P(T \mid D)P(D) + P(T \mid D^c)P(D^c)} \\ &= \frac{0.95 \cdot 0.01}{0.95 \cdot 0.01 + 0.05 \cdot 0.99} \\ &\approx 0.16 \end{aligned}$$

So there is only a 16% chance that Fred has conditionitis, given that he tested positive, even though the test seems to be quite reliable!

Most people find it surprising to learn that the conditional probability of having the disease given a positive test result is only 16%, even though the test is 95% accurate (see Gigerenzer and Hoffrage [13]). The key to understanding this surprisingly low posterior probability of having the disease is to realize that there are two factors at play: the evidence from the test, and our prior information about the prevalence of the disease.

Although the test provides evidence in favor of disease, conditionitis is also a rare condition! The conditional probability $P(D \mid T)$ reflects a balance between these two factors, appropriately weighing the rarity of the disease against the rarity of a mistaken test result.

For further intuition, consider a population of 10000 people as illustrated in Figure 2.4, where 100 have conditionitis and 9900 don't; this corresponds to a 1% disease rate. If we tested everybody in the population, we'd expect that out of the 100 diseased individuals, 95 would test positive and 5 would test negative. Out of the 9900 healthy individuals, we'd expect $(0.95)(9900) \approx 9405$ to test negative and 495 to test positive.




FIGURE 2.4

Testing for a rare disease in a population of 10000 people, where the prevalence of the disease is 1% and the true positive and true negative rates are both equal to 95%. Bubbles are not to scale.

Now let's focus in on those individuals who test positive; that is, let's condition on a positive test result. The 95 true positives (i.e., the individuals who test positive and have the disease) are far outnumbered by the 495 false positives (i.e., the individuals who test positive despite not having the disease). So most people who test positive for the disease don't actually have the disease!

Example 2.3.10 (Six-fingered man). A crime has been committed in a certain country. The perpetrator is one (and only one) of the n men who live in the country. Initially, these n men are all deemed equally likely to be the perpetrator. An eyewitness then reports that the crime was committed by a man with six fingers on his right hand.

Let p_0 be the probability that an innocent man has six fingers on his right hand, and p_1 be the probability that the perpetrator has six fingers on his right hand, with $p_0 < p_1$. (We may have $p_1 < 1$, since eyewitnesses are not 100% reliable.) Let $a = p_0/p_1$ and $b = (1 - p_1) / (1 - p_0)$.

Rugen lives in the country. He is found to have six fingers on his right hand.

- Given this information, what is the probability that Rugen is the perpetrator?
- Now suppose that all n men who live in the country have their hands checked, and Rugen is the only one with six fingers on his right hand. Given this information, what is the probability that Rugen is the perpetrator?

Solution:

(a) Let R be the event that Rugen is guilty and M be the event that he has six fingers on his right hand. By Bayes' rule and LOTP,

$$P(R | M) = \frac{P(M|R)P(R)}{P(M|R)P(R) + P(M|R^c)P(R^c)} = \frac{p_1 \cdot \frac{1}{n}}{p_1 \cdot \frac{1}{n} + p_0 \left(1 - \frac{1}{n}\right)} = \frac{1}{1 + a(n-1)}.$$

(b) Let N be the event that none of the men in the country other than Rugen have six fingers on their right hands. With notation as above,

$$\begin{aligned}
P(R \mid M, N) &= \frac{P(M, N \mid R)P(R)}{P(M, N \mid R)P(R) + P(M, N \mid R^c)P(R^c)} \\
&= \frac{p_1(1-p_0)^{n-1} \cdot \frac{1}{n}}{p_1(1-p_0)^{n-1} \cdot \frac{1}{n} + p_0(1-p_1)(1-p_0)^{n-2}(1-\frac{1}{n})} \\
&= \frac{1}{1 + ab(n-1)}
\end{aligned}$$

4.2.4 2.4 Conditional Probabilities are Probabilities

When we condition on an event E , we update our beliefs to be consistent with this knowledge, effectively putting ourselves in a universe where we know that E occurred. Within our new universe, however, the laws of probability operate just as before. Conditional probability satisfies all the properties of probability! Therefore, any of the results we have derived about probability are still valid if we replace all unconditional probabilities with probabilities conditional on E . In particular:

- Conditional probabilities are between 0 and 1 .
- $P(S \mid E) = 1, P(\emptyset \mid E) = 0$.
- If $A_1, A_2, \dots$ are disjoint, then $P(\cup_{j=1}^{\infty} A_j \mid E) = \sum_{j=1}^{\infty} P(A_j \mid E)$.
- $P(A^c \mid E) = 1 - P(A \mid E)$.
- Inclusion-exclusion: $P(A \cup B \mid E) = P(A \mid E) + P(B \mid E) - P(A \cap B \mid E)$.

2.4.1. When we write $P(A \mid E)$, it does not mean that $A \mid E$ is an event and we're taking its probability; $A \mid E$ is not an event. Rather, $P(\cdot \mid E)$ is a probability function which assigns probabilities in accordance with the knowledge that E has occurred, and $P(\cdot)$ is a different probability function which assigns probabilities without regard for whether E has occurred or not. When we take an event A and plug it into the $P(\cdot)$ function, we'll get a number, $P(A)$; when we plug it into the $P(\cdot \mid E)$ function, we'll get another number, $P(A \mid E)$, which incorporates the information (if any) provided by knowing that E occurred.

To prove mathematically that conditional probabilities are probabilities, fix an event E with $P(E) > 0$, and for any event A , define $\tilde{P}(A) = P(A \mid E)$. This notation helps emphasize the fact that we are fixing E and treating $P(\cdot \mid E)$ as our new probability function. We just need to check the two axioms of probability. First,

$$\tilde{P}(\emptyset) = P(\emptyset \mid E) = \frac{P(\emptyset \cap E)}{P(E)} = 0, \tilde{P}(S) = P(S \mid E) = \frac{P(S \cap E)}{P(E)} = 1$$

Second, if $A_1, A_2, \dots$ are disjoint events, then

$$\tilde{P}(A_1 \cup A_2 \cup \dots) = \frac{P((A_1 \cap E) \cup (A_2 \cap E) \cup \dots)}{P(E)} = \frac{\sum_{j=1}^{\infty} P(A_j \cap E)}{P(E)} = \sum_{j=1}^{\infty} \tilde{P}(A_j).$$

So $\tilde{P}$ satisfies the axioms of probability.

Conversely, all probabilities can be thought of as conditional probabilities: whenever we make a probability statement, there is always some background information that we are conditioning on, even if we don't state it explicitly. Consider the rain example from the beginning of this chapter. It would be natural to base the initial probability of rain today, $P(R)$, on the fraction of days in the past on which it rained. But which days in the past should we look at? If it's November 1, should we only count past rainy days in

autumn, thus conditioning on the season? What about conditioning on the exact month, or the exact day? We could ask the same about location: should we look at days when it rained in our exact location, or is it enough for it to have rained somewhere nearby?

In order to determine the seemingly unconditional probability $P(R)$, we actually have to make decisions about what background information to condition on! These choices require careful thought and different people may come up with different prior probabilities $P(R)$ (though everyone can agree on how to update based on new evidence).

Since all probabilities are conditional on background information, we can imagine that there is always a vertical conditioning bar, with background knowledge K to the right of the vertical bar. Then the unconditional probability $P(A)$ is just shorthand for $P(A | K)$; the background knowledge is absorbed into the letter P instead of being written explicitly.

To summarize our discussion in a nutshell:

Conditional probabilities are probabilities, and all probabilities are conditional.

We now state conditional forms of Bayes' rule and the law of total probability. These are obtained by taking the ordinary forms of Bayes' rule and LOTP and adding E to the right of the vertical bar everywhere.

Theorem 2.4.2 (Bayes' rule with extra conditioning). Provided that $P(A \cap E) > 0$ and $P(B \cap E) > 0$, we have

$$P(A | B, E) = \frac{P(B | A, E)P(A | E)}{P(B | E)}$$

Theorem 2.4.3 (LOTP with extra conditioning). Let $A_1, \dots, A_n$ be a partition of S . Provided that $P(A_i \cap E) > 0$ for all i , we have

$$P(B | E) = \sum_{i=1}^n P(B | A_i, E) P(A_i | E).$$

The extra conditioning forms of Bayes' rule and LOTP can be proved similarly to how we verified that $\tilde{P}$ satisfies the axioms of probability, but they also follow directly from the "metatheorem" that conditional probabilities are probabilities.

Example 2.4.4 (Random coin, continued). Continuing with the scenario from Example 2.3.7, suppose that we have now seen our chosen coin land Heads three times. If we toss the coin a fourth time, what is the probability that it will land Heads once more?

Solution:

As before, let A be the event that the chosen coin lands Heads three times, and define a new event H for the chosen coin landing Heads on the fourth toss. We are interested in $P(H | A)$. It would be very helpful to know whether we have the fair coin. LOTP with extra conditioning gives us $P(H | A)$ as a weighted average of $P(H | F, A)$ and $P(H | F^c, A)$, and within these two conditional probabilities we do know whether we have the fair coin:

$$\begin{aligned} P(H | A) &= P(H | F, A)P(F | A) + P(H | F^c, A)P(F^c | A) \\ &\approx \frac{1}{2} \cdot 0.23 + \frac{3}{4} \cdot (1 - 0.23) \\ &\approx 0.69. \end{aligned}$$

The posterior probabilities $P(F | A)$ and $P(F^c | A)$ are from our answer to Example 2.3.7.

An equivalent way to solve this problem is to define a new probability function $\tilde{P}$ such that for any event B , $\tilde{P}(B) = P(B | A)$. This new function assigns probabilities that are updated with the knowledge that A occurred. Then by the ordinary law of total probability,

$$\tilde{P}(H) = \tilde{P}(H | F)\tilde{P}(F) + \tilde{P}(H | F^c)\tilde{P}(F^c)$$

which is exactly the same as our use of LOTP with extra conditioning. This once again illustrates the principle that conditional probabilities are probabilities.

Example 2.4.5 (Unanimous agreement). The article "Why too much evidence can be a bad thing" by Lisa Zyga [30] says:

Under ancient Jewish law, if a suspect on trial was unanimously found guilty by all judges, then the suspect was acquitted. This reasoning sounds counterintuitive, but the legislators of the time had noticed that unanimous agreement often indicates the presence of systemic error in the judicial process.

There are n judges deciding a case. The suspect has prior probability p of being guilty. Each judge votes whether to convict or acquit the suspect. With probability s , a systemic error occurs (e.g., the defense is incompetent). If a systemic error occurs, then the judges unanimously vote to convict (i.e., all n judges vote to convict).

Whether a systemic error occurs is independent of whether the suspect is guilty. Given that a systemic error doesn't occur and that the suspect is guilty, each judge has probability c of voting to convict, independently. Given that a systemic error doesn't occur and that the suspect is not guilty, each judge has probability w of voting to convict, independently. Suppose that

$$0 < p < 1, 0 < s < 1, \text{ and } 0 < w < \frac{1}{2} < c < 1.$$

- (a) For this part only, suppose that exactly k out of n judges vote to convict, where $k < n$. Given this information, find the probability that the suspect is guilty.
- (b) Now suppose that all n judges vote to convict. Given this information, find the probability that the suspect is guilty.
- (c) Is the answer to (b), viewed as a function of n , an increasing function? Give a short, intuitive explanation in words.

Solution:

(a) Since $k < n$, a systemic error didn't occur. We will implicitly condition on this in this part. Let G be the event that the suspect is guilty and X be the number of judges who vote to convict. Using Bayes' rule, LOTP, and the Binomial PMF,

$$P(G | X = k) = \frac{P(X = k | G)P(G)}{P(X = k)} = \frac{pc^k(1-c)^{n-k}}{pc^k(1-c)^{n-k} + (1-p)w^k(1-w)^{n-k}}.$$

(b) Let U be the event $X = n$ and B be the event that a systemic error occurs. Then

$$P(G | U) = \frac{P(U | G)P(G)}{P(U)} = \frac{pP(U | G)}{pP(U | G) + (1-p)P(U | G^c)}.$$

By LOTP with extra conditioning,

$$\begin{aligned} P(U | G) &= P(U | G, B)P(B | G) + P(U | G, B^c)P(B^c | G) = s + (1-s)c^n \\ P(U | G^c) &= P(U | G^c, B)P(B | G^c) + P(U | G^c, B^c)P(B^c | G^c) = s + (1-s)w^n \end{aligned}$$

Thus,

$$P(G | U) = \frac{p(s + (1 - s)c^n)}{p(s + (1 - s)c^n) + (1 - p)(s + (1 - s)w^n)}$$

(c) No, since a large value of n yields a high chance of systemic error, and if a systemic error occurs then the judges' votes are uninformative about whether the suspect is guilty. The answer to (b) reverts to the prior probability p as $n \rightarrow \infty$.

We often want to condition on more than one piece of information, and we now have several ways of doing that. For example, here are some approaches for finding $P(A | B, C)$:

1. We can think of B, C as the single event $B \cap C$ and use the definition of conditional probability to get

$$P(A | B, C) = \frac{P(A, B, C)}{P(B, C)}$$

This is a natural approach if it's easiest to think about B and C in tandem. We can then try to evaluate the numerator and denominator. For example, we can use LOTP in both the numerator and the denominator, or we can write the numerator as $P(B, C | A)P(A)$ (which would give us a version of Bayes' rule) and use LOTP to help with the denominator.

2. We can use Bayes' rule with extra conditioning on C to get

$$P(A | B, C) = \frac{P(B | A, C)P(A | C)}{P(B | C)}$$

This is a natural approach if we want to think of everything in our problem as being conditioned on C .

3. We can use Bayes' rule with extra conditioning on B to get

$$P(A | B, C) = \frac{P(C | A, B)P(A | B)}{P(C | B)}.$$

This is the same as the previous approach, except with the roles of B and C swapped. We mention it separately just to emphasize that it's a bad idea to plug into a formula without thinking about which event should play which role.

It is both challenging and powerful that there are a variety of ways to approach this kind of conditioning problem.

4.2.5 2.5 Independence of Events

We have now seen several examples where conditioning on one event changes our beliefs about the probability of another event. The situation where events provide no information about each other is called independence.

Definition 2.5.1 (Independence of two events). Events A and B are independent if

$$P(A \cap B) = P(A)P(B).$$

If $P(A) > 0$ and $P(B) > 0$, then this is equivalent to

$$P(A | B) = P(A)$$

and also equivalent to $P(B | A) = P(B)$.

In words, two events are independent if we can obtain the probability of their intersection by multiplying their individual probabilities. Alternatively, A and B are independent if learning

that B occurred gives us no information that would change our probabilities for A occurring (and vice versa).

Note that independence is a symmetric relation: if A is independent of B , then B is independent of A .

2.5.2. Independence is completely different from disjointness. If A and B are disjoint, then $P(A \cap B) = 0$, so disjoint events can be independent only if $P(A) = 0$ or $P(B) = 0$. Knowing that A occurs tells us that B definitely did not occur, so A clearly conveys information about B , meaning the two events are not independent (except if A or B already has zero probability).

Intuitively, it makes sense that if A provides no information about whether or not B occurred, then it also provides no information about whether or not B^c occurred. We now prove a handy result along those lines.

Proposition 2.5.3. If A and B are independent, then A and B^c are independent, A^c and B are independent, and A^c and B^c are independent.

Proof. Let A and B be independent. We will first show that A and B^c are independent. If $P(A) = 0$, then A is independent of every event, including B^c . So assume $P(A) \neq 0$. Then

$$P(B^c | A) = 1 - P(B | A) = 1 - P(B) = P(B^c),$$

so A and B^c are independent. Swapping the roles of A and B , we have that A^c and B are independent. Using the fact that A, B independent implies A, B^c independent, with A^c playing the role of A , we also have that A^c and B^c are independent.

We also often need to talk about independence of three or more events.

Definition 2.5.4 (Independence of three events). Events A, B , and C are said to be independent if all of the following equations hold:

$$\begin{aligned} P(A \cap B) &= P(A)P(B), \\ P(A \cap C) &= P(A)P(C), \\ P(B \cap C) &= P(B)P(C), \\ P(A \cap B \cap C) &= P(A)P(B)P(C). \end{aligned}$$

If the first three conditions hold, we say that A, B , and C are pairwise independent. Pairwise independence does not imply independence: it is possible that just learning about A or just learning about B is of no use in predicting whether C occurred, but learning that both A and B occurred could still be highly relevant for C . Here is a simple example of this distinction.

Example 2.5.5 (Pairwise independence doesn't imply independence). Consider two fair, independent coin tosses, and let A be the event that the first is Heads, B the event that the second is Heads, and C the event that both tosses have the same result. Then A, B , and C are pairwise independent but not independent, since $P(A \cap B \cap C) = 1/4$ while $P(A)P(B)P(C) = 1/8$. The point is that just knowing about A or just knowing about B tells us nothing about C , but knowing what happened with both A and B gives us information about C (in fact, in this case it gives us perfect information about C).

On the other hand, $P(A \cap B \cap C) = P(A)P(B)P(C)$ does not imply pairwise independence; this can be seen quickly by looking at the extreme case $P(A) = 0$, when the equation becomes $0 = 0$, which tells us nothing about B and C .

We can define independence of any number of events similarly. Intuitively, the idea is that knowing what happened with any particular subset of the events gives us no information about what happened with the events not in that subset.

Definition 2.5.6 (Independence of many events). For n events $A_1, A_2, \dots, A_n$ to be independent, we require any pair to satisfy $P(A_i \cap A_j) = P(A_i)P(A_j)$ (for $i \neq j$), any triplet to satisfy $P(A_i \cap A_j \cap A_k) = P(A_i)P(A_j)P(A_k)$ (for i, j, k distinct), and similarly for all quadruplets, quintuplets, and so on. This can quickly become unwieldy, but later we will discuss other ways

to think about independence. For infinitely many events, we say that they are independent if every finite subset of the events is independent.

Conditional independence is defined analogously to independence.

Definition 2.5.7 (Conditional independence). Events A and B are said to be conditionally independent given E if $P(A \cap B \mid E) = P(A \mid E)P(B \mid E)$.

2.5.8. It is easy to make terrible blunders stemming from confusing independence and conditional independence. Two events can be conditionally independent given E , but not independent given E^c . Two events can be conditionally independent given E , but not independent. Two events can be independent, but not conditionally independent given E .

In particular, $P(A, B) = P(A)P(B)$ does not imply $P(A, B \mid E) = P(A \mid E)P(B \mid E)$; we can't just insert "given E " everywhere, as we did in going from LOTP to LOTP with extra conditioning. This is because LOTP always holds (it is a consequence of the axioms of probability), whereas $P(A, B)$ may or may not equal $P(A)P(B)$, depending on what A and B are.

The next few examples illustrate these distinctions. Great care is needed in working with conditional probabilities and conditional independence!

Example 2.5.9 (Conditional independence given E vs. given E^c). Suppose there are two types of classes: good classes and bad classes. In good classes, if you work hard, you are very likely to get an A. In bad classes, the professor randomly assigns grades to students regardless of their effort. Let G be the event that a class is good, W be the event that you work hard, and A be the event that you receive an A. Then W and A are conditionally independent given G^c , but they are not conditionally independent given G .

Example 2.5.10 (Conditional independence doesn't imply independence). Returning once more to the scenario from Example 2.3.7, suppose we have chosen either a fair coin or a biased coin with probability $3/4$ of Heads, but we do not know which one we have chosen. We flip the coin a number of times. Conditional on choosing the fair coin, the coin tosses are independent, with each toss having probability $1/2$ of Heads. Similarly, conditional on choosing the biased coin, the tosses are independent, each with probability $3/4$ of Heads.

However, the coin tosses are not unconditionally independent, because if we don't know which coin we've chosen, then observing the sequence of tosses gives us information about whether we have the fair coin or the biased coin in our hand. This in turn helps us to predict the outcomes of future tosses from the same coin.

To state this formally, let F be the event that we've chosen the fair coin, and let A_1 and A_2 be the events that the first and second coin tosses land Heads. Conditional on F , A_1 and A_2 are independent, but A_1 and A_2 are not unconditionally independent because A_1 provides information about A_2 . $\square$

Example 2.5.11 (Independence doesn't imply conditional independence). My friends Alice and Bob are the only two people who ever call me on the phone. Each day, they decide independently whether to call me that day. Let A be the event that Alice calls me next Friday and B be the event that Bob calls me next Friday. Assume A and B are unconditionally independent with $P(A) > 0$ and $P(B) > 0$.

However, given that I receive exactly one call next Friday, A and B are no longer independent: the call is from Alice if and only if it is not from Bob. In other words, letting C be the event that I receive exactly one call next Friday, $P(B \mid C) > 0$ while $P(B \mid A, C) = 0$, so A and B are not conditionally independent given C . $\square$

Example 2.5.12. (Why is the baby crying?) A certain baby cries if and only if she is hungry, tired, or both. Let C be the event that the baby is crying, H be the event that she is hungry, and T be the event that she is tired. Let $P(C) = c$, $P(H) = h$, and $P(T) = t$, where none of c, h, t are equal to 0 or 1. Let H and T be independent.

(a) Find c , in terms of h and t .

- (b) Find $P(H | C)$, $P(T | C)$, and $P(H, T | C)$.
 (c) Are H and T conditionally independent given C ? Explain in two ways: algebraically using the quantities from (b), and with an intuitive explanation in words.

Solution:

- (a) Since H and T are independent, we have

$$P(C) = P(H \cup T) = P(H) + P(T) - P(H \cap T) = h + t - ht.$$

- (b) By Bayes' rule,

$$\begin{aligned} P(H | C) &= \frac{P(C | H)P(H)}{P(C)} = \frac{h}{c} \\ P(T | C) &= \frac{P(C | T)P(T)}{P(C)} = \frac{t}{c} \\ P(H, T | C) &= \frac{P(C | H, T)P(H, T)}{P(C)} = \frac{ht}{c} \end{aligned}$$

- (c) No, H and T are not conditionally independent given C , since

$$P(H, T | C) = \frac{ht}{c} < \frac{ht}{c^2} = P(H | C)P(T | C).$$

We can also see intuitively why they are not conditionally independent given C : if the baby is crying but not hungry, she must be tired. $\square$

4.2.6 2.6 Coherency of Bayes' Rule

An important property of Bayes' rule is that it is coherent: if we receive multiple pieces of information and wish to update our probabilities to incorporate all the information, it does not matter whether we update sequentially, taking each piece of evidence into account one at a time, or simultaneously, using all the evidence at once. Suppose, for example, that we're conducting a weeklong experiment that yields data at the end of each day. We could use Bayes' rule every day to update our probabilities based on the data from that day. Or we could go on vacation for the week, come back on Friday afternoon, and update using the entire week's worth of data. Either method will give the same result.

Let's look at a concrete application of this principle.

Example 2.6.1 (Testing for a rare disease, continued). Fred, who tested positive for conditionitis in Example 2.3.9, decides to get tested a second time. The new test is independent of the original test (given his disease status) and has the same sensitivity and specificity. Unfortunately for Fred, he tests positive a second time. Find the probability that Fred has the disease, given the evidence, in two ways: in one step, conditioning on both test results simultaneously, and in two steps, first updating the probabilities based on the first test result, and then updating again based on the second test result.

Solution:

Let D be the event that he has the disease, T_1 that the first test result is positive, and T_2 that the second test result is positive. In Example 2.3.9, we used Bayes' rule and the law of total probability to find $P(D | T_1)$. Another quick solution uses the odds form of Bayes' rule:

$$\frac{P(D | T_1)}{P(D^c | T_1)} = \frac{P(D)}{P(D^c)} \frac{P(T_1 | D)}{P(T_1 | D^c)} = \frac{1}{99} \cdot \frac{0.95}{0.05} \approx 0.19.$$

Since $P(D | T_1) / (1 - P(D | T_1)) = 0.19$, we have $P(D | T_1) = 0.19 / (1 + 0.19) \approx 0.16$, in agreement with our answer from before. The odds form of Bayes' rule is faster in this case because we don't need to compute the unconditional probability $P(T_1)$ in the denominator of the ordinary Bayes' rule. Now, again using the odds form of Bayes' rule, let's find out what happens if Fred tests positive a second time.

One-step method: Updating based on both test results at once, we have

$$\begin{aligned} \frac{P(D | T_1 \cap T_2)}{P(D^c | T_1 \cap T_2)} &= \frac{P(D)}{P(D^c)} \frac{P(T_1 \cap T_2 | D)}{P(T_1 \cap T_2 | D^c)} \\ &= \frac{1}{99} \cdot \frac{0.95^2}{0.05^2} = \frac{361}{99} \approx 3.646, \end{aligned}$$

which corresponds to a probability of 0.78 .

Two-step method: After the first test, the posterior odds of Fred having the disease are

$$\frac{P(D | T_1)}{P(D^c | T_1)} = \frac{1}{99} \cdot \frac{0.95}{0.05} \approx 0.19,$$

from the above. These posterior odds become the new prior odds, and then updating based on the second test gives

$$\begin{aligned} \frac{P(D | T_1 \cap T_2)}{P(D^c | T_1 \cap T_2)} &= \frac{P(D | T_1)}{P(D^c | T_1)} \frac{P(T_2 | D, T_1)}{P(T_2 | D^c, T_1)} \\ &= \left(\frac{1}{99} \cdot \frac{0.95}{0.05} \right) \frac{0.95}{0.05} = \frac{361}{99} \approx 3.646, \end{aligned}$$

which is the same result as above.

Note that with a second positive test result, the probability that Fred has the disease jumps from 0.16 to 0.78 , making us much more confident that Fred is actually afflicted with conditionitis. The moral of the story is that getting a second opinion is a good idea! $\square$


4.2.7 2.7 Conditioning as a Problem-solving Tool

Conditioning is a powerful tool for solving problems because it lets us engage in wishful thinking: when we encounter a problem that would be made easier if only we knew whether E happened or not, we can condition on E and then on E^c , consider these possibilities separately, then combine them using LOTP.

2.7.1 Strategy: Condition on What you Wish you Knew

Example 2.7.1 (Monty Hall). On the game show Let's Make a Deal, hosted by Monty Hall, a contestant chooses one of three closed doors, two of which have a goat behind them and one of which has a car. Monty, who knows where the car is, then opens one of the two remaining doors. The door he opens always has a goat behind it (he never reveals the car!). If he has a choice, then he picks a door at random with equal probabilities. Monty then offers the contestant the option of switching to the other unopened door. If the contestant's goal is to get the car, should

she switch doors?



Solution:

Let's label the doors 1 through 3. Without loss of generality, we can assume the contestant picked door 1 (if she didn't pick door 1, we could simply relabel the doors, or rewrite this solution with the door numbers permuted). Monty opens a door, revealing a goat. As the contestant decides whether or not to switch to the remaining unopened door, what does she really wish she knew? Naturally, her decision would be a lot easier if she knew where the car was! This suggests that we should condition on the location of the car. Let C_i be the event that the car is behind door i , for $i = 1, 2, 3$. By the law of total probability,

$$P(\text{get car}) = P(\text{get car} \mid C_1) \cdot \frac{1}{3} + P(\text{get car} \mid C_2) \cdot \frac{1}{3} + P(\text{get car} \mid C_3) \cdot \frac{1}{3}.$$

Suppose the contestant employs the switching strategy. If the car is behind door 1, then switching will fail, so $P(\text{get car} \mid C_1) = 0$. If the car is behind door 2 or 3, then because Monty always reveals a goat, the remaining unopened door must contain the car, so switching will succeed. Thus,

$$P(\text{get car}) = 0 \cdot \frac{1}{3} + 1 \cdot \frac{1}{3} + 1 \cdot \frac{1}{3} = \frac{2}{3},$$

so the switching strategy succeeds $2/3$ of the time. The contestant should switch to the other door.

Figure 2.5 is a tree diagram of the argument we have just outlined: using the switching strategy, the contestant will win as long as the car is behind doors 2 or 3, which has probability $2/3$. We can also give an intuitive frequentist argument in favor of switching. Imagine playing this game 1000 times. Typically, about 333 times your initial guess for the car's location will be correct, in which case switching will fail. The other 667 or so times, you will win by switching.

There's a subtlety though, which is that when the contestant chooses whether to switch, she also knows which door Monty opened. We showed that the unconditional probability of success is $2/3$ (when following the switching strategy), but let's also show that the conditional probability of success for switching, given the information that Monty provides, is also $2/3$.

Let M_j be the event that Monty opens door j , for $j = 2, 3$. Then

$$P(\text{get car}) = P(\text{get car} \mid M_2) P(M_2) + P(\text{get car} \mid M_3) P(M_3),$$




FIGURE 2.5

Tree diagram of Monty Hall problem. Switching gets the car 2/3 of the time.

where by symmetry $P(M_2) = P(M_3) = 1/2$ and $P(\text{get car} | M_2) = P(\text{get car} | M_3)$. The symmetry here is that there is nothing in the statement of the problem that distinguishes between door 2 and door 3; in contrast, Problem 40 considers a scenario where Monty enjoys opening door 2 more than he enjoys opening door 3.

Let $x = P(\text{get car} | M_2) = P(\text{get car} | M_3)$. Plugging in what we know,

$$\frac{2}{3} = P(\text{get car}) = \frac{x}{2} + \frac{x}{2} = x,$$

as claimed.

Bayes' rule also works nicely for finding the conditional probability of success using the switching strategy, given the evidence. Suppose that Monty opens door 2. Using the notation and results above,

$$P(C_1 | M_2) = \frac{P(M_2 | C_1) P(C_1)}{P(M_2)} = \frac{(1/2)(1/3)}{1/2} = \frac{1}{3}.$$

So given that Monty opens door 2, there is a 1/3 chance that the contestant's original choice of door has the car, which means that there is a 2/3 chance that the switching strategy will succeed.

Many people, upon seeing this problem for the first time, argue that there is no advantage to switching: "There are two doors remaining, and one of them has the car, so the chances are 50-50." After the last chapter, we recognize that this argument misapplies the naive definition of probability. Yet the naive definition, even when inappropriate, has a powerful hold on people's intuitions. When Marilyn vos Savant presented a correct solution to the Monty Hall problem in her column

for Parade magazine in 1990, she received thousands upon thousands of letters from readers (even mathematicians) insisting that she was wrong.

To build correct intuition, let's consider an extreme case. Suppose that there are a million doors, 999,999 of which contain goats and 1 of which has a car. After the contestant's initial pick, Monty opens 999,998 doors with goats behind them and offers the choice to switch. In this extreme case, it becomes clear that the probabilities are not 50 – 50 for the two unopened doors; very few people would stubbornly stick with their original choice. The same is true for the three-door case.


Just as we had to make assumptions about how we came across the random girl in Example 2.2.6, here the 2/3 success rate of the switching strategy depends on the assumptions we make

about how Monty decides which door to open. In the exercises, we consider several variants and generalizations of the Monty Hall problem, some of which change the desirability of the switching strategy. $\square$

2.7.2 Strategy: Condition on the First Step

In problems with a recursive structure, it can often be useful to condition on the first step of the experiment. The next two examples apply this strategy, which we call first-step analysis.

Example 2.7.2 (Branching process). A single amoeba, Bobo, lives in a pond. After one minute Bobo will either die, split into two amoebas, or stay the same, with equal probability, and in subsequent minutes all living amoebas will behave the same way, independently. What is the probability that the amoeba population will eventually die out?



Solution:

Let D be the event that the population eventually dies out; we want to find $P(D)$. We proceed by conditioning on the outcome at the first step: let B_i be the event that Bobo turns into i amoebas after the first minute, for $i = 0, 1, 2$. We know $P(D | B_0) = 1$ and $P(D | B_1) = P(D)$ (if Bobo stays the same, we're back to where we started). If Bobo splits into two, then we just have two independent versions of our original problem! We need both of the offspring to eventually die out, so $P(D | B_2) = P(D)^2$. Now we have exhausted all the possible cases and can combine them with the law of total probability:

$$\begin{aligned} P(D) &= P(D | B_0) \cdot \frac{1}{3} + P(D | B_1) \cdot \frac{1}{3} + P(D | B_2) \cdot \frac{1}{3} \\ &= 1 \cdot \frac{1}{3} + P(D) \cdot \frac{1}{3} + P(D)^2 \cdot \frac{1}{3} \end{aligned}$$


Solving for $P(D)$ gives $P(D) = 1$: the amoeba population will die out with probability 1.

The strategy of first-step analysis works here because the problem is self-similar in nature: when Bobo continues as a single amoeba or splits into two, we end up with another version or another two versions of our original problem. Conditioning on the first step allows us to express $P(D)$ in terms of itself.

Example 2.7.3 (Gambler's ruin). Two gamblers, A and B, make a sequence of \$1 bets. In each bet, gambler A has probability p of winning, and gambler B has probability $q = 1 - p$ of winning. Gambler A starts with i dollars and gambler B starts with $N - i$ dollars; the total

wealth between the two remains constant since every time A loses a dollar, the dollar goes to B, and vice versa.

We can visualize this game as a random walk on the integers between 0 and N , where p is the probability of going to the right in a given step: imagine a person who starts at position i and, at each time step, moves one step to the right with probability p and one step to the left with probability $q = 1 - p$. The game ends when either A or B is ruined, i.e., when the random walk reaches 0 or N . What is the probability that A wins the game (walking away with all the money)?



Solution:

We recognize that this game, like Bobo's reproductive process, has a recursive structure: after the first step, it's exactly the same game, except that A's wealth is now either $i + 1$ or $i - 1$. Let p_i be the probability that A wins the game, given that A starts with i dollars. We will use first-step analysis to solve for the p_i . Let W be the event that A wins the game. By LOTP, conditioning on the outcome of the first round, we have

$$\begin{aligned} p_i &= P(W \mid \text{A starts at } i, \text{ wins round 1}) \cdot p + P(W \mid \text{A starts at } i, \text{ loses round 1}) \cdot q \\ &= P(W \mid \text{A starts at } i + 1) \cdot p + P(W \mid \text{A starts at } i - 1) \cdot q \\ &= p_{i+1} \cdot p + p_{i-1} \cdot q \end{aligned}$$

This must be true for all i from 1 to $N - 1$, and we also have the boundary conditions $p_0 = 0$ and $p_N = 1$. Now we can solve this equation, called a difference equation, to obtain the p_i . Section A. 4 of the math appendix discusses how to solve difference equations, so we will omit some of the steps here.

The characteristic equation of the difference equation is $px^2 - x + q = 0$, which has roots 1 and q/p . If $p \neq 1/2$, these roots are distinct, and the general solution is

$$p_i = a \cdot 1^i + b \cdot \left(\frac{q}{p}\right)^i$$

Using the boundary conditions $p_0 = 0$ and $p_N = 1$, we get

$$a = -b = \frac{1}{1 - \left(\frac{q}{p}\right)^N}$$

and we simply plug these back in to get the specific solution. If $p = 1/2$, the roots of the characteristic polynomial are not distinct, so the general solution is

$$p_i = a \cdot 1^i + b \cdot i \cdot 1^i$$

The boundary conditions give $a = 0$ and $b = 1/N$.

In summary, the probability of A winning with a starting wealth of i is

$$p_i = \begin{cases} \frac{1 - \left(\frac{q}{p}\right)^i}{1 - \left(\frac{q}{p}\right)^N} & \text{if } p \neq 1/2 \\ \frac{i}{N} & \text{if } p = 1/2 \end{cases}$$

The $p = 1/2$ case is consistent with the $p \neq 1/2$ case, in the sense that

$$\lim_{p \rightarrow \frac{1}{2}} \frac{1 - \left(\frac{q}{p}\right)^i}{1 - \left(\frac{q}{p}\right)^N} = \frac{i}{N}$$

To see this, let $x = q/p$ and let x approach 1. By L'Hôpital's rule,

$$\lim_{x \rightarrow 1} \frac{1 - x^i}{1 - x^N} = \lim_{x \rightarrow 1} \frac{ix^{i-1}}{Nx^{N-1}} = \frac{i}{N}$$

The answer for the $p = 1/2$ case has a simple interpretation: A's probability of winning equals the proportion of the wealth that A starts out with. So if $p = 1/2$ and A starts out with much less money than B, then A's chance of winning the game is low. Having $p < 1/2$ may also make A's chance of winning low, even if p is only a little bit less than $1/2$ and the players start out with the same amount of money. For example, if $p = 0.49$ and each player starts out with \$100, then A has only about a 1.8% chance of winning the game.

We have focused on the probability of A winning the game, but what about B ? Rather than starting from scratch, we can use symmetry: aside from notation, there is nothing in the description of the game to distinguish A from B . By symmetry, the probability of B winning from a starting wealth of $N - i$ is obtained by switching the roles of q and p , and of i and $N - i$. This gives

$$P(\text{B wins} \mid \text{B starts at } N - i) = \begin{cases} \frac{1 - \left(\frac{p}{q}\right)^{N-i}}{1 - \left(\frac{p}{q}\right)^N} & \text{if } p \neq 1/2, \\ \frac{N-i}{N} & \text{if } p = 1/2. \end{cases}$$

It can then be verified that for all i and all p , $P(\text{A wins}) + P(\text{B wins}) = 1$, so the game is guaranteed to end: the probability is 0 that it will oscillate forever.

4.2.8 2.8 Pitfalls and Paradoxes

The next two examples are fallacies of conditional thinking that have arisen in the legal context. The prosecutor's fallacy is the confusion of $P(A \mid B)$ with $P(B \mid A)$; the defense attorney's fallacy is the failure to condition on all the evidence.

2.8.1 (Prosecutor's fallacy). In 1998, Sally Clark was tried for murder after two of her sons died shortly after birth. During the trial, an expert witness for the prosecution testified that the probability of a newborn dying of sudden infant death syndrome (SIDS) was $1/8500$, so the probability of two deaths due to SIDS in one family was $(1/8500)^2$, or about one in 73 million. Therefore, he continued, the probability of Clark's innocence was one in 73 million.

There are at least two major problems with this line of reasoning. First, the expert witness found the probability of the intersection of "first son dies of SIDS" and "second son dies of SIDS" by multiplying the individual event probabilities; as we know, this is only valid if deaths due to SIDS are independent within a family. This independence would not hold if genetic or other family-specific risk factors cause all newborns within certain families to be at increased risk of SIDS.

Second, the so-called expert has confused two different conditional probabilities: $P(\text{innocence}|\text{evidence})$ is different from $P(\text{evidence}|\text{innocence})$. The witness claims that the probability of observing two newborn deaths if the defendant were innocent is extremely low; that is, $P(\text{evidence}|\text{innocence})$ is small. What we are interested in, however, is $P(\text{innocence}|\text{evidence})$, the probability that the defendant is innocent given all the evidence. By Bayes' rule,

$$P(\text{innocence} \mid \text{evidence}) = \frac{P(\text{evidence} \mid \text{innocence})P(\text{innocence})}{P(\text{evidence})},$$

so to calculate the conditional probability of innocence given the evidence, we must take into account $P(\text{innocence})$, the prior probability of innocence. This probability is extremely high: although double deaths due to SIDS are rare, so are double infanticides! Expanding the denominator as

$$P(\text{evidence} \mid \text{innocence})P(\text{innocence}) + P(\text{evidence} \mid \text{guilt})P(\text{guilt}),$$

note that if $P(\text{guilt})$ is small enough so that the second term is negligible compared to the first term, then the denominator of $P(\text{innocence}|\text{evidence})$ is approximately equal to the numerator, making $P(\text{innocence}|\text{evidence})$ close to 1.

The posterior probability of innocence given the evidence depends strongly on both $P(\text{evidence}|\text{innocence})$, which is very low, and $P(\text{innocence})$, which is very high. The expert's probability of $(1/8500)^2$, questionable in and of itself, is only part of the equation.

Sadly, Clark was convicted of murder and sent to prison, partly based on the expert's wrongheaded testimony, and spent over three years in jail before her conviction was overturned. The outcry over the misuse of conditional probability in the Sally Clark case led to the review of hundreds of other cases where similar faulty logic was used by the prosecution.

2.8.2 (Defense attorney's fallacy). A woman has been murdered, and her husband is put on trial for this crime. Evidence comes to light that the defendant had a history of abusing his wife. The defense attorney argues that the evidence of abuse should be excluded on grounds of irrelevance, since only 1 in 10,000 men with wives they abuse subsequently murder their wives. Should the judge grant the defense attorney's motion to bar this evidence from trial?

Suppose that the defense attorney's 1-in-10,000 figure is correct, and further assume the following for a relevant population of husbands and wives: 1 in 10 husbands abuse their wives, 1 in 5 murdered wives were murdered by their husbands, and 50% of husbands who murder their wives previously abused them. Also, assume that if the husband of a murdered wife is not guilty of the murder, then the probability that he abused his wife reverts to the unconditional probability of abuse.

How to define the "relevant population" and how to estimate such probabilities are difficult issues. For example, should we look at citywide, statewide, national, or international statistics? How should we account for unreported abuse and unsolved murders? What if murder rates are changing over time? For this problem, assume that a reasonable choice of the relevant population has been agreed on, and that the stated probabilities are known to be correct.

Let A be the event that the husband commits abuse against his wife, and let G be the event that the husband is guilty. The defense's argument is that $P(G \mid A) = 1/10,000$, so guilt is still extremely unlikely conditional on a previous history of abuse.

However, the defense attorney fails to condition on a crucial fact: in this case, we know that the wife was murdered. Therefore, the relevant probability is not $P(G \mid A)$, but $P(G \mid A, M)$, where M is the event that the wife was murdered.

Bayes' rule with extra conditioning gives

$$\begin{aligned}P(G \mid A, M) &= \frac{P(A \mid G, M)P(G \mid M)}{P(A \mid G, M)P(G \mid M) + P(A \mid G^c, M)P(G^c \mid M)} \\&= \frac{0.5 \cdot 0.2}{0.5 \cdot 0.2 + 0.1 \cdot 0.8} \\&= \frac{5}{9}\end{aligned}$$

So the posterior probability of guilt, $P(G \mid A, M)$, is over 5,000 times as large as the quantity $P(G \mid A)$ that the defense attorney focused on. Conditioning on the evidence of abuse increases the probability of guilt from $P(G \mid M) = 0.2$ to $P(G \mid A, M) \approx 0.56$, so the defendant's history of abuse gives very important information, contrary to the defense attorney's argument.

In the above calculation of $P(G \mid A, M)$, we did not use the defense attorney's $P(G \mid A)$ number anywhere; it is irrelevant to our calculation because it does not account for the fact that the wife was murdered. We must condition on all the evidence.

We end this chapter with a paradox about conditional probability and aggregation of data.

Example 2.8.3 (Simpson's paradox). Two doctors, Dr. Hibbert and Dr. Nick, each perform two types of surgeries: heart surgery and Band-Aid removal. Each surgery can be either a success or a failure. The two doctors' respective records are given in the following tables, and shown graphically in Figure 2.6, where white dots represent successful surgeries and black dots represent failed surgeries.

	Heart	Band-Aid
Success	70	10
Failure	20	0

Dr. Hibbert

	Heart	Band-Aid
Success	2	81
Failure	8	9

Dr. Nick

Dr. Hibbert had a higher success rate than Dr. Nick in heart surgeries: 70 out of 90 versus 2 out of 10. Dr. Hibbert also had a higher success rate in Band-Aid removal: 10 out of 10 versus 81 out of 90. But if we aggregate across the two types of surgeries to compare overall surgery success rates, Dr. Hibbert was successful in 80 out of 100 surgeries while Dr. Nick was successful in 83 out of 100 surgeries: Dr. Nick's overall success rate is higher!




FIGURE 2.6

An example of Simpson's paradox. White dots represent successful surgeries and black dots represent failed surgeries. Dr. Hibbert is better in both types of surgery but has a lower overall success rate, because he is performing the harder type of surgery much more often than Dr. Nick is.

What's happening is that Dr. Hibbert, presumably due to his reputation as the superior doctor, is performing a greater number of heart surgeries, which are inherently riskier than Band-Aid removals. His overall success rate is lower not because of lesser skill on any particular type of surgery, but because a larger fraction of his surgeries are risky.

Let's use event notation to make this precise. For events A , B , and C , we say that we have a Simpson's paradox if

$$P(A | B, C) < P(A | B^c, C)$$

$$P(A | B, C^c) < P(A | B^c, C^c),$$

but

$$P(A | B) > P(A | B^c).$$

In this case, let A be the event of a successful surgery, B be the event that Dr. Nick is the surgeon, and C be the event that the surgery is a heart surgery. The conditions for Simpson's paradox are fulfilled because the probability of a successful surgery is lower under Dr. Nick than under Dr. Hibbert whether we condition on heart surgery or on Band-Aid removal, but the overall probability of success is higher for Dr. Nick.

The law of total probability tells us mathematically why this can happen:

$$P(A | B) = P(A | C, B)P(C | B) + P(A | C^c, B)P(C^c | B)$$

$$P(A | B^c) = P(A | C, B^c)P(C | B^c) + P(A | C^c, B^c)P(C^c | B^c).$$

The above equations express $P(A | B)$ as a weighted average of $P(A | C, B)$ and $P(A | C^c, B)$, and $P(A | B^c)$ as a weighted average of $P(A | C, B^c)$ and $P(A | C^c, B^c)$. If the corresponding weights were the same in both of these weighted averages, then Simpson's paradox could not occur. But the weights here are different:

$$P(C | B) < P(C | B^c) \text{ and } P(C^c | B) > P(C^c | B^c)$$

since Dr. Nick is much less likely than Dr. Hibbert to be performing a heart surgery. Although we have

$$P(A \mid C, B) < P(A \mid C, B^c)$$

and

$$P(A \mid C^c, B) < P(A \mid C^c, B^c),$$

the fact that the weights are so different results in the inequality flipping when we do not condition on whether or not C occurred:

$$P(A \mid B) > P(A \mid B^c)$$

Numerically, the two weighted averages are

$$\begin{aligned} P(A \mid B) &= 0.83 = (2/10) \cdot 0.1 + (81/90) \cdot 0.9 \\ P(A \mid B^c) &= 0.80 = (70/90) \cdot 0.9 + (10/10) \cdot 0.1. \end{aligned}$$

The first equation (corresponding to Dr. Nick) puts much more weight on the second term (corresponding to the easier surgery) than does the second equation.

Aggregation across different types of surgeries presents a misleading picture of the doctors' abilities because we lose the information about which doctor tends to perform which type of surgery. When we think confounding variables like surgery type could be at play, we should examine the disaggregated data to see what is really going on.

Simpson's paradox arises in many real-world contexts. In the following examples, you should try to identify the events A , B , and C that create the paradox.

- Gender discrimination in college admissions: In the 1970s, men were significantly more likely than women to be admitted for graduate study at the University of California, Berkeley, leading to charges of gender discrimination. Yet within most individual departments, women were admitted at a higher rate than men. It was found that women tended to apply to the departments with more competitive admissions, while men tended to apply to less competitive departments.
- Baseball batting averages: It is possible for player 1 to have a higher batting average than player 2 in the first half of a baseball season and a higher batting average than player 2 in the second half of the season, yet have a lower overall batting average for the entire season. It depends on how many at-bats the players have in each half of the season. (An at-bat is when it's a player's turn to try to hit the ball; the player's batting average is the number of hits the player gets divided by the player's number of at-bats.)
- Health effects of smoking: Cochran [4] found that within any age group, cigarette smokers had higher mortality rates than cigar smokers, but because cigarette smokers were on average younger than cigar smokers, overall mortality rates were lower for cigarette smokers.

4.2.9 2.9 Recap

The conditional probability of A given B is

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)}$$

Conditional probability has exactly the same properties as probability, but $P(\cdot \mid B)$ updates our uncertainty about events to reflect the observed evidence B . Events whose probabilities

are unchanged after observing the evidence B are said to be independent of B . Two events can also be conditionally independent given a third event E . Conditional independence does not imply unconditional independence, nor does unconditional independence imply conditional independence.

Two important results about conditional probability are Bayes' rule, which relates $P(A | B)$ to $P(B | A)$, and the law of total probability, which allows us to get unconditional probabilities by partitioning the sample space and calculating conditional probabilities within each slice of the partition.

Bayes' rule says that

$$P(A | B) = \frac{P(B | A)P(A)}{P(B)}$$

while LOTP says that

$$P(B) = \sum_{i=1}^n P(B | A_i) P(A_i)$$

for any partition $A_1, \dots, A_n$ of the sample space. Bayes' rule and LOTP are often used in tandem.

Conditioning is extremely helpful for problem-solving because it allows us to break a problem into smaller pieces, consider all possible cases separately, and then combine them. When using this strategy, we should try to condition on the information that, if known, would make the problem simpler, hence the saying condition on what you wish you knew. When a problem involves multiple stages, it can be helpful to condition on the first step to obtain a recursive relationship.

Common mistakes in thinking conditionally include:

- confusion of the prior probability $P(A)$ with the posterior probability $P(A | B)$;
- the prosecutor's fallacy, confusing $P(A | B)$ with $P(B | A)$;
- the defense attorney's fallacy, failing to condition on all the evidence;
- unawareness of Simpson's paradox and the importance of thinking carefully about whether to aggregate data.

Figure 2.7 illustrates how probabilities can be updated as new evidence comes in sequentially. Imagine that there is some event A that we are interested in. On Monday morning, for example, our prior probability for A is $P(A)$. If we observe on Monday afternoon that B occurred, then we can use Bayes' rule (or the definition of conditional probability) to compute the posterior probability $P(A | B)$.

We use this posterior probability for A as the new prior on Tuesday morning, and then we continue to collect evidence. Suppose that on Tuesday we observe that C occurred. Then we can compute the new posterior probability $P(A | B, C)$ in various ways (in this context, probably the most natural way is to use Bayes' rule with extra conditioning on B). This in turn becomes the new prior if we are going to continue to collect evidence.

4.2.10 2.10 R

Simulating the Frequentist Interpretation

Recall that the frequentist interpretation of conditional probability based on a large number n of repetitions of an experiment is $P(A | B) \approx n_{AB}/n_B$, where n_{AB} is the number of times that

$A \cap B$ occurs and n_B is the number of times that B occurs. Let's try this out by simulation, and verify the results of Example 2.2.5. So let's simulate n families, each with two children.

```
n <- 10^5
child1 <- sample(2,n,replace=TRUE)
child2 <- sample(2,n,replace=TRUE)
```

Here `child1` is a vector of length n , where each element is a 1 or a 2. Letting 1 stand for "girl" and 2 stand for "boy", this vector represents the gender of the elder child in each of the n families. Similarly, `child2` represents the gender of the younger child in each family.

Alternatively, we could have used

```
sample(c("girl","boy"),n,replace=TRUE)
```

but it is more convenient working with numerical values.

Let A be the event that both children are girls and B the event that the elder is a girl.




FIGURE 2.7

Conditional probability tells us how to update probabilities as new evidence comes in. Shown are the probabilities for an event A initially, after obtaining one piece of evidence B , and after obtaining a second piece of evidence C . The posterior for A after observing the first piece of evidence becomes the new prior before observing the second piece of evidence. After both B and C are observed, a new posterior for A can be found in various ways. This then becomes the new prior if more evidence will be collected.

Following the frequentist interpretation, we count the number of repetitions where B occurred and name it `n.b`, and we also count the number of repetitions where $A \cap B$ occurred and name it `n.ab`. Finally, we divide `n.ab` by `n.b` to approximate $P(A | B)$.

```
n.b <- sum(child1==1)
n.ab <- sum(child1==1 & child2==1)
n.ab/n.b
```

The ampersand & is an elementwise AND, so `n.ab` is the number of families where both the first child and the second child are girls. When we ran this code, we got 0.50 , which agrees with $P(\text{both girls} \mid \text{elder is a girl}) = 1/2$.

Now let A be the event that both children are girls and B the event that at least one of the children is a girl. Then $A \cap B$ is the same, but `n.b` needs to count the number of families where at least one child is a girl. This is accomplished with the elementwise OR operator `|` (this is not a conditioning bar; it is an inclusive OR, returning TRUE if at least one element is TRUE).

```
n.b <- sum(child1==1 | child2==1)
n.ab <- sum(child1==1 & child2==1)
n.ab/n.b
```

We got 0.33 , which agrees with $P(\text{both girls} \mid \text{at least one girl}) = 1/3$.

Monty Hall Simulation

Many long, bitter debates about the Monty Hall problem could have been averted by trying it out with a simulation. To study how well the never-switch strategy performs, let's generate 10^5 runs of the Monty Hall game. To simplify notation, assume the contestant always chooses door 1 . Then we can generate a vector specifying which door has the car for each repetition:

```
n <- 10^5
cardoor <- sample(3,n,replace=TRUE)
```

At this point we could generate the vector specifying which doors Monty opens, but that's unnecessary since the never-switch strategy succeeds if and only if door 1 has the car! So the fraction of times when the never-switch strategy succeeds is `sum(cardoor == 1)/n`, which was 0.334 in our simulation. This is very close to the true value, $1/3$.

What if we want to play the Monty Hall game interactively? We can do this by programming a function. Entering the following code in R defines a function called `monty`, which can then be invoked by entering the command `monty()` any time you feel like playing the game!

```
monty <- function() {
  doors <- 1:3
  # randomly pick where the car is
  cardoor <- sample(doors,1)
  # prompt player
  print("Monty Hall says 'Pick a door, any door!'")
  # receive the player's choice of door (should be 1,2, or 3)
  chosen <- scan(what = integer(), nlines = 1, quiet = TRUE)
  # pick Monty's door (can't be the player's door or the car door)
  if (chosen != cardoor) montydoor <- doors[-c(chosen, cardoor)]
  else montydoor <- sample(doors[-chosen],1)
  # find out whether the player wants to switch doors
  print(paste("Monty opens door ", montydoor, "!", sep=""))
  print("Would you like to switch (y/n)?")
  reply <- scan(what = character(), nlines = 1, quiet = TRUE)
  # interpret what player wrote as "yes" if it starts with "y"
  if (substr(reply,1,1) == "y"){
    chosen <- doors[-c(chosen,montydoor)]
  }
}
```

```
# announce the result of the game!
if (chosen == cardoor) print("You won!")
else print("You lost!")
}
```

The print command prints its argument to the screen. We combine this with the paste command since if we used `print("Monty opens door montydoor")` then R would literally print "Monty opens door montydoor". The scan command interactively requests input from the user; we use `what = integer()` when we want the user to enter an integer and `what = character()` when we want the user to enter text. Using `substr(reply, 1,1)` extracts the first character of reply, in case the user replies with "yes" or "yep" or "yeah!" rather than with "y".

4.2.11 2.11 Exercises

Exercises marked with (S) have detailed solutions at <http://stat110.net>.

Conditioning on Evidence

1. (S) A spam filter is designed by looking at commonly occurring phrases in spam. Suppose that 80% of email is spam. In 10% of the spam emails, the phrase "free money" is used, whereas this phrase is only used in 1% of non-spam emails. A new email has just arrived, which does mention "free money". What is the probability that it is spam?
2. (S) A woman is pregnant with twin boys. Twins may be either identical or fraternal. Suppose that $1/3$ of twins born are identical, that identical twins have a 50% chance of being both boys and a 50% chance of being both girls, and that for fraternal twins each twin independently has a 50% chance of being a boy and a 50% chance of being a girl. Given the above information, what is the probability that the woman's twins are identical?
3. According to the CDC (Centers for Disease Control and Prevention), men who smoke are 23 times more likely to develop lung cancer than men who don't smoke. Also according to the CDC, 21.6% of men in the U.S. smoke. What is the probability that a man in the U.S. is a smoker, given that he develops lung cancer?
4. Fred is answering a multiple-choice problem on an exam, and has to choose one of n options (exactly one of which is correct). Let K be the event that he knows the answer, and R be the event that he gets the problem right (either through knowledge or through luck). Suppose that if he knows the right answer he will definitely get the problem right, but if he does not know then he will guess completely randomly. Let $P(K) = p$.
 - (a) Find $P(K | R)$ (in terms of p and n).
 - (b) Show that $P(K | R) \geq p$, and explain why this makes sense intuitively. When (if ever) does $P(K | R)$ equal p ?
5. Three cards are dealt from a standard, well-shuffled deck. The first two cards are flipped over, revealing the Ace of Spades as the first card and the 8 of Clubs as the second card. Given this information, find the probability that the third card is an ace in two ways: using the definition of conditional probability, and by symmetry.
6. A hat contains 100 coins, where 99 are fair but one is double-headed (always landing Heads). A coin is chosen uniformly at random. The chosen coin is flipped 7 times, and it lands Heads all 7 times. Given this information, what is the probability that the chosen

coin is double-headed? (Of course, another approach here would be to look at both sides of the coin-but this is a metaphorical coin.)

7. A hat contains 100 coins, where at least 99 are fair, but there may be one that is doubleheaded (always landing Heads); if there is no such coin, then all 100 are fair. Let D be the event that there is such a coin, and suppose that $P(D) = 1/2$. A coin is chosen uniformly at random. The chosen coin is flipped 7 times, and it lands Heads all 7 times.
 - (a) Given this information, what is the probability that one of the coins is doubleheaded?
 - (b) Given this information, what is the probability that the chosen coin is doubleheaded?
8. The screens used for a certain type of cell phone are manufactured by 3 companies, A, B, and C. The proportions of screens supplied by A, B, and C are 0.5, 0.3, and 0.2, respectively, and their screens are defective with probabilities 0.01, 0.02, and 0.03, respectively. Given that the screen on such a phone is defective, what is the probability that Company A manufactured it?
9.
 - (a) Show that if events A_1 and A_2 have the same prior probability $P(A_1) = P(A_2)$, A_1 implies B , and A_2 implies B , then A_1 and A_2 have the same posterior probability $P(A_1 | B) = P(A_2 | B)$ if it is observed that B occurred.
 - (b) Explain why (a) makes sense intuitively, and give a concrete example.
10. Fred is working on a major project. In planning the project, two milestones are set up, with dates by which they should be accomplished. This serves as a way to track Fred's progress. Let A_1 be the event that Fred completes the first milestone on time, A_2 be the event that he completes the second milestone on time, and A_3 be the event that he completes the project on time.

Suppose that $P(A_{j+1} | A_j) = 0.8$ but $P(A_{j+1} | A_j^c) = 0.3$ for $j = 1, 2$, since if Fred falls behind on his schedule it will be hard for him to get caught up. Also, assume that the second milestone supersedes the first, in the sense that once we know whether he is on time in completing the second milestone, it no longer matters what happened with the first milestone. We can express this by saying that A_1 and A_3 are conditionally independent given A_2 and they're also conditionally independent given A_2^c .

 - (a) Find the probability that Fred will finish the project on time, given that he completes the first milestone on time. Also find the probability that Fred will finish the project on time, given that he is late for the first milestone.
 - (b) Suppose that $P(A_1) = 0.75$. Find the probability that Fred will finish the project on time.
11. An exit poll in an election is a survey taken of voters just after they have voted. One major use of exit polls has been so that news organizations can try to figure out as soon as possible who won the election, before the votes are officially counted. This has been notoriously inaccurate in various elections, sometimes because of selection bias: the sample of people who are invited to and agree to participate in the survey may not be similar enough to the overall population of voters.

Consider an election with two candidates, Candidate A and Candidate B. Every voter is invited to participate in an exit poll, where they are asked whom they voted for; some accept and some refuse. For a randomly selected voter, let A be the event that they voted for A, and W be the event that they are willing to participate in the exit poll. Suppose that $P(W | A) = 0.7$ but $P(W | A^c) = 0.3$. In the exit poll, 60% of the respondents say they voted for A (assume that they are all honest), suggesting a comfortable victory for A. Find $P(A)$, the true proportion of people who voted for A.

12. Alice is trying to communicate with Bob, by sending a message (encoded in binary) across a channel.
- (a) Suppose for this part that she sends only one bit (a 0 or 1), with equal probabilities. If she sends a 0, there is a 5% chance of an error occurring, resulting in Bob receiving a 1; if she sends a 1, there is a 10% chance of an error occurring, resulting in Bob receiving a 0. Given that Bob receives a 1, what is the probability that Alice actually sent a 1?
- (b) To reduce the chance of miscommunication, Alice and Bob decide to use a repetition code. Again Alice wants to convey a 0 or a 1, but this time she repeats it two more times, so that she sends 000 to convey 0 and 111 to convey 1. Bob will decode the message by going with what the majority of the bits were. Assume that the error probabilities are as in (a), with error events for different bits independent of each other. Given that Bob receives 110, what is the probability that Alice intended to convey a 1?
13. Company A has just developed a diagnostic test for a certain disease. The disease afflicts 1% of the population. As defined in Example 2.3.9, the sensitivity of the test is the probability of someone testing positive, given that they have the disease, and the specificity of the test is the probability that of someone testing negative, given that they don't have the disease. Assume that, as in Example 2.3.9, the sensitivity and specificity are both 0.95.
- Company B, which is a rival of Company A, offers a competing test for the disease. Company B claims that their test is faster and less expensive to perform than Company A's test, is less painful (Company A's test requires an incision), and yet has a higher overall success rate, where overall success rate is defined as the probability that a random person gets diagnosed correctly.
- (a) It turns out that Company B's test can be described and performed very simply: no matter who the patient is, diagnose that they do not have the disease. Check whether Company B's claim about overall success rates is true.
- (b) Explain why Company A's test may still be useful.
- (c) Company A wants to develop a new test such that the overall success rate is higher than that of Company B's test. If the sensitivity and specificity are equal, how high does the sensitivity have to be to achieve their goal? If (amazingly) they can get the sensitivity equal to 1, how high does the specificity have to be to achieve their goal? If (amazingly) they can get the specificity equal to 1, how high does the sensitivity have to be to achieve their goal?
14. Consider the following scenario, from Tversky and Kahneman [27]:

Let A be the event that before the end of next year, Peter will have installed a burglar alarm system in his home. Let B denote the event that Peter's home will be burglarized before the end of next year.

- (a) Intuitively, which do you think is bigger, $P(A | B)$ or $P(A | B^c)$? Explain your intuition.
- (b) Intuitively, which do you think is bigger, $P(B | A)$ or $P(B | A^c)$? Explain your intuition.
- (c) Show that for any events A and B (with probabilities not equal to 0 or 1), the inequality $P(A | B) > P(A | B^c)$ is equivalent to $P(B | A) > P(B | A^c)$.
- (d) Tversky and Kahneman report that 131 out of 162 people whom they posed (a) and (b) to said that $P(A | B) > P(A | B^c)$ and $P(B | A) < P(B | A^c)$. What is a plausible explanation for why this was such a popular opinion despite (c) showing that it is impossible for these inequalities both to hold?
15. Let A and B be events with $0 < P(A \cap B) < P(A) < P(B) < P(A \cup B) < 1$. You are hoping that both A and B occurred. Which of the following pieces of information would you be happiest to observe: that A occurred, that B occurred, or that $A \cup B$ occurred?

16. Show that $P(A | B) \leq P(A)$ implies $P(A | B^c) \geq P(A)$, and give an intuitive explanation of why this makes sense.

17. In deterministic logic, the statement " A implies B " is equivalent to its contrapositive, "not B implies not A ". In this problem we will consider analogous statements in probability, the logic of uncertainty. Let A and B be events with probabilities not equal to 0 or 1.

(a) Show that if $P(B | A) = 1$, then $P(A^c | B^c) = 1$.

Hint: Apply Bayes' rule and LOTP.

(b) Show however that the result in (a) does not hold in general if $=$ is replaced by $\approx$. In particular, find an example where $P(B | A)$ is very close to 1 but $P(A^c | B^c)$ is very close to 0.

Hint: What happens if A and B are independent?

18. Show that if $P(A) = 1$, then $P(A | B) = 1$ for any B with $P(B) > 0$. Intuitively, this says that if someone dogmatically believes something with absolute certainty, then no amount of evidence will change their mind. The principle of avoiding assigning probabilities of 0 or 1 to any event (except for mathematical certainties) was named Cromwell's rule by the statistician Dennis Lindley, due to Cromwell saying to the Church of Scotland, "Think it possible you may be mistaken."

Hint: Write $P(B) = P(B \cap A) + P(B \cap A^c)$, and then show that $P(B \cap A^c) = 0$.

19. Explain the following Sherlock Holmes saying in terms of conditional probability, carefully distinguishing between prior and posterior probabilities: "It is an old maxim of mine that when you have excluded the impossible, whatever remains, however improbable, must be the truth."

20. The Jack of Spades (with cider), Jack of Hearts (with tarts), Queen of Spades (with a wink), and Queen of Hearts (without tarts) are taken from a deck of cards. These four cards are shuffled, and then two are dealt. Note: Literary references to cider, tarts, and winks do not need to be considered when solving this problem.

(a) Find the probability that both of these two cards are queens, given that the first card dealt is a queen.

(b) Find the probability that both are queens, given that at least one is a queen.

(c) Find the probability that both are queens, given that one is the Queen of Hearts.

21. A fair coin is flipped 3 times. The toss results are recorded on separate slips of paper (writing "H" if Heads and "T" if Tails), and the 3 slips of paper are thrown into a hat.

(a) Find the probability that all 3 tosses landed Heads, given that at least 2 were Heads.

(b) Two of the slips of paper are randomly drawn from the hat, and both show the letter H. Given this information, what is the probability that all 3 tosses landed Heads?

22. (S) A bag contains one marble which is either green or blue, with equal probabilities. A green marble is put in the bag (so there are 2 marbles now), and then a random marble is taken out. The marble taken out is green. What is the probability that the remaining marble is also green?

23. (S) Let G be the event that a certain individual is guilty of a certain robbery. In gathering evidence, it is learned that an event E_1 occurred, and a little later it is also learned that another event E_2 also occurred. Is it possible that individually, these pieces of evidence increase the chance of guilt (so $P(G | E_1) > P(G)$ and $P(G | E_2) > P(G)$), but together they decrease the chance of guilt (so $P(G | E_1, E_2) < P(G)$)?

24. Is it possible to have events A_1, A_2, B, C with $P(A_1 | B) > P(A_1 | C)$ and $P(A_2 | B) > P(A_2 | C)$, yet $P(A_1 \cup A_2 | B) < P(A_1 \cup A_2 | C)$? If so, find an example (with a "story" interpreting the events, as well as giving specific numbers); otherwise, show that it is impossible for this phenomenon to happen.

25. (S) A crime is committed by one of two suspects, A and B . Initially, there is equal evidence against both of them. In further investigation at the crime scene, it is found that the guilty party had a blood type found in 10% of the population. Suspect A does match this blood type,

whereas the blood type of Suspect B is unknown.

- (a) Given this new information, what is the probability that A is the guilty party?
- (b) Given this new information, what is the probability that B 's blood type matches that found at the crime scene?

26. (S) To battle against spam, Bob installs two anti-spam programs. An email arrives, which is either legitimate (event L) or spam (event L^c), and which program j marks as legitimate (event M_j) or marks as spam (event M_j^c) for $j \in \{1, 2\}$. Assume that 10% of Bob's email is legitimate and that the two programs are each "90% accurate" in the sense that $P(M_j | L) = P(M_j^c | L^c) = 9/10$. Also assume that given whether an email is spam, the two programs' outputs are conditionally independent.

- (a) Find the probability that the email is legitimate, given that the 1st program marks it as legitimate (simplify).
- (b) Find the probability that the email is legitimate, given that both programs mark it as legitimate (simplify).
- (c) Bob runs the 1st program and M_1 occurs. He updates his probabilities and then runs the 2nd program. Let $\tilde{P}(A) = P(A | M_1)$ be the updated probability function after running the 1st program. Explain briefly in words whether or not $\tilde{P}(L | M_2) = P(L | M_1 \cap M_2)$: is conditioning on $M_1 \cap M_2$ in one step equivalent to first conditioning on M_1 , then updating probabilities, and then conditioning on M_2 ?

27. Suppose that there are 5 blood types in the population, named type 1 through type 5, with probabilities $p_1, p_2, \dots, p_5$. A crime was committed by two individuals. A suspect, who has blood type 1, has prior probability p of being guilty. At the crime scene, blood evidence is collected, which shows that one of the criminals has type 1 and the other has type 2.

Find the posterior probability that the suspect is guilty, given the evidence. Does the evidence make it more likely or less likely that the suspect is guilty, or does this depend on the values of the parameters $p, p_1, \dots, p_5$? If it depends on these values, give a simple criterion for when the evidence makes it more likely that the suspect is guilty.

28. Fred has just tested positive for a certain disease.

- (a) Given this information, find the posterior odds that he has the disease, in terms of the prior odds, the sensitivity of the test, and the specificity of the test.
- (b) Not surprisingly, Fred is much more interested in P (have disease | test positive), known as the positive predictive value, than in the sensitivity P (test positive | have disease). A handy rule of thumb in biostatistics and epidemiology is as follows:

For a rare disease and a reasonably good test, specificity matters much more than sensitivity in determining the positive predictive value.

Explain intuitively why this rule of thumb works. For this part you can make up some specific numbers and interpret probabilities in a frequentist way as proportions in a large population, e.g., assume the disease afflicts 1% of a population of 10000 people and then consider various possibilities for the sensitivity and specificity.

29. A family has two children. Let C be a characteristic that a child can have, and assume that each child has characteristic C with probability p , independently of each other and of gender. For example, C could be the characteristic "born in winter" as in Example 2.2.7. Under the assumptions of Example 2.2.5, show that the probability that both children are girls given that at least one is a girl with characteristic C is $\frac{2-p}{4-p}$. Note that this is $1/3$ if $p = 1$ (agreeing with the first part of Example 2.2.5) and approaches $1/2$ from below as $p \rightarrow 0$ (agreeing with Example 2.2.7).

Independence and Conditional Independence

30. (S) A family has 3 children, creatively named A, B , and C .

- (a) Discuss intuitively (but clearly) whether the event " A is older than B " is independent

of the event " A is older than C ".

(b) Find the probability that A is older than B , given that A is older than C .

31. (S) Is it possible that an event is independent of itself? If so, when is this the case?

32. (S) Consider four nonstandard dice (the Efron dice), whose sides are labeled as follows (the 6 sides on each die are equally likely).

A: 4, 4, 4, 4, 0, 0

B: 3, 3, 3, 3, 3, 3

C: 6, 6, 2, 2, 2, 2

D: 5, 5, 5, 1, 1, 1

These four dice are each rolled once. Let A be the result for die A, B be the result for die B, etc.

(a) Find $P(A > B)$, $P(B > C)$, $P(C > D)$, and $P(D > A)$.

(b) Is the event $A > B$ independent of the event $B > C$? Is the event $B > C$ independent of the event $C > D$? Explain.

33. Alice, Bob, and 100 other people live in a small town. Let C be the set consisting of the 100 other people, let A be the set of people in C who are friends with Alice, and let B be the set of people in C who are friends with Bob. Suppose that for each person in C , Alice is friends with that person with probability $1/2$, and likewise for Bob, with all of these friendship statuses independent.

(a) Let $D \subseteq C$. Find $P(A = D)$.

(b) Find $P(A \subseteq B)$.

(c) Find $P(A \cup B = C)$.

34. Suppose that there are two types of drivers: good drivers and bad drivers. Let G be the event that a certain man is a good driver, A be the event that he gets into a car accident next year, and B be the event that he gets into a car accident the following year. Let $P(G) = g$ and $P(A | G) = P(B | G) = p_1$, $P(A | G^c) = P(B | G^c) = p_2$, with $p_1 < p_2$. Suppose that given the information of whether or not the man is a good driver, A and B are independent (for simplicity and to avoid being morbid, assume that the accidents being considered are minor and wouldn't make the man unable to drive).

(a) Explain intuitively whether or not A and B are independent.

(b) Find $P(G | A^c)$.

(c) Find $P(B | A^c)$.

35. (S) You are going to play 2 games of chess with an opponent whom you have never played against before (for the sake of this problem). Your opponent is equally likely to be a beginner, intermediate, or a master. Depending on which, your chances of winning an individual game are 90%, 50%, or 30%, respectively.

(a) What is your probability of winning the first game?

(b) Congratulations: you won the first game! Given this information, what is the probability that you will also win the second game (assume that, given the skill level of your opponent, the outcomes of the games are independent)?

(c) Explain the distinction between assuming that the outcomes of the games are independent and assuming that they are conditionally independent given the opponent's skill level. Which of these assumptions seems more reasonable, and why?

36. (a) Suppose that in the population of college applicants, being good at baseball is independent of having a good math score on a certain standardized test (with respect to some measure of "good"). A certain college has a simple admissions procedure: admit an applicant if and only if the applicant is good at baseball or has a good math score on the test.

Give an intuitive explanation of why it makes sense that among students that the college admits, having a good math score is negatively associated with being good at baseball, i.e.,

conditioning on having a good math score decreases the chance of being good at baseball.

(b) Show that if A and B are independent and $C = A \cup B$, then A and B are conditionally dependent given C (as long as $P(A \cap B) > 0$ and $P(A \cup B) < 1$), with

$$P(A | B, C) < P(A | C).$$

This phenomenon is known as Berkson's paradox, especially in the context of admissions to a school, hospital, etc.

37. Two different diseases cause a certain weird symptom; anyone who has either or both of these diseases will experience the symptom. Let D_1 be the event of having the first disease, D_2 be the event of having the second disease, and W be the event of having the weird symptom. Suppose that D_1 and D_2 are independent with $P(D_j) = p_j$, and that a person with neither of these diseases will have the weird symptom with probability w_0 . Let $q_j = 1 - p_j$, and assume that $0 < p_j < 1$.

(a) Find $P(W)$.

(b) Find $P(D_1 | W)$, $P(D_2 | W)$, and $P(D_1, D_2 | W)$.

(c) Determine algebraically whether or not D_1 and D_2 are conditionally independent given W .

(d) Suppose for this part only that $w_0 = 0$. Give a clear, convincing intuitive explanation in words of whether D_1 and D_2 are conditionally independent given W .

38. We want to design a spam filter for email. As described in Exercise 1, a major strategy is to find phrases that are much more likely to appear in a spam email than in a nonspam email. In that exercise, we only consider one such phrase: "free money". More realistically, suppose that we have created a list of 100 words or phrases that are much more likely to be used in spam than in non-spam.

Let W_j be the event that an email contains the j th word or phrase on the list. Let

$$p = P(\text{spam}), p_j = P(W_j | \text{spam}), r_j = P(W_j | \text{not spam}),$$

where "spam" is shorthand for the event that the email is spam.

Assume that $W_1, \dots, W_{100}$ are conditionally independent given that the email is spam, and conditionally independent given that it is not spam. A method for classifying emails (or other objects) based on this kind of assumption is called a naive Bayes classifier. (Here "naive" refers to the fact that the conditional independence is a strong assumption, not to Bayes being naive. The assumption may or may not be realistic, but naive Bayes classifiers sometimes work well in practice even if the assumption is not realistic.)

Under this assumption we know, for example, that

$$P(W_1, W_2, W_3^c, W_4^c, \dots, W_{100}^c | \text{spam}) = p_1 p_2 (1 - p_3) (1 - p_4) \dots (1 - p_{100}).$$

Without the naive Bayes assumption, there would be vastly more statistical and computational difficulties since we would need to consider $2^{100} \approx 1.3 \times 10^{30}$ events of the form $A_1 \cap A_2 \dots \cap A_{100}$ with each A_j equal to either W_j or W_j^c . A new email has just arrived, and it includes the 23rd, 64th, and 65th words or phrases on the list (but not the other 97). So we want to compute

$$P(\text{spam} | W_1^c, \dots, W_{22}^c, W_{23}, W_{24}^c, \dots, W_{63}^c, W_{64}, W_{65}, W_{66}^c, \dots, W_{100}^c).$$

Note that we need to condition on all the evidence, not just the fact that $W_{23} \cap W_{64} \cap W_{65}$ occurred. Find the conditional probability that the new email is spam (in terms of p and the p_j and r_j).

Monty Hall

39. (S) (a) Consider the following 7 -door version of the Monty Hall problem. There are 7 doors, behind one of which there is a car (which you want), and behind the rest of which there are goats (which you don't want). Initially, all possibilities are equally likely for where the car is. You choose a door. Monty Hall then opens 3 goat doors, and offers you the option of switching to any of the remaining 3 doors.
Assume that Monty Hall knows which door has the car, will always open 3 goat doors and offer the option of switching, and that Monty chooses with equal probabilities from all his choices of which goat doors to open. Should you switch? What is your probability of success if you switch to one of the remaining 3 doors?
- (b) Generalize the above to a Monty Hall problem where there are $n \geq 3$ doors, of which Monty opens m goat doors, with $1 \leq m \leq n - 2$.
40. (S) Consider the Monty Hall problem, except that Monty enjoys opening door 2 more than he enjoys opening door 3, and if he has a choice between opening these two doors, he opens door 2 with probability p , where $\frac{1}{2} \leq p \leq 1$.
To recap: there are three doors, behind one of which there is a car (which you want), and behind the other two of which there are goats (which you don't want). Initially, all possibilities are equally likely for where the car is. You choose a door, which for concreteness we assume is door 1. Monty Hall then opens a door to reveal a goat, and offers you the option of switching. Assume that Monty Hall knows which door has the car, will always open a goat door and offer the option of switching, and as above assume that if Monty Hall has a choice between opening door 2 and door 3, he chooses door 2 with probability p (with $\frac{1}{2} \leq p \leq 1$).
- (a) Find the unconditional probability that the strategy of always switching succeeds (unconditional in the sense that we do not condition on which of doors 2 or 3 Monty opens).
- (b) Find the probability that the strategy of always switching succeeds, given that Monty opens door 2 .
- (c) Find the probability that the strategy of always switching succeeds, given that Monty opens door 3 .
41. The ratings of Monty Hall's show have dropped slightly, and a panicking executive producer complains to Monty that the part of the show where he opens a door lacks suspense: Monty always opens a door with a goat. Monty replies that the reason is so that the game is never spoiled by him revealing the car, but he agrees to update the game as follows.

Before each show, Monty secretly flips a coin with probability p of Heads. If the coin lands Heads, Monty resolves to open a door with a goat (with equal probabilities if there is a choice). Otherwise, Monty resolves to open a random door, with equal probabilities. Of course, Monty will not open the door that the contestant initially chooses. The contestant knows p but does not know the outcome of the coin flip. When the show starts, the contestant chooses a door. Monty (who knows where the car is) then opens a door. If the car is revealed, the game is over; if a goat is revealed, the contestant is offered the option of switching. Now suppose it turns out that the contestant chooses door 1 and then Monty opens door 2, revealing a goat. What is the contestant's probability of success if they switch to door 3 ?

42. Consider the following variation of the Monty Hall problem, where in some situations Monty may not open a door and give the contestant the choice of whether to switch doors. Specifically, there are 3 doors, with 2 containing goats and 1 containing a car. The car is equally likely to be anywhere, and Monty knows where the car is. Let $0 \leq p \leq 1$.

The contestant chooses a door. If this initial choice has the car, Monty will open another door, revealing a goat (choosing with equal probabilities among his two choices of door), and then offer the contestant the choice of whether to switch to the other unopened door.

If the contestant's initial choice has a goat, then with probability p Monty will open another door, revealing a goat, and then offer the contestant the choice of whether to switch to the other unopened door; but with probability $1 - p$, Monty will not open a door, and the contestant must stick with their initial choice.

The contestant decides in advance to use the following strategy: initially choose door 1. Then, if Monty opens a door and offers the choice of whether to switch, do switch.

(a) Find the unconditional probability that the contestant will get the car. Also, check what your answer reduces to in the extreme cases $p = 0$ and $p = 1$, and briefly explain why your answer makes sense in these two cases.

(b) Monty now opens door 2, revealing a goat. So the contestant switches to door 3. Given this information, find the conditional probability that the contestant will get the car.

43. You are the contestant on the Monty Hall show. Monty is trying out a new version of his game, with rules as follows. You get to choose one of three doors. One door has a car behind it, another has a computer, and the other door has a goat (with all permutations equally likely). Monty, who knows which prize is behind each door, will open a door (but not the one you chose) and then let you choose whether to switch from your current choice to the other unopened door.

Assume that you prefer the car to the computer, the computer to the goat, and (by transitivity) the car to the goat.

(a) Suppose for this part only that Monty always opens the door that reveals your less preferred prize out of the two alternatives, e.g., if he is faced with the choice between revealing the goat or the computer, he will reveal the goat. Monty opens a door, revealing a goat (this is again for this part only). Given this information, should you switch? If you do switch, what is your probability of success in getting the car?

(b) Now suppose that Monty reveals your less preferred prize with probability p , and your more preferred prize with probability $q = 1 - p$. Monty opens a door, revealing a computer. Given this information, should you switch (your answer can depend on p)? If you do switch, what is your probability of success in getting the car (in terms of p)?

44. Monty Hall has introduced a new twist in his game, by generalizing the assumption that the initial probabilities for where the car is are $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Specifically, there are three doors, behind one of which there is a car (which the contestant wants), and behind the other two of which there are goats (which the contestant doesn't want). Initially, door i has probability p_i of having the car, where p_1, p_2, p_3 are known constants such that $0 < p_1 \leq p_2 \leq p_3 < 1$ and $p_1 + p_2 + p_3 = 1$. The contestant chooses a door. Then Monty opens a door (other than the one the contestant chose) and offers the contestant the option of switching to the other unopened door.

(a) Assume for this part that Monty knows in advance which door has the car. He always opens a door to reveal a goat, and if he has a choice of which door to open he chooses with equal probabilities. Suppose for this part that the contestant initially chooses door 3, and then Monty opens door 2, revealing a goat. Given the above information, find the conditional probability that door 3 has the car. Should the contestant switch doors? (If whether to switch depends on the p_i 's, give a fully simplified criterion in terms of the p_i 's.)

(b) Now assume instead that Monty does not know in advance where the car is. He randomly chooses which door to open (other than the one the contestant chose), with equal probabilities. (The game is spoiled if he reveals the car.) Suppose again that the contestant initially chooses door 3, and then Monty opens door 2, revealing a goat. Given the above information, find the conditional probability that door 3 has the car. Should the contestant switch doors? (If whether to switch depends on the p_i 's, give a fully simplified criterion

in terms of the p_i 's.)

(c) Repeat (a), except with the contestant initially choosing door 1 rather than door 3 .

(d) Repeat (b), except with the contestant initially choosing door 1 rather than door 3.

45. Monty Hall is trying out a new version of his game. In this version, instead of there always being 1 car and 2 goats, the prizes behind the doors are generated independently, with each door having probability p of having a car and $q = 1 - p$ of having a goat. In detail: There are three doors, behind each of which there is one prize: either a car or a goat. For each door, there is probability p that there is a car behind it and $q = 1 - p$ that there is a goat, independent of the other doors.

The contestant chooses a door. Monty, who knows the contents of each door, then opens one of the two remaining doors. In choosing which door to open, Monty will always reveal a goat if possible. If both of the remaining doors have the same kind of prize, Monty chooses randomly (with equal probabilities). After opening a door, Monty offers the contestant the option of switching to the other unopened door.

The contestant decides in advance to use the following strategy: first choose door 1 . Then, after Monty opens a door, switch to the other unopened door.

(a) Find the unconditional probability that the contestant will get a car.

(b) Monty now opens door 2, revealing a goat. Given this information, find the conditional probability that the contestant will get a car.

46. Monty Hall is trying out a new version of his game, with rules as follows. The contestant gets to choose one of four doors. One door has a car behind it, another has an apple, another has a book, and another has a goat. All 24 permutations for which door has which prize are equally likely. In order from least preferred to most preferred, the contestant's preferences are: goat, apple, book, car.

Monty, who knows which prize is behind each door, will open a door (other than the contestant's initial choice) and then let the contestant choose whether to switch to another unopened door. Monty will reveal the least preferred prize (among the 3 doors other than the contestant's initial choice) with probability p , the intermediately preferred prize with probability $1 - p$, and the most preferred prize never.

The contestant decides in advance to use the following strategy: Initially choose door 1. After Monty opens a door, switch to one of the other two unopened doors, randomly choosing between them (with probability $1/2$ each).

(a) Find the unconditional probability that the contestant will get the car.

Hint: Condition on where the car is.

(b) Find the unconditional probability that Monty will reveal the apple.

Hint: Condition on what is behind door 1.

(c) Monty now opens a door, revealing the apple. Given this information, find the conditional probability that the contestant will get the car.

47. You are the contestant on Monty Hall's game show. Hoping to double the excitement of the game, Monty will offer you two opportunities to switch to another door. Specifically, the new rules are as follows. There are four doors. Behind one door there is a car (which you want); behind the other three doors there are goats (which you don't want). Initially, all possibilities are equally likely for where the car is. Monty knows where the car is, and when he has a choice of which door to open, he chooses with equal probabilities.

You choose a door, which for concreteness we assume is door 1 . Monty opens a door (other than door 1), revealing a goat, and then offers you the option to switch to another door. Monty then opens another door (other than your currently selected door), revealing another goat. So now there are two open doors (with goats) and two unopened doors. Again Monty offers you the option to switch.

You decide in advance to use one of the following four strategies: stay-stay, stay-switch, switch-

stay, switch-switch, where, for example, "stay-switch" means that the first time Monty offers you the choice of switching, you stay with your current selection, but then the second time Monty offers you the choice, you do switch doors. In each part below the goal is to find or compare unconditional probabilities, i.e., from a vantage point of before the game has started.

- (a) Find the probability of winning the car if you follow the stay-stay strategy.
- (b) Find the probability of winning the car if you follow the stay-switch strategy.
- (c) Find the probability of winning the car if you follow the switch-stay strategy.
- (d) Find the probability of winning the car if you follow the switch-switch strategy.
- (e) Which of these four strategies is the best?

First-step Analysis and Gambler's Ruin

- 48. (S) A fair die is rolled repeatedly, and a running total is kept (which is, at each time, the total of all the rolls up until that time). Let p_n be the probability that the running total is ever exactly n (assume the die will always be rolled enough times so that the running total will eventually exceed n , but it may or may not ever equal n).
 - (a) Write down a recursive equation for p_n (relating p_n to earlier terms p_k in a simple way). Your equation should be true for all positive integers n , so give a definition of p_0 and p_k for $k < 0$ so that the recursive equation is true for small values of n .
 - (b) Find p_7 .
 - (c) Give an intuitive explanation for the fact that $p_n \rightarrow 1/3.5 = 2/7$ as $n \rightarrow \infty$.
- 49. A sequence of $n \geq 1$ independent trials is performed, where each trial ends in "success" or "failure" (but not both). Let p_i be the probability of success in the i th trial, $q_i = 1 - p_i$, and $b_i = q_i - 1/2$, for $i = 1, 2, \dots, n$. Let A_n be the event that the number of successful trials is even.
 - (a) Show that for $n = 2$, $P(A_2) = 1/2 + 2b_1b_2$.
 - (b) Show by induction that

$$P(A_n) = 1/2 + 2^{n-1}b_1b_2\dots b_n$$

(This result is very useful in cryptography. Also, note that it implies that if n coins are flipped, then the probability of an even number of Heads is $1/2$ if and only if at least one of the coins is fair.) Hint: Group some trials into a supertrial.

- (c) Check directly that the result of (b) is true in the following simple cases: $p_i = 1/2$ for some i ; $p_i = 0$ for all i ; $p_i = 1$ for all i .

- 50. (S) Calvin and Hobbes play a match consisting of a series of games, where Calvin has probability p of winning each game (independently). They play with a "win by two" rule: the first player to win two games more than his opponent wins the match. Find the probability that Calvin wins the match (in terms of p), in two different ways:

- (a) by conditioning, using the law of total probability.
- (b) by interpreting the problem as a gambler's ruin problem.

- 51. (S) A gambler repeatedly plays a game where in each round, he wins a dollar with probability $1/3$ and loses a dollar with probability $2/3$. His strategy is "quit when he is ahead by \$2". Suppose that he starts with a million dollars. Show that the probability that he'll ever be ahead by \$2 is less than $1/4$.

- 52. As in the gambler's ruin problem, two gamblers, A and B, make a series of bets, until one of the gamblers goes bankrupt. Let A start out with i dollars and B start out with $N - i$ dollars, and let p be the probability of A winning a bet, with $0 < p < \frac{1}{2}$. Each bet is for $\frac{1}{k}$ dollars, with k a positive integer, e.g., $k = 1$ is the original gambler's ruin problem and $k = 20$ means they're betting nickels. Find the probability that A wins the game, and determine what

happens to this as $k \rightarrow \infty$.

53. There are 100 equally spaced points around a circle. At 99 of the points, there are sheep, and at 1 point, there is a wolf. At each time step, the wolf randomly moves either clockwise or counterclockwise by 1 point. If there is a sheep at that point, he eats it. The sheep don't move. What is the probability that the sheep who is initially opposite the wolf is the last one remaining?

54. An immortal drunk man wanders around randomly on the integers. He starts at the origin, and at each step he moves 1 unit to the right or 1 unit to the left, with probabilities p and $q = 1 - p$ respectively, independently of all his previous steps. Let S_n be his position after n steps.

(a) Let p_k be the probability that the drunk ever reaches the value k , for all $k \geq 0$. Write down a difference equation for p_k (you do not need to solve it for this part).

(b) Find p_k , fully simplified; be sure to consider all 3 cases: $p < 1/2$, $p = 1/2$, and $p > 1/2$. Feel free to assume that if $A_1, A_2, \dots$ are events with $A_j \subseteq A_{j+1}$ for all j , then $P(A_n) \rightarrow P(\cup_{j=1}^{\infty} A_j)$ as $n \rightarrow \infty$ (because it is true; this is known as continuity of probability).

Simpson's Paradox

55. (S) (a) Is it possible to have events A, B, C such that $P(A | C) < P(B | C)$ and $P(A | C^c) < P(B | C^c)$, yet $P(A) > P(B)$? That is, A is less likely than B given that C is true, and also less likely than B given that C is false, yet A is more likely than B if we're given no information about C . Show this is impossible (with a short proof) or find a counterexample (with a story interpreting A, B, C).

(b) If the scenario in (a) is possible, is it a special case of Simpson's paradox, equivalent to Simpson's paradox, or neither? If it is impossible, explain intuitively why it is impossible even though Simpson's paradox is possible.

56. (S) Consider the following conversation from an episode of The Simpsons:

Lisa: Dad, I think he's an ivory dealer! His boots are ivory, his hat is ivory, and I'm pretty sure that check is ivory.

Homer: Lisa, a guy who has lots of ivory is less likely to hurt Stampy than a guy whose ivory supplies are low.

Here Homer and Lisa are debating the question of whether or not the man (named Blackheart) is likely to hurt Stampy the Elephant if they sell Stampy to him. They clearly disagree about how to use their observations about Blackheart to learn about the probability (conditional on the evidence) that Blackheart will hurt Stampy.

(a) Define clear notation for the various events of interest here.

(b) Express Lisa's and Homer's arguments (Lisa's is partly implicit) as conditional probability statements in terms of your notation from (a).

(c) Assume it is true that someone who has a lot of a commodity will have less desire to acquire more of the commodity. Explain what is wrong with Homer's reasoning that the evidence about Blackheart makes it less likely that he will harm Stampy.

57. (a) There are two crimson jars (labeled C_1 and C_2) and two mauve jars (labeled M_1 and M_2). Each jar contains a mixture of green gummi bears and red gummi bears. Show by example that it is possible that C_1 has a much higher percentage of green gummi bears than M_1 , and C_2 has a much higher percentage of green gummi bears than M_2 , yet if the contents of C_1 and C_2 are merged into a new jar and likewise for M_1 and M_2 , then the combination of C_1 and C_2 has a lower percentage of green gummi bears than the combination of M_1 and M_2 .

(b) Explain how (a) relates to Simpson's paradox, both intuitively and by explicitly defining events A, B, C as in the statement of Simpson's paradox.

58. As explained in this chapter, Simpson's paradox says that it is possible to have events A, B, C such that $P(A | B, C) < P(A | B^c, C)$ and $P(A | B, C^c) < P(A | B^c, C^c)$, yet $P(A | B) > P(A | B^c)$.

(a) Can Simpson's paradox occur if A and B are independent? If so, give a concrete example (with both numbers and an interpretation); if not, prove that it is impossible.

(b) Can Simpson's paradox occur if A and C are independent? If so, give a concrete example (with both numbers and an interpretation); if not, prove that it is impossible.

(c) Can Simpson's paradox occur if B and C are independent? If so, give a concrete example (with both numbers and an interpretation); if not, prove that it is impossible.

59. (S) The book *Red State, Blue State, Rich State, Poor State* by Andrew Gelman [12] discusses the following election phenomenon: within any U.S. state, a wealthy voter is more likely to vote for a Republican than a poor voter, yet the wealthier states tend to favor Democratic candidates!

(a) Assume for simplicity that there are only 2 states (called Red and Blue), each of which has 100 people, and that each person is either rich or poor, and either a Democrat or a Republican. Make up numbers consistent with the above, showing how this phenomenon is possible, by giving a 2×2 table for each state (listing how many people in each state are rich Democrats, etc.). So within each state, a rich voter is more likely to vote for a Republican than a poor voter, but the percentage of Democrats is higher in the state with the higher percentage of rich people than in the state with the lower percentage of rich people.

(b) In the setup of (a) (not necessarily with the numbers you made up there), let D be the event that a randomly chosen person is a Democrat (with all 200 people equally likely), and B be the event that the person lives in the Blue State. Suppose that 10 people move from the Blue State to the Red State. Write P_{old} and P_{new} for probabilities before and after they move. Assume that people do not change parties, so we have $P_{\text{new}}(D) = P_{\text{old}}(D)$. Is it possible that both $P_{\text{new}}(D | B) > P_{\text{old}}(D | B)$ and $P_{\text{new}}(D | B^c) > P_{\text{old}}(D | B^c)$ are true? If so, explain how it is possible and why it does not contradict the law of total probability $P(D) = P(D | B)P(B) + P(D | B^c)P(B^c)$; if not, show that it is impossible.

Mixed Practice

60. A patient is being given a blood test for the disease conditionitis. Let p be the prior probability that the patient has conditionitis. The blood sample is sent to one of two labs for analysis, lab A or lab B. The choice of which lab to use is made randomly, independent of the patient's disease status, with probability $1/2$ for each lab.

For lab A, the probability of someone testing positive given that they do have the disease is a_1 , and the probability of someone testing negative given that they do not have the disease is a_2 . The corresponding probabilities for lab B are b_1 and b_2 .

(a) Find the probability that the patient has the disease, given that they tested positive.

(b) Find the probability that the patient's blood sample was analyzed by lab A, given that the patient tested positive.

61. Fred decides to take a series of n tests, to diagnose whether he has a certain disease (any individual test is not perfectly reliable, so he hopes to reduce his uncertainty by taking multiple tests). Let D be the event that he has the disease, $p = P(D)$ be the prior probability that he has the disease, and $q = 1 - p$. Let T_j be the event that he tests positive on the j th test.

(a) Assume for this part that the test results are conditionally independent given Fred's disease status. Let $a = P(T_j | D)$ and $b = P(T_j | D^c)$, where a and b don't depend on j . Find the posterior probability that Fred has the disease, given that he tests positive on all n of the n tests.

(b) Suppose that Fred tests positive on all n tests. However, some people have a certain gene

that makes them always test positive. Let G be the event that Fred has the gene. Assume that $P(G) = 1/2$ and that D and G are independent. If Fred does not have the gene, then the test results are conditionally independent given his disease status. Let $a_0 = P(T_j | D, G^c)$ and $b_0 = P(T_j | D^c, G^c)$, where a_0 and b_0 don't depend on j . Find the posterior probability that Fred has the disease, given that he tests positive on all n of the tests.

62. A certain hereditary disease can be passed from a mother to her children. Given that the mother has the disease, her children independently will have it with probability $1/2$. Given that she doesn't have the disease, her children won't have it either. A certain mother, who has probability $1/3$ of having the disease, has two children.

(a) Find the probability that neither child has the disease.

(b) Is whether the elder child has the disease independent of whether the younger child has the disease? Explain.

(c) The elder child is found not to have the disease. A week later, the younger child is also found not to have the disease. Given this information, find the probability that the mother has the disease.

63. Three fair coins are tossed at the same time. Explain what is wrong with the following argument: "there is a 50% chance that the three coins all landed the same way, since obviously it is possible to find two coins that match, and then the other coin has a 50% chance of matching those two".

64. An urn contains red, green, and blue balls. Let r, g, b be the proportions of red, green, blue balls, respectively ($r + g + b = 1$).

(a) Balls are drawn randomly with replacement. Find the probability that the first time a green ball is drawn is before the first time a blue ball is drawn.

Hint: Explain how this relates to finding the probability that a draw is green, given that it is either green or blue.

(b) Balls are drawn randomly without replacement. Find the probability that the first time a green ball is drawn is before the first time a blue ball is drawn. Is the answer the same or different than the answer in (a)?

Hint: Imagine the balls all lined up, in the order in which they will be drawn. Note that where the red balls are standing in this line is irrelevant.

(c) Generalize the result from (a) to the following setting. Independent trials are performed, and the outcome of each trial is classified as being exactly one of type 1, type 2, ..., or type n , with probabilities $p_1, p_2, \dots, p_n$, respectively. Find the probability that the first trial to result in type i comes before the first trial to result in type j , for $i \neq j$.

65. Marilyn vos Savant was asked the following question for her column in Parade:

You're at a party with 199 other guests when robbers break in and announce that they are going to rob one of you. They put 199 blank pieces of paper in a hat, plus one marked "you lose." Each guest must draw, and the person who draws "you lose" will get robbed. The robbers offer you the option of drawing first, last, or at any time in between. When would you take your turn?

The draws are made without replacement, and for (a) are uniformly random.

(a) Determine whether it is optimal to draw first, last, or somewhere in between (or whether it does not matter), to maximize the probability of not being robbed. Give a clear, concise, and compelling explanation.

(b) More generally, suppose that there is one "you lose" piece of paper, with "weight" v , and there are n blank pieces of paper, each with "weight" w . At each stage, draws are made with probability proportional to weight, i.e., the probability of drawing a particular piece of paper is its weight divided by the sum of the weights of all the remaining pieces of paper. Determine whether it is better to draw first or second (or whether it does not matter); here $v > 0, w > 0$, and $n \geq 1$ are known constants.

66. A fair die is rolled repeatedly, until the running total is at least 100 (at which point the rolling stops). Find the most likely value of the final running total (i.e., the value of the running total at the first time when it is at least 100).

Hint: Consider the possibilities for what the running total is just before the last roll.

67. Homer has a box of donuts, which currently contains exactly c chocolate, g glazed, and j jelly donuts. Homer eats donuts one after another, each time choosing uniformly at random from the remaining donuts.

(a) Find the probability that the last donut remaining in the box is a chocolate donut.

(b) Find the probability of the following event: glazed is the first type of donut that Homer runs out of, and then jelly is the second type of donut that he runs out of.

Hint: Consider the last donut remaining, and the last donut that is either glazed or jelly.

68. Let D be the event that a person develops a certain disease, and C be the event that the person was exposed to a certain substance (e.g., D may correspond to lung cancer and C may correspond to smoking cigarettes). We are interested in whether exposure to the substance is related to developing the disease (and if so, how they are related). The odds ratio is a very widely used measure in epidemiology of the association between disease and exposure, defined as

$$\text{OR} = \frac{\text{odds}(D | C)}{\text{odds}(D | C^c)}$$

where conditional odds are defined analogously to unconditional odds: $\text{odds}(A | B) = \frac{P(A|B)}{P(A^c|B)}$. The relative risk of the disease for someone exposed to the substance, another widely used measure, is

$$\text{RR} = \frac{P(D | C)}{P(D | C^c)}$$

The relative risk is especially easy to interpret, e.g., $\text{RR} = 2$ says that someone exposed to the substance is twice as likely to develop the disease as someone who isn't exposed (though this does not necessarily mean that the substance causes the increased chance of getting the disease, nor is there necessarily a causal interpretation for the odds ratio).

(a) Show that if the disease is rare, both for exposed people and for unexposed people, then the relative risk is approximately equal to the odds ratio.

(b) Let p_{ij} for $i = 0, 1$ and $j = 0, 1$ be the probabilities in the following 2×2 table.

	D	D^c
C	p_{11}	p_{10}
C^c	p_{01}	p_{00}

For example, $p_{10} = P(C, D^c)$. Show that the odds ratio can be expressed as a crossproduct ratio, in the sense that

$$\text{OR} = \frac{p_{11}p_{00}}{p_{10}p_{01}}$$

(c) Show that the odds ratio has the neat symmetry property that the roles of C and D can be swapped without changing the value:

$$\text{OR} = \frac{\text{odds}(C | D)}{\text{odds}(C | D^c)}$$

This property is one of the main reasons why the odds ratio is so widely used, since it turns out that it allows the odds ratio to be estimated in a wide variety of problems where relative

risk would be hard to estimate well.

69. A researcher wants to estimate the percentage of people in some population who have used illegal drugs, by conducting a survey. Concerned that a lot of people would lie when asked a sensitive question like "Have you ever used illegal drugs?", the researcher uses a method known as randomized response. A hat is filled with slips of paper, each of which says either "I have used illegal drugs" or "I have not used illegal drugs". Let p be the proportion of slips of paper that say "I have used illegal drugs" (p is chosen by the researcher in advance).

Each participant chooses a random slip of paper from the hat and answers (truthfully) "yes" or "no" to whether the statement on that slip is true. The slip is then returned to the hat. The researcher does not know which type of slip the participant had. Let y be the probability that a participant will say "yes", and d be the probability that a participant has used illegal drugs.

(a) Find y , in terms of d and p .

(b) What would be the worst possible choice of p that the researcher could make in designing the survey? Explain.

(c) Now consider the following alternative system. Suppose that proportion p of the slips of paper say "I have used illegal drugs", but that now the remaining $1 - p$ say "I was born in winter" rather than "I have not used illegal drugs". Assume that $1/4$ of people are born in winter, and that a person's season of birth is independent of whether they have used illegal drugs. Find d , in terms of y and p .

70. At the beginning of the play *Rosencrantz and Guildenstern Are Dead* by Tom Stoppard [25], Guildenstern is spinning coins and Rosencrantz is betting on the outcome for each. The coins have been landing Heads over and over again, prompting the following remark:

Guildenstern: A weaker man might be moved to re-examine his faith, if in nothing else at least in the law of probability.

The coin spins have resulted in Heads 92 times in a row.

(a) Fred and his friend are watching the play. Upon seeing the events described above, they have the following conversation:

Fred: That outcome would be incredibly unlikely with fair coins. They must be using trick coins (maybe with double-headed coins), or the experiment must have been rigged somehow (maybe with magnets).

Fred's friend: It's true that the string HH...H of length 92 is very unlikely; the chance is $1/2^{92} \approx 2 \times 10^{-28}$ with fair coins. But any other specific string of H's and T's with length 92 has exactly the same probability! The reason the outcome seems extremely unlikely is that the number of possible outcomes grows exponentially as the number of spins grows, so any outcome would seem extremely unlikely. You could just as well have made the same argument even without looking at the results of their experiment, which means you really don't have evidence against the coins being fair.

Discuss these comments, to help Fred and his friend resolve their debate.

(b) Suppose there are only two possibilities: either the coins are all fair (and spun fairly), or double-headed coins are being used (in which case the probability of Heads is 1). Let p be the prior probability that the coins are fair. Find the posterior probability that the coins are fair, given that they landed Heads in 92 out of 92 trials.

(c) Continuing from (b), for which values of p is the posterior probability that the coins are fair greater than 0.5? For which values of p is it less than 0.05?

71. There are n types of toys, which you are collecting one by one. Each time you buy a toy, it is randomly determined which type it has, with equal probabilities. Let p_{ij} be the probability that just after you have bought your i th toy, you have exactly j toy types in your collection, for $i \geq 1$ and $0 \leq j \leq n$. (This problem is in the setting of the coupon collector problem, a famous problem which we study in Example 4.3.12.)

(a) Find a recursive equation expressing p_{ij} in terms of $p_{i-1,j}$ and $p_{i-1,j-1}$, for $i \geq 2$ and

$1 \leq j \leq n$.

(b) Describe how the recursion from (a) can be used to calculate p_{ij} .

72. A/B testing is a form of randomized experiment that is used by many companies to learn about how customers will react to different treatments. For example, a company may want to see how users will respond to a new feature on their website (compared with how users respond to the current version of the website) or compare two different advertisements.

As the name suggests, two different treatments, Treatment A and Treatment B, are being studied. Users arrive one by one, and upon arrival are randomly assigned to one of the two treatments. The trial for each user is classified as "success" (e.g., the user made a purchase) or "failure". The probability that the n th user receives Treatment A is allowed to depend on the outcomes for the previous users. This set-up is known as a two-armed bandit.

Many algorithms for how to randomize the treatment assignments have been studied. Here is an especially simple (but fickle) algorithm, called a stay-with-a-winner procedure:

- (i) Randomly assign the first user to Treatment A or Treatment B, with equal probabilities.
- (ii) If the trial for the n th user is a success, stay with the same treatment for the $(n + 1)$ st user; otherwise, switch to the other treatment for the $(n + 1)$ st user.

Let a be the probability of success for Treatment A, and b be the probability of success for Treatment B. Assume that $a \neq b$, but that a and b are unknown (which is why the test is needed). Let p_n be the probability of success on the n th trial and a_n be the probability that Treatment A is assigned on the n th trial (using the above algorithm).

(a) Show that

$$\begin{aligned}p_n &= (a - b)a_n + b, \\a_{n+1} &= (a + b - 1)a_n + 1 - b.\end{aligned}$$

(b) Use the results from (a) to show that p_{n+1} satisfies the following recursive equation:

$$p_{n+1} = (a + b - 1)p_n + a + b - 2ab$$

(c) Use the result from (b) to find the long-run probability of success for this algorithm, $\lim_{n \rightarrow \infty} p_n$, assuming that this limit exists.

73. In humans (and many other organisms), genes come in pairs. A certain gene comes in two types (alleles): type a and type A . The genotype of a person for that gene is the types of the two genes in the pair: AA , Aa , or aa (aA is equivalent to Aa). Assume that the Hardy-Weinberg law applies here, which means that the frequencies of AA , Aa , aa in the population are p^2 , $2p(1 - p)$, $(1 - p)^2$ respectively, for some p with $0 < p < 1$.

When a woman and a man have a child, the child's gene pair has one gene contributed by each parent. Suppose that the mother is equally likely to contribute either of the two genes in her gene pair, and likewise for the father, independently. Also suppose that the genotypes of the parents are independent of each other (with probabilities given by the Hardy-Weinberg law).

(a) Find the probabilities of each possible genotype (AA , Aa , aa) for a child of two random parents. Explain what this says about stability of the Hardy-Weinberg law from one generation to the next.

Hint: Condition on the genotypes of the parents.

(b) A person of type AA or aa is called homozygous (for the gene under consideration), and a person of type Aa is called heterozygous (for that gene). Find the probability that a child is homozygous, given that both parents are homozygous. Also, find the probability that a child is heterozygous, given that both parents are heterozygous.

(c) Suppose that having genotype aa results in a distinctive physical characteristic, so it is easy to tell by looking at someone whether or not they have that genotype. A mother and father, neither of whom are of type aa , have a child. The child is also not of type aa . Given this information, find the probability that the child is heterozygous.

Hint: Use the definition of conditional probability. Then expand both the numerator and the denominator using LOTP, conditioning on the genotypes of the parents.

74. A standard deck of cards will be shuffled and then the cards will be turned over one at a time until the first ace is revealed. Let B be the event that the next card in the deck will also be an ace.

(a) Intuitively, how do you think $P(B)$ compares in size with $1/13$ (the overall proportion of aces in a deck of cards)? Explain your intuition. (Give an intuitive discussion rather than a mathematical calculation; the goal here is to describe your intuition explicitly.)

(b) Let C_j be the event that the first ace is at position j in the deck. Find $P(B | C_j)$ in terms of j , fully simplified.

(c) Using the law of total probability, find an expression for $P(B)$ as a sum. (The sum can be left unsimplified, but it should be something that could easily be computed in software such as R that can calculate sums.)

(d) Find a fully simplified expression for $P(B)$ using a symmetry argument.

Hint: If you were deciding whether to bet on the next card after the first ace being an ace or to bet on the last card in the deck being an ace, would you have a preference?

4.3 Random Variables and Their Distributions

In this chapter, we introduce random variables, an incredibly useful concept that simplifies notation and expands our ability to quantify uncertainty and summarize the results of experiments. Random variables are essential throughout the rest of this book, and throughout statistics, so it is crucial to think through what they mean, both intuitively and mathematically.

4.3.1 3.1 Random Variables

To see why our current notation can quickly become unwieldy, consider again the gambler's ruin problem from Chapter 2. In this problem, we may be very interested in how much wealth each gambler has at any particular time. So we could make up notation like letting A_{jk} be the event that gambler A has exactly j dollars after k rounds, and similarly defining an event B_{jk} for gambler B, for all j and k .

This is already too complicated. Furthermore, we may also be interested in other quantities, such as the difference in their wealths (gambler A's minus gambler B's) after k rounds, or the duration of the game (the number of rounds until one player is bankrupt). Expressing the event "the duration of the game is r rounds" in terms of the A_{jk} and B_{jk} would involve a long, awkward string of unions and intersections. And then what if we want to express gambler A's wealth as the equivalent amount in euros rather than dollars? We can multiply a number in dollars by a currency exchange rate, but we can't multiply an event by an exchange rate.

Instead of having convoluted notation that obscures how the quantities of interest are related, wouldn't it be nice if we could say something like the following?

Let X_k be the wealth of gambler A after k rounds. Then $Y_k = N - X_k$ is the wealth of gambler B after k rounds (where N is the fixed total wealth); $X_k - Y_k = 2X_k - N$ is the difference in wealths after k rounds; $c_k X_k$ is the wealth of gambler A in euros after k rounds, where c_k is the euros per dollar exchange rate after k rounds; and the duration is $R = \min \{n : X_n = 0 \text{ or } Y_n = 0\}$.

The notion of a random variable will allow us to do exactly this! It needs to be introduced carefully though, to make it both conceptually and technically correct. Sometimes a definition of "random variable" is given that is a barely paraphrased version of "a random variable is a variable that takes on random values", but such a feeble attempt at a definition fails to say where the randomness come from. Nor does it help us to

derive properties of random variables: we're familiar with working with algebraic equations like $x^2 + y^2 = 1$, but what are the valid mathematical operations if x and y are random variables? To make the notion of random variable precise, we define it as a function mapping the sample space to the real line. (See the math appendix for review of some concepts about functions.)




FIGURE 3.1 A random variable maps the sample space into the real line. The r.v. X depicted here is defined on a sample space with 6 elements, and has possible values 0, 1, and 4. The randomness comes from choosing a random pebble according to the probability function P for the sample space.

Definition 3.1.1 (Random variable). Given an experiment with sample space S , a random variable (r.v.) is a function from the sample space S to the real numbers $\mathbb{R}$. It is common, but not required, to denote random variables by capital letters.

Thus, a random variable X assigns a numerical value $X(s)$ to each possible outcome s of the experiment. The randomness comes from the fact that we have a random experiment (with probabilities described by the probability function P); the mapping itself is deterministic, as illustrated in Figure 3.1. The same r.v. is shown in a simpler way in the left panel of Figure 3.2, in which we inscribe the values inside the pebbles.

This definition is abstract but fundamental; one of the most important skills to develop when studying probability and statistics is the ability to go back and forth between abstract ideas and concrete examples. Relatedly, it is important to work on recognizing the essential pattern or structure of a problem and how it connects to problems you have studied previously. We will often discuss stories that involve tossing coins or drawing balls from urns because they are simple, convenient scenarios to work with, but many other problems are isomorphic: they have the same essential structure, but in a different guise.

To start, let's consider a coin-tossing example. The structure of the problem is that we have a sequence of trials where there are two possible outcomes for each trial. Here we think of the possible outcomes as H (Heads) and T (Tails), but we could just as well think of them as "success" and "failure" or as 1 and 0, for example.

Example 3.1.2 (Coin tosses). Consider an experiment where we toss a fair coin twice. The sample space consists of four possible outcomes: $S = \{HH, HT, TH, TT\}$. Here are some random variables on this space (for practice, you can think up some of your own). Each r.v. is a numerical summary of some aspect of the experiment.

- Let X be the number of Heads. This is a random variable with possible values 0, 1, and 2. Viewed as a function, X assigns the value 2 to the outcome HH , 1 to the outcomes HT and TH , and 0 to the outcome TT . That is,

$$X(HH) = 2, X(HT) = X(TH) = 1, X(TT) = 0.$$

- Let Y be the number of Tails. In terms of X , we have $Y = 2 - X$. In other words, Y and $2 - X$ are the same r.v.: $Y(s) = 2 - X(s)$ for all s .
- Let I be 1 if the first toss lands Heads and 0 otherwise. Then I assigns the value 1 to the outcomes HH and HT and 0 to the outcomes TH and TT . This r.v. is an example of what is called an indicator random variable since it indicates whether the first toss lands Heads, using 1 to mean "yes" and 0 to mean "no".

We can also encode the sample space as $\{(1, 1), (1, 0), (0, 1), (0, 0)\}$, where 1 is the code for Heads and 0 is the code for Tails. Then we can give explicit formulas for X, Y, I :

$$X(s_1, s_2) = s_1 + s_2, Y(s_1, s_2) = 2 - s_1 - s_2, I(s_1, s_2) = s_1$$

where for simplicity we write $X(s_1, s_2)$ to mean $X((s_1, s_2))$, etc.

For most r.v.s we will consider, it is tedious or infeasible to write down an explicit formula in this way. Fortunately, it is usually unnecessary to do so, since (as we saw in this example) there are other ways to define an r.v., and (as we will see throughout the rest of this book) there are many ways to study the properties of an r.v. other than by doing computations with an explicit formula for what it maps each outcome s to.

As in the previous chapters, for a sample space with a finite number of outcomes we can visualize the outcomes as pebbles, with the mass of a pebble corresponding to its probability, such that the total mass of the pebbles is 1. A random variable simply labels each pebble with a number. Figure 3.2 shows two random variables defined on the same sample space: the pebbles or outcomes are the same, but the real numbers assigned to the outcomes are different.




FIGURE 3.2

Two random variables defined on the same sample space.

As we've mentioned earlier, the source of the randomness in a random variable is the experiment itself, in which a sample outcome $s \in S$ is chosen according to a probability function P . Before we perform the experiment, the outcome s has not yet been realized, so we don't know the value of X , though we could calculate the probability that X will take on a given value or range of values. After we perform the experiment and the outcome s has been realized, the random variable crystallizes into the numerical value $X(s)$.

Random variables provide numerical summaries of the experiment in question. This is very handy because the sample space of an experiment is often incredibly complicated or high-dimensional, and the outcomes $s \in S$ may be non-numeric. For example, the experiment may be to collect a random sample of people in a certain city and ask them various questions, which may have numeric (e.g., age or height) or non-numeric (e.g., political party or favorite movie) answers. The fact that r.v.s take on numerical values is a very convenient simplification compared to having to work with the full complexity of S at all times.

4.3.2 3.2 Distributions and Probability Mass Functions

There are two main types of random variables used in practice: discrete r.v.s and continuous r.v.s. In this chapter and the next, our focus is on discrete r.v.s. Continuous r.v.s are introduced in Chapter 5.

Definition 3.2.1 (Discrete random variable). A random variable X is said to be discrete if there is a finite list of values $a_1, a_2, \dots, a_n$ or an infinite list of values $a_1, a_2, \dots$ such that $P(X = a_j \text{ for some } j) = 1$. If X is a discrete r.v., then the finite or countably infinite set of values x such that $P(X = x) > 0$ is called the support of X .

Most commonly in applications, the support of a discrete r.v. is a set of integers. In contrast, a continuous r.v. can take on any real value in an interval (possibly even the entire real line); such r.v.s are defined more precisely in Chapter 5. It is also possible to have an r.v. that is a hybrid of discrete and continuous, such as by flipping a coin and then generating a discrete r.v. if the coin lands Heads and generating a continuous r.v. if the coin lands Tails. But the starting point for understanding such r.v.s is to understand discrete and continuous r.v.s.

Given a random variable, we would like to be able to describe its behavior using the language of probability. For example, we might want to answer questions about the probability that the r.v. will fall into a given range: if L is the lifetime earnings of a randomly chosen U.S. college graduate, what is the probability that L exceeds a million dollars? If M is the number of major earthquakes in California in the next five years, what is the probability that M equals 0?

The distribution of a random variable provides the answers to these questions; it specifies the probabilities of all events associated with the r.v., such as the probability of it equaling 3 and the probability of it being at least 110. We will see that there are several equivalent ways to express the distribution of an r.v. For a discrete r.v., the most natural way to do so is with a probability mass function, which we now define.

Definition 3.2.2 (Probability mass function). The probability mass function (PMF) of a discrete r.v. X is the function p_X given by $p_X(x) = P(X = x)$. Note that this is positive if x is in the support of X , and 0 otherwise.

3.2.3. In writing $P(X = x)$, we are using $X = x$ to denote an event, consisting of all outcomes s to which X assigns the number x . This event is also written as $\{X = x\}$; formally, $\{X = x\}$ is defined as $\{s \in S : X(s) = x\}$, but writing $\{X = x\}$ is shorter and more intuitive. Going back to Example 3.1.2, if X is the number of Heads in two fair coin tosses, then $\{X = 1\}$ consists of the sample outcomes HT and TH , which are the two outcomes to which X assigns the number 1. Since $\{HT, TH\}$ is a subset of the sample space, it is an event. So it makes sense to talk about $P(X = 1)$, or more generally, $P(X = x)$. If $\{X = x\}$ were anything other than an event, it would make no sense to calculate its probability! It does not make sense to write " $P(X)$ "; we can only take the probability of an event, not of an r.v.

Let's look at a few examples of PMFs.

Example 3.2.4 (Coin tosses continued). In this example we'll find the PMFs of all the random variables in Example 3.1.2, the example with two fair coin tosses. Here are the r.v.s we defined, along with their PMFs:

- X , the number of Heads. Since X equals 0 if TT occurs, 1 if HT or TH occurs, and 2 if HH occurs, the PMF of X is the function p_X given by

$$p_X(0) = P(X = 0) = 1/4,$$

$$p_X(1) = P(X = 1) = 1/2,$$

$$p_X(2) = P(X = 2) = 1/4,$$

and $p_X(x) = 0$ for all other values of x .

- $Y = 2 - X$, the number of Tails. Reasoning as above or using the fact that

$$P(Y = y) = P(2 - X = y) = P(X = 2 - y) = p_X(2 - y),$$

the PMF of Y is

$$p_Y(0) = P(Y = 0) = 1/4$$

$$p_Y(1) = P(Y = 1) = 1/2$$

$$p_Y(2) = P(Y = 2) = 1/4$$

and $p_Y(y) = 0$ for all other values of y .

Note that X and Y have the same PMF (that is, p_X and p_Y are the same function) even though X and Y are not the same r.v. (that is, X and Y are two different functions from $\{HH, HT, TH, TT\}$ to the real line).

- I , the indicator of the first toss landing Heads. Since I equals 0 if TH or TT occurs and 1 if HH or HT occurs, the PMF of I is

$$p_I(0) = P(I = 0) = 1/2$$

$$p_I(1) = P(I = 1) = 1/2$$

and $p_I(i) = 0$ for all other values of i .




FIGURE 3.3

Left to right: PMFs of X, Y , and I , with X the number of Heads in two fair coin tosses, Y the number of Tails, and I the indicator of Heads on the first toss.

The PMFs of X, Y , and I are plotted in Figure 3.3. Vertical bars are drawn to make it easier to compare the heights of different points. $\square$

Example 3.2.5 (Sum of die rolls). We roll two fair 6-sided dice. Let $T = X + Y$ be the total of the two rolls, where X and Y are the individual rolls. The sample space of this experiment has 36 equally likely outcomes:

$$S = \{(1, 1), (1, 2), \dots, (6, 5), (6, 6)\}.$$

For example, 7 of the 36 outcomes s are shown in the table below, along with the corresponding values of X, Y , and T . After the experiment is performed, we observe values for X and Y , and then the observed value of T is the sum of those values.

s	X	Y	$X + Y$
(1, 2)	1	2	3
(1, 6)	1	6	7
(2, 5)	2	5	7
(3, 1)	3	1	4
(4, 3)	4	3	7
(5, 4)	5	4	9
(6, 6)	6	6	12

Since the dice are fair, the PMF of X is

$$P(X = j) = 1/6,$$

for $j = 1, 2, \dots, 6$ (and $P(X = j) = 0$ otherwise); we say that X has a Discrete Uniform distribution on $1, 2, \dots, 6$. Similarly, Y is also Discrete Uniform on $1, 2, \dots, 6$.

Note that Y has the same distribution as X but is not the same random variable as X . In fact, we have

$$P(X = Y) = 6/36 = 1/6.$$

Two more r.v.s in this experiment with the same distribution as X are $7 - X$ and $7 - Y$. To see this, we can use the fact that for a standard die, $7 - X$ is the value on the bottom if X is the value on the top. If the top value is equally likely to be any of the numbers $1, 2, \dots, 6$, then so is the bottom value. Note that even though $7 - X$ has the same distribution as X , it is never equal to X in a run of the experiment!

Let's now find the PMF of T . By the naive definition of probability,

$$P(T = 2) = P(T = 12) = 1/36$$

$$P(T = 3) = P(T = 11) = 2/36$$

$$P(T = 4) = P(T = 10) = 3/36$$

$$P(T = 5) = P(T = 9) = 4/36$$

$$P(T = 6) = P(T = 8) = 5/36$$

$$P(T = 7) = 6/36$$

For all other values of t , $P(T = t) = 0$. We can see directly that the support of T is $\{2, 3, \dots, 12\}$ just by looking at the possible totals for two dice, but as a check, note that

$$P(T = 2) + P(T = 3) + \dots + P(T = 12) = 1,$$

which shows that all possibilities have been accounted for. The symmetry property of T that appears above, $P(T = t) = P(T = 14 - t)$, makes sense since each outcome $\{X = x, Y = y\}$ which makes $T = t$ has a corresponding outcome $\{X = 7 - x, Y = 7 - y\}$ of the same probability which makes $T = 14 - t$.




FIGURE 3.4
PMF of the sum of two die rolls.

The PMF of T is plotted in Figure 3.4; it has a triangular shape, and the symmetry noted above is very visible. $\square$

Example 3.2.6 (Children in a U.S. household). Suppose we choose a household in the United States at random. Let X be the number of children in the chosen household. Since X can only take on integer values, it is a discrete r.v. The probability that X takes on the value x is proportional to the number of households in the United States with x children.

Using data from the 2010 General Social Survey [23], we can approximate the proportion of households with 0 children, 1 child, 2 children, etc., and hence approximate the PMF of X , which is plotted in Figure 3.5. $\square$

We will now state the properties of a valid PMF.

Theorem 3.2.7 (Valid PMFs). Let X be a discrete r.v. with support $x_1, x_2, \dots$ (assume these values are distinct and, for notational simplicity, that the support is countably infinite; the analogous results hold if the support is finite). The PMF p_X of X must satisfy the following two criteria:

- Nonnegative: $p_X(x) > 0$ if $x = x_j$ for some j , and $p_X(x) = 0$ otherwise;
- Sums to 1: $\sum_{j=1}^{\infty} p_X(x_j) = 1$.




FIGURE 3.5

PMF of the number of children in a randomly selected U.S. household.

Proof. The first criterion is true since probability is nonnegative. The second is true since X must take on some value, and the events $\{X = x_j\}$ are disjoint, so

$$\sum_{j=1}^{\infty} P(X = x_j) = P\left(\bigcup_{j=1}^{\infty} \{X = x_j\}\right) = P(X = x_1 \text{ or } X = x_2 \text{ or } \dots) = 1.$$

Conversely, if distinct values $x_1, x_2, \dots$ are specified and we have a function satisfying the two criteria above, then this function is the PMF of some r.v.; we will show how to construct such an r.v. in Chapter 5.

We claimed earlier that the PMF is one way of expressing the distribution of a discrete r.v. This is because once we know the PMF of X , we can calculate the probability that X will fall into a given subset of the real numbers by summing over the appropriate values of x , as the next example shows.

Example 3.2.8. Returning to Example 3.2.5, let T be the sum of two fair die rolls. We have already calculated the PMF of T . Now suppose we're interested in the probability that T is in the interval $[1, 4]$. There are only three values in the interval $[1, 4]$ that T can take on, namely, 2, 3, and 4. We know the probability of each of these values from the PMF of T , so

$$P(1 \leq T \leq 4) = P(T = 2) + P(T = 3) + P(T = 4) = 6/36.$$

In general, given a discrete r.v. X and a set B of real numbers, if we know the PMF of X we can find $P(X \in B)$, the probability that X is in B , by summing up the heights of the vertical bars at points in B in the plot of the PMF of X . Knowing the PMF of a discrete r.v. determines its distribution.

4.3.3 3.3 Bernoulli and Binomial

Some distributions are so ubiquitous in probability and statistics that they have their own names. We will introduce these named distributions throughout the book, starting with a very simple but useful case: an r.v. that can take on only two possible values, 0 and 1.

Definition 3.3.1 (Bernoulli distribution). An r.v. X is said to have the Bernoulli distribution with parameter p if $P(X = 1) = p$ and $P(X = 0) = 1 - p$, where $0 < p < 1$. We write this as $X \sim \text{Bern}(p)$. The symbol $\sim$ is read "is distributed as".

Any r.v. whose possible values are 0 and 1 has a $\text{Bern}(p)$ distribution, with p the probability of the r.v. equaling 1. This number p in $\text{Bern}(p)$ is called the parameter of the distribution; it determines which specific Bernoulli distribution we have. Thus there is not just one Bernoulli distribution, but rather a family of Bernoulli distributions, indexed by p . For example, if $X \sim \text{Bern}(1/3)$, it would be correct but incomplete to say " X is Bernoulli"; to fully specify the distribution of X , we should both say its name (Bernoulli) and its parameter value ($1/3$), which is the point of the notation $X \sim \text{Bern}(1/3)$.

Any event has a Bernoulli r.v. that is naturally associated with it, equal to 1 if the event happens and 0 otherwise. This is called the indicator random variable of the event; we will see that such r.v.s are extremely useful.

Definition 3.3.2 (Indicator random variable). The indicator random variable of an event A is the r.v. which equals 1 if A occurs and 0 otherwise. We will denote the indicator r.v. of A by I_A or $I(A)$. Note that $I_A \sim \text{Bern}(p)$ with $p = P(A)$.

We often imagine Bernoulli r.v.s using coin tosses, but this is just convenient language for discussing the following general story.

Story 3.3.3 (Bernoulli trial). An experiment that can result in either a "success" or a "failure" (but not both) is called a Bernoulli trial. A Bernoulli random variable can be thought of as the indicator of success in a Bernoulli trial: it equals 1 if success occurs and 0 if failure occurs in the trial.

Because of this story, the parameter p is often called the success probability of the $\text{Bern}(p)$ distribution. Once we start thinking about Bernoulli trials, it's hard not to start thinking about what happens when we have more than one trial.

Story 3.3.4 (Binomial distribution). Suppose that n independent Bernoulli trials are performed, each with the same success probability p . Let X be the number of successes. The distribution of X is called the Binomial distribution with parameters n and p . We write $X \sim \text{Bin}(n, p)$ to mean that X has the Binomial distribution with parameters n and p , where n is a positive integer and $0 < p < 1$.

Notice that we define the Binomial distribution not by its PMF, but by a story about the type of experiment that could give rise to a random variable with a Binomial distribution. The most famous distributions in statistics all have stories which explain why they are so often used as models for data, or as the building blocks for more complicated distributions.

Thinking about the named distributions first and foremost in terms of their stories has many benefits. It facilitates pattern recognition, allowing us to see when two problems are essentially identical in structure; it often leads to cleaner solutions that avoid PMF calculations altogether; and it helps us understand how the named distributions are connected to one another. Here it is clear that $\text{Bern}(p)$ is the same distribution as $\text{Bin}(1, p)$: the Bernoulli is a special case of the Binomial.

Using the story definition of the Binomial, let's find its PMF.

Theorem 3.3.5 (Binomial PMF). If $X \sim \text{Bin}(n, p)$, then the PMF of X is

$$P(X = k) = \binom{n}{k} p^k (1 - p)^{n-k}$$

for $k = 0, 1, \dots, n$ (and $P(X = k) = 0$ otherwise).

3.3.6. To save writing, it is often left implicit that a PMF is zero wherever it is not specified to be nonzero, but in any case it is important to understand what the support of a random variable is, and good practice to check that PMFs are valid. If two discrete r.v.s have the same PMF, then they also must have the same support. So we sometimes refer to the support of a discrete distribution; this is the support of any r.v. with that distribution.

Proof. An experiment consisting of n independent Bernoulli trials produces a sequence of successes and failures. The probability of any specific sequence of k successes and $n - k$ failures is $p^k(1 - p)^{n-k}$. There are $\binom{n}{k}$ such sequences, since we just need to select where the successes are. Therefore, letting X be the number of successes,

$$P(X = k) = \binom{n}{k} p^k (1 - p)^{n-k}$$

for $k = 0, 1, \dots, n$, and $P(X = k) = 0$ otherwise. This is a valid PMF because it is nonnegative and it sums to 1 by the binomial theorem.

Figure 3.6 shows plots of the Binomial PMF for various values of n and p . Note that the PMF of the $\text{Bin}(10, 1/2)$ distribution is symmetric about 5, but when the success probability is not $1/2$, the PMF is skewed. For a fixed number of trials n , X tends to be larger when the success probability is high and lower when the success probability is low, as we would expect from the story of the Binomial distribution. Also recall that in any PMF plot, the sum of the heights of the vertical bars must be 1.

We've used Story 3.3.4 to find the $\text{Bin}(n, p)$ PMF. The story also gives us a straightforward proof of the fact that if X is Binomial, then $n - X$ is also Binomial.




FIGURE 3.6

Some Binomial PMFs. In the lower left, we plot the $\text{Bin}(100, 0.03)$ PMF between 0 and 10 only, as the probability of more than 10 successes is close to 0.

Theorem 3.3.7. Let $X \sim \text{Bin}(n, p)$, and $q = 1 - p$ (we often use q to denote the failure probability of a Bernoulli trial). Then $n - X \sim \text{Bin}(n, q)$.

Proof. Using the story of the Binomial, interpret X as the number of successes in n independent Bernoulli trials. Then $n - X$ is the number of failures in those trials. Interchanging the roles of success and failure, we have $n - X \sim \text{Bin}(n, q)$. Alternatively, we can check that $n - X$ has the $\text{Bin}(n, q)$ PMF. Let $Y = n - X$. The PMF of Y is

$$P(Y = k) = P(X = n - k) = \binom{n}{n - k} p^{n - k} q^k = \binom{n}{k} q^k p^{n - k},$$

for $k = 0, 1, \dots, n$.

Corollary 3.3.8. Let $X \sim \text{Bin}(n, p)$ with $p = 1/2$ and n even. Then the distribution of X is symmetric about $n/2$, in the sense that $P(X = n/2 + j) = P(X = n/2 - j)$ for all nonnegative integers j .

Proof. By Theorem 3.3.7, $n - X$ is also $\text{Bin}(n, 1/2)$, so

$$P(X = k) = P(n - X = k) = P(X = n - k)$$

for all nonnegative integers k . Letting $k = n/2 + j$, the desired result follows. This explains why the $\text{Bin}(10, 1/2)$ PMF is symmetric about 5 in Figure 3.6.

Example 3.3.9 (Coin tosses continued). Going back to Example 3.1.2, we now know that $X \sim \text{Bin}(2, 1/2)$, $Y \sim \text{Bin}(2, 1/2)$, and $I \sim \text{Bern}(1/2)$. Consistent with Theorem 3.3.7, X and $Y = 2 - X$ have the same distribution, and consistent with Corollary 3.3.8, the distribution of X (and of Y) is symmetric about 1.

4.3.4 3.4 Hypergeometric

If we have an urn filled with w white and b black balls, then drawing n balls out of the urn with replacement yields a $\text{Bin}(n, w/(w + b))$ distribution for the number of white balls obtained in n trials, since the draws are independent Bernoulli trials, each with probability $w/(w + b)$ of success. If we instead sample without replacement, as illustrated in Figure 3.7, then the number of white balls follows a Hypergeometric distribution.

Story 3.4.1 (Hypergeometric distribution). Consider an urn with w white balls and b black balls. We draw n balls out of the urn at random without replacement, such that all $\binom{w + b}{n}$ samples are equally likely. Let X be the number of white balls in the sample. Then X is said to have the Hypergeometric distribution with parameters w, b , and n ; we denote this by $X \sim \text{HGeom}(w, b, n)$.




FIGURE 3.7

Hypergeometric story. An urn contains $w = 6$ white balls and $b = 4$ black balls. We sample $n = 5$ without replacement. The number X of white balls in the sample is Hypergeometric; here we observe $X = 3$.

As with the Binomial distribution, we can obtain the PMF of the Hypergeometric distribution from the story.

Theorem 3.4.2 (Hypergeometric PMF). If $X \sim \text{HGeom}(w, b, n)$, then the PMF of X is

$$P(X = k) = \frac{\binom{w}{k} \binom{b}{n-k}}{\binom{w+b}{n}},$$

for integers k satisfying $0 \leq k \leq w$ and $0 \leq n - k \leq b$, and $P(X = k) = 0$ otherwise.

Proof. To get $P(X = k)$, we first count the number of possible ways to draw exactly k white balls and $n - k$ black balls from the urn (without distinguishing between different orderings for getting the same set of balls). If $k > w$ or $n - k > b$, then the draw is impossible. Otherwise, there are $\binom{w}{k} \binom{b}{n-k}$ ways to draw k white and $n - k$ black balls by the multiplication rule, and there are $\binom{w+b}{n}$ total ways to draw n balls. Since all samples are equally likely, the naive definition of probability gives

$$P(X = k) = \frac{\binom{w}{k} \binom{b}{n-k}}{\binom{w+b}{n}}$$

for integers k satisfying $0 \leq k \leq w$ and $0 \leq n - k \leq b$. This PMF is valid because the numerator, summed over all k , equals $\binom{w+b}{n}$ by Vandermonde's identity (Example 1.5.3), so the PMF sums to 1. $\square$

The Hypergeometric distribution comes up in many scenarios which, on the surface, have little in common with white and black balls in an urn. The essential structure of the Hypergeometric story is that items in a population are classified using two sets of tags: in the urn story, each ball is either white or black (this is the first set of tags), and each ball is either sampled or not sampled (this is the second set of tags). Furthermore, at least one of these sets of tags is assigned completely at random (in the urn story, the balls are sampled randomly, with all sets of the correct size equally likely). Then $X \sim \text{HGeom}(w, b, n)$ represents the number of twice-tagged items: in the urn story, balls that are both white and sampled.

The next two examples show seemingly dissimilar scenarios that are nonetheless isomorphic to the urn story.

Example 3.4.3 (Elk capture-recapture). A forest has N elk. Today, m of the elk are captured, tagged, and released into the wild. At a later date, n elk are recaptured at random. Assume that the recaptured elk are equally likely to be any set of n of the elk, e.g., an elk that has been captured does not learn how to avoid being captured again.

By the story of the Hypergeometric, the number of tagged elk in the recaptured sample is $\text{HGeom}(m, N - m, n)$. The m tagged elk in this story correspond to the white balls and the $N - m$ untagged elk correspond to the black balls. Instead of sampling n balls from the urn, we recapture n elk from the forest. $\square$

Example 3.4.4 (Aces in a poker hand). In a five-card hand drawn at random from a well-shuffled standard deck, the number of aces in the hand has the $\text{HGeom}(4, 48, 5)$ distribution, which can be seen by thinking of the aces as white balls and the nonaces as black balls. Using the Hypergeometric PMF, the probability that the hand has exactly three aces is

$$\frac{\binom{4}{3} \binom{48}{2}}{\binom{52}{5}} \approx 0.0017.$$

$\square$

The following table summarizes how the above examples can be thought of in terms of two sets of tags. In each example, the r.v. of interest is the number of items falling into both the second and the fourth columns: white and sampled, tagged and recaptured, ace and in one's hand.

Story	First set of tags		Second set of tags	
urn	white	black	sampled	not sampled
elk	tagged	untagged	recaptured	not recaptured
cards	ace	not ace	in hand	not in hand

The next theorem describes a symmetry between two Hypergeometric distributions with different parameters; the proof follows from swapping the two sets of tags in the Hypergeometric story.

Theorem 3.4.5. The $\text{HGeom}(w, b, n)$ and $\text{HGeom}(n, w + b - n, w)$ distributions are identical. That is, if $X \sim \text{HGeom}(w, b, n)$ and $Y \sim \text{HGeom}(n, w + b - n, w)$, then X and Y have the same distribution.

Proof. Using the story of the Hypergeometric, imagine an urn with w white balls, b black balls, and a sample of size n made without replacement. Let $X \sim \text{HGeom}(w, b, n)$ be the number of white balls in the sample, thinking of white/black as the first set of tags and sampled/not sampled as the second set of tags. Let $Y \sim \text{HGeom}(n, w + b - n, w)$ be the number of sampled balls among the white balls, thinking of sampled/not sampled as the first set of tags and white/black as

the second set of tags. Both X and Y count the number of white sampled balls, so they have the same distribution.

Alternatively, we can check algebraically that X and Y have the same PMF:

$$P(X = k) = \frac{\binom{w}{k} \binom{b}{n-k}}{\binom{w+b}{n}} = \frac{w!b!n!(w+b-n)!}{k!(w+b)!(w-k)!(n-k)!(b-n+k)!},$$

$$P(Y = k) = \frac{\binom{n}{k} \binom{w+b-n}{w-k}}{\binom{w+b}{w}} = \frac{w!b!n!(w+b-n)!}{k!(w+b)!(w-k)!(n-k)!(b-n+k)!}.$$

We prefer the story proof because it is less tedious and more memorable.

3.4.6 (Binomial vs. Hypergeometric). The Binomial and Hypergeometric distributions are often confused. Both are discrete distributions taking on integer values between 0 and n for some n , and both can be interpreted as the number of successes in n Bernoulli trials (for the Hypergeometric, each tagged elk in the recaptured sample can be considered a success and each untagged elk a failure). However, a crucial part of the Binomial story is that the Bernoulli trials involved are independent. The Bernoulli trials in the Hypergeometric story are dependent, since the sampling is done without replacement: knowing that one elk in our sample is tagged decreases the probability that the second elk will also be tagged.

4.3.5 3.5 Discrete Uniform

A very simple story, closely connected to the naive definition of probability, describes picking a random number from some finite set of possibilities.

Story 3.5.1 (Discrete Uniform distribution). Let C be a finite, nonempty set of numbers. Choose one of these numbers uniformly at random (i.e., all values in C are equally likely). Call the chosen number X . Then X is said to have the Discrete Uniform distribution with parameter C ; we denote this by $X \sim \text{DUnif}(C)$.

The PMF of $X \sim \text{DUnif}(C)$ is

$$P(X = x) = \frac{1}{|C|}$$

for $x \in C$ (and 0 otherwise), since a PMF must sum to 1. As with questions based on the naive definition of probability, questions based on a Discrete Uniform distribution reduce to counting problems. Specifically, for $X \sim \text{DUnif}(C)$ and any $A \subseteq C$, we have

$$P(X \in A) = \frac{|A|}{|C|}$$

Example 3.5.2 (Random slips of paper). There are 100 slips of paper in a hat, each of which has one of the numbers $1, 2, \dots, 100$ written on it, with no number appearing more than once. Five of the slips are drawn, one at a time.

First consider random sampling with replacement (with equal probabilities).

- (a) What is the distribution of how many of the drawn slips have a value of at least 80 written on them?
- (b) What is the distribution of the value of the j th draw (for $1 \leq j \leq 5$)?
- (c) What is the probability that the number 100 is drawn at least once?

Now consider random sampling without replacement (with all sets of five slips equally likely to be chosen).

- (d) What is the distribution of how many of the drawn slips have a value of at least 80 written on them?
- (e) What is the distribution of the value of the j th draw (for $1 \leq j \leq 5$)?
- (f) What is the probability that the number 100 is drawn in the sample?

Solution:

- (a) By the story of the Binomial, the distribution is $\text{Bin}(5, 0.21)$.
- (b) Let X_j be the value of the j th draw. By symmetry, $X_j \sim \text{DUnif}(1, 2, \dots, 100)$. There aren't certain slips that love being chosen on the j th draw and others that avoid being chosen then; all are equally likely.
- (c) Taking complements,

$$P(X_j = 100 \text{ for at least one } j) = 1 - P(X_1 \neq 100, \dots, X_5 \neq 100).$$

By the naive definition of probability, this is

$$1 - (99/100)^5 \approx 0.049$$

This solution just uses new notation for concepts from Chapter 1. It is useful to have this new notation since it is compact and flexible. In the above calculation, it is important to see why

$$P(X_1 \neq 100, \dots, X_5 \neq 100) = P(X_1 \neq 100) \dots P(X_5 \neq 100).$$

This follows from the naive definition in this case, but a more general way to think about such statements is through independence of r.v.s, a concept discussed in detail in Section 3.8.

- (d) By the story of the Hypergeometric, the distribution is $\text{HGeom}(21, 79, 5)$.
- (e) Let Y_j be the value of the j th draw. By symmetry, $Y_j \sim \text{DUnif}(1, 2, \dots, 100)$.

Learning any Y_i gives information about the other values (so $Y_1, \dots, Y_5$ are not independent, as defined in Section 3.8), but symmetry still holds since, unconditionally, the j th slip drawn is equally likely to be any of the slips. This is the unconditional distribution of Y_j : we are working from a vantage point before drawing any of the slips.

For further insight into why each of $Y_1, \dots, Y_5$ is Discrete Uniform and how to think about Y_j unconditionally, imagine that instead of one person drawing five slips, one at a time, there are five people who draw one slip each, all reaching into the hat simultaneously, with all possibilities equally likely for who gets which slip. This formulation does not change the problem in any important way, and it helps avoid getting distracted by irrelevant chronological details. Label the five people $1, 2, \dots, 5$ in some way, e.g., from youngest to oldest, and let Z_j be the value

drawn by person j . By symmetry, $Z_j \sim \text{DUnif}(1, 2, \dots, 100)$ for each j ; the Z_j 's are dependent but, looked at individually, each person is drawing a uniformly random slip.

(f) The events $Y_1 = 100, \dots, Y_5 = 100$ are disjoint since we are now sampling without replacement, so

$$P(Y_j = 100 \text{ for some } j) = P(Y_1 = 100) + \dots + P(Y_5 = 100) = 0.05.$$

Sanity check: This answer makes sense intuitively since we can just as well think of first choosing five random slips out of 100 blank slips and then randomly writing the numbers from 1 to 100 on the slips, which gives a $5/100$ chance that the number 100 is on one of the five chosen slips.

It would be bizarre if the answer to (c) were greater than or equal to the answer to (f), since sampling without replacement makes it easier to find the number 100. (For the same reason, when searching for a lost possession it makes more sense to sample locations without replacement than with replacement.) But it makes sense that the answer to (c) is only slightly less than the answer to (f), since it is unlikely in (c) that the same slip will be sampled more than once (though by the birthday problem it's less unlikely than many people would guess).

More generally, if k slips are drawn without replacement, where $0 \leq k \leq 100$, then the same reasoning gives that the probability of drawing the number 100 is $k/100$. Note that this makes sense in the extreme case $k = 100$, since in that case we draw all of the slips.

4.3.6 3.6 Cumulative Distribution Functions

Another function that describes the distribution of an r.v. is the cumulative distribution function (CDF). Unlike the PMF, which only discrete r.v.s possess, the CDF is defined for all r.v.s.

Definition 3.6.1. The cumulative distribution function (CDF) of an r.v. X is the function F_X given by $F_X(x) = P(X \leq x)$. When there is no risk of ambiguity, we sometimes drop the subscript and just write F (or some other letter) for a CDF.

The next example demonstrates that for discrete r.v.s, we can freely convert between CDF and PMF.

Example 3.6.2. Let $X \sim \text{Bin}(4, 1/2)$. Figure 3.8 shows the PMF and CDF of X .




FIGURE 3.8 Bin(4, 1/2) PMF and CDF. The height of the vertical bar $P(X = 2)$ in the PMF is also the height of the jump in the CDF at 2 .

- From PMF to CDF: To find $P(X \leq 1.5)$, which is the CDF evaluated at 1.5 , we sum the PMF over all values of the support that are less than or equal to 1.5 :

$$P(X \leq 1.5) = P(X = 0) + P(X = 1) = \left(\frac{1}{2}\right)^4 + 4 \left(\frac{1}{2}\right)^4 = \frac{5}{16}.$$

Similarly, the value of the CDF at an arbitrary point x is the sum of the heights of the vertical bars of the PMF at values less than or equal to x .

- From CDF to PMF: The CDF of a discrete r.v. consists of jumps and flat regions. The height of a jump in the CDF at x is equal to the value of the PMF at x . For example, in Figure 3.8, the height of the jump in the CDF at 2 is the same as the height of the corresponding vertical bar in the PMF; this is indicated in the figure with curly braces. The flat regions of the CDF correspond to values outside the support of X , so the PMF is equal to 0 in those regions. $\square$

Valid CDFs satisfy the following criteria.

Theorem 3.6.3 (Valid CDFs). Any CDF F has the following properties.

- Increasing: If $x_1 \leq x_2$, then $F(x_1) \leq F(x_2)$.
- Right-continuous: As in Figure 3.8, the CDF is continuous except possibly for having some jumps. Wherever there is a jump, the CDF is continuous from the right. That is, for any a , we have

$$F(a) = \lim_{x \rightarrow a^+} F(x).$$

- Convergence to 0 and 1 in the limits:

$$\lim_{x \rightarrow -\infty} F(x) = 0 \text{ and } \lim_{x \rightarrow \infty} F(x) = 1.$$

Proof. The above criteria are true for all CDFs, but for simplicity we will only prove it for the case where F is the CDF of a discrete r.v. X whose possible values are $0, 1, 2, \dots$. As an example of how to visualize the criteria, consider Figure 3.8: the CDF shown there is increasing (with some flat regions), continuous from the right (it is continuous except at jumps, and each jump has an open dot at the bottom and a closed dot at the top), and it converges to 0 as $x \rightarrow -\infty$ and to 1 as $x \rightarrow \infty$ (in this example, it reaches 0 and 1 ; in some examples, one or both of these values may be approached but never reached).

The first criterion is true since the event $\{X \leq x_1\}$ is a subset of the event $\{X \leq x_2\}$, so $P(X \leq x_1) \leq P(X \leq x_2)$.

For the second criterion, note that

$$P(X \leq x) = P(X \leq \lfloor x \rfloor),$$

where $\lfloor x \rfloor$ is the greatest integer less than or equal to x . For example, $P(X \leq 4.9) = P(X \leq 4)$ since X is integer-valued. So $F(a + b) = F(a)$ for any $b > 0$ that is small enough so that $a + b < \lfloor a \rfloor + 1$, e.g., for $a = 4.9$, this holds for $0 < b < 0.1$. This implies $F(a) = \lim_{x \rightarrow a^+} F(x)$ (in fact, it's much stronger since it says $F(x)$ equals $F(a)$ when x is close enough to a and on the right).

For the third criterion, we have $F(x) = 0$ for $x < 0$, and

$$\lim_{x \rightarrow \infty} F(x) = \lim_{x \rightarrow \infty} P(X \leq \lfloor x \rfloor) = \lim_{x \rightarrow \infty} \sum_{n=0}^{\lfloor x \rfloor} P(X = n) = \sum_{n=0}^{\infty} P(X = n) = 1$$

The converse is true too: we will show in Chapter 5 that given any function F meeting these criteria, we can construct a random variable whose CDF is F .

To recap, we have now seen three equivalent ways of expressing the distribution of a random variable. Two of these are the PMF and the CDF: we know these two functions contain the same information, since we can always figure out the CDF from the PMF and vice versa.

Generally the PMF is easier to work with for discrete r.v.s, since evaluating the CDF requires a summation.

A third way to describe a distribution is with a story that explains (in a precise way) how the distribution can arise. We used the stories of the Binomial and Hypergeometric distributions to derive the corresponding PMFs. Thus the story and the PMF also contain the same information, though we can often achieve more intuitive proofs with the story than with PMF calculations.

4.3.7 3.7 Functions of Random Variables

In this section we will discuss what it means to take a function of a random variable, and we will build understanding for why a function of a random variable is a random variable. That is, if X is a random variable, then X^2 , e^X , and $\sin(X)$ are also random variables, as is $g(X)$ for any function $g : \mathbb{R} \rightarrow \mathbb{R}$.

For example, imagine that two basketball teams (A and B) are playing a seven game match, and let X be the number of wins for team A (so $X \sim \text{Bin}(7, 1/2)$ if the teams are evenly matched and the games are independent). Let $g(x) = 7 - x$, and let $h(x) = 1$ if $x \geq 4$ and $h(x) = 0$ if $x < 4$. Then $g(X) = 7 - X$ is the number of wins for team B , and $h(X)$ is the indicator of team A winning the majority of the games. Since X is an r.v., both $g(X)$ and $h(X)$ are also r.v.s.

To see how to define functions of an r.v. formally, let's rewind a bit. At the beginning of this chapter, we considered a random variable X defined on a sample space with 6 elements. Figure 3.1 used arrows to illustrate how X maps each pebble in the sample space to a real number, and the left half of Figure 3.2 showed how we can equivalently imagine X writing a real number inside each pebble.

Now we can, if we want, apply the same function g to all the numbers inside the pebbles. Instead of the numbers $X(s_1)$ through $X(s_6)$, we now have the numbers $g(X(s_1))$ through $g(X(s_6))$, which gives a new mapping from sample outcomes to real numbers—we've created a new random variable, $g(X)$.

Definition 3.7.1 (Function of an r.v.). For an experiment with sample space S , an r.v. X , and a function $g : \mathbb{R} \rightarrow \mathbb{R}$, $g(X)$ is the r.v. that maps s to $g(X(s))$ for all $s \in S$.

Taking $g(x) = \sqrt{x}$ for concreteness, Figure 3.9 shows that $g(X)$ is the composition of the functions X and g , saying "first apply X , then apply g ". Figure 3.10 represents $g(X)$ more succinctly by directly labeling the sample outcomes. Both figures show us that $g(X)$ is an r.v.; if X crystallizes to 4 , then $g(X)$ crystallizes to 2 .

Given a discrete r.v. X with a known PMF, how can we find the PMF of $Y = g(X)$? In the case where g is a one-to-one function, the answer is straightforward: the support of Y is the set of all $g(x)$ with x in the support of X , and

$$P(Y = g(x)) = P(g(X) = g(x)) = P(X = x)$$




FIGURE 3.9
The r.v. X is defined on a sample space with 6 elements, and has possible values 0 , 1 , and 4 . The function g is the square root function. Composing X and g gives the random variable $g(X)=\sqrt{X}$, which has possible values 0,1 , and 2 . }




FIGURE 3.10
Since $g(X)=\sqrt{X}$ labels each pebble with a number, it is an r.v. }

The case where $Y = g(X)$ with g one-to-one is illustrated in the following tables; the idea is that if the distinct possible values of X are $x_1, x_2, \dots$ with probabilities $p_1, p_2, \dots$ (respectively), then the distinct possible values of Y are $g(x_1), g(x_2), \dots$, with the same list $p_1, p_2, \dots$ of probabilities.

x	$P(X = x)$
x_1	p_1
x_2	p_2
x_3	p_3
$\vdots$	$\vdots$

PMF of X , in table form

y	$P(Y = y)$
$g(x_1)$	p_1
$g(x_2)$	p_2
$g(x_3)$	p_3
$\vdots$	$\vdots$

PMF of Y , in table form

This suggests a strategy for finding the PMF of an r.v. with an unfamiliar distribution: try to express the r.v. as a one-to-one function of an r.v. with a known distribution. The next example illustrates this method.

Example 3.7.2 (Random walk). A particle moves n steps on a number line. The particle starts at 0, and at each step it moves 1 unit to the right or to the left, with equal probabilities. Assume all steps are independent. Let Y be the particle's position after n steps. Find the PMF of Y .

Solution:

Consider each step to be a Bernoulli trial, where right is considered a success and left is considered a failure. Then the number of steps the particle takes to the right is a $\text{Bin}(n, 1/2)$ random variable, which we can name X . If $X = j$, then the particle has taken j steps to the right and $n - j$ steps to the left, giving a final position of $j - (n - j) = 2j - n$. So we can express Y as a one-to-one function of X , namely, $Y = 2X - n$. Since X takes values in $\{0, 1, 2, \dots, n\}$, Y takes values in $\{-n, 2 - n, 4 - n, \dots, n\}$.

The PMF of Y can then be found from the PMF of X :

$$P(Y = k) = P(2X - n = k) = P(X = (n + k)/2) = \binom{n}{\frac{n+k}{2}} \left(\frac{1}{2}\right)^n,$$

if k is an integer between $-n$ and n (inclusive) such that $n + k$ is an even number.

If g is not one-to-one, then for a given y , there may be multiple values of x such that $g(x) = y$. To compute $P(g(X) = y)$, we need to sum up the probabilities of X taking on any of these candidate values of x .

Theorem 3.7.3 (PMF of $g(X)$). Let X be a discrete r.v. and $g : \mathbb{R} \rightarrow \mathbb{R}$. Then the support of $g(X)$ is the set of all y such that $g(x) = y$ for at least one x in the support of X , and the PMF of $g(X)$ is

$$P(g(X) = y) = \sum_{x:g(x)=y} P(X = x),$$

for all y in the support of $g(X)$.

Example 3.7.4. Continuing as in the previous example, let D be the particle's distance from the origin after n steps. Assume that n is even. Find the PMF of D .

Solution:

We can write $D = |Y|$; this is a function of Y , but it isn't one-to-one. The event $D = 0$ is the same as the event $Y = 0$. For $k = 2, 4, \dots, n$, the event $D = k$ is the same as the event $\{Y = k\} \cup \{Y = -k\}$. So the PMF of D is

$$P(D = 0) = \binom{n}{\frac{n}{2}} \left(\frac{1}{2}\right)^n,$$

$$P(D = k) = P(Y = k) + P(Y = -k) = 2 \binom{n}{\frac{n+k}{2}} \left(\frac{1}{2}\right)^n,$$

for $k = 2, 4, \dots, n$. In the final step we used symmetry (imagine a new random walk that moves left each time our random walk moves right, and vice versa) to see that $P(Y = k) = P(Y = -k)$.

The same reasoning we have used to handle functions of one random variable can be extended to deal with functions of multiple random variables. We have already seen an example of this with the addition function (which maps two numbers x, y to their sum $x + y$): in Example 3.2.5, we saw how to view $T = X + Y$ as an r.v. in its own right, where X and Y are obtained by rolling dice.

Definition 3.7.5 (Function of two r.v.s). Given an experiment with sample space S , if X and Y are r.v.s that map $s \in S$ to $X(s)$ and $Y(s)$ respectively, then $g(X, Y)$ is the r.v. that maps s to $g(X(s), Y(s))$.

Note that we are assuming that X and Y are defined on the same sample space S . Usually we assume that S is chosen to be rich enough to encompass whatever r.v.s we wish to work with. For example, if X is based on a coin flip and Y is based on a die roll, and we initially were using the sample space $S_1 = \{H, T\}$ for X and the sample space $S_2 = \{1, 2, 3, 4, 5, 6\}$ for Y , we can easily redefine X and Y so that both are defined on the richer space $S = S_1 \times S_2 = \{(s_1, s_2) : s_1 \in S_1, s_2 \in S_2\}$.

One way to understand the mapping from S to $\mathbb{R}$ represented by the r.v. $g(X, Y)$ is with a table displaying the values of X, Y , and $g(X, Y)$ under various possible outcomes. Interpreting $X + Y$ as an r.v. is intuitive: if we observe $X = x$ and $Y = y$, then $X + Y$ crystallizes to $x + y$. For a less familiar example like $\max(X, Y)$, students often are unsure how to interpret it as an r.v. But the idea is the same: if we observe $X = x$ and $Y = y$, then $\max(X, Y)$ crystallizes to $\max(x, y)$.

Example 3.7.6 (Maximum of two die rolls). We roll two fair 6-sided dice. Let X be the number on the first die and Y the number on the second die. The following table gives the values of X, Y , and $\max(X, Y)$ under 7 of the 36 outcomes in the sample space, analogously to the table in Example 3.2.5.

s	X	Y	$\max(X, Y)$
(1, 2)	1	2	2
(1, 6)	1	6	6
(2, 5)	2	5	5
(3, 1)	3	1	3
(4, 3)	4	3	4
(5, 4)	5	4	5
(6, 6)	6	6	6

So $\max(X, Y)$ assigns a numerical value to each sample outcome. The PMF is

$$P(\max(X, Y) = 1) = 1/36,$$

$$P(\max(X, Y) = 2) = 3/36,$$

$$P(\max(X, Y) = 3) = 5/36,$$

$$P(\max(X, Y) = 4) = 7/36,$$

$$P(\max(X, Y) = 5) = 9/36,$$

$$P(\max(X, Y) = 6) = 11/36.$$

These probabilities can be obtained by tabulating the values of $\max(x, y)$ in a 6×6 grid and counting how many times each value appears in the grid, or with calculations such as

$$\begin{aligned} P(\max(X, Y) = 5) &= P(X = 5, Y \leq 4) + P(X \leq 4, Y = 5) + P(X = 5, Y = 5) \\ &= 2P(X = 5, Y \leq 4) + 1/36 \\ &= 2(4/36) + 1/36 = 9/36 \end{aligned}$$

□

3.7.7 (Category errors and sympathetic magic). Many common mistakes in probability can be traced to confusing two of the following fundamental objects with each other: distributions, random variables, events, and numbers. Such mistakes are examples of category errors. In general, a category error is a mistake that doesn't just happen to be wrong, but in fact is necessarily wrong since it is based on the wrong category of object. For example, answering the question "How many people live in Boston?" with " -42 " or " π " or "pink elephants" would be a category error—we may not know the population size of a city, but we do know that it is a nonnegative integer at any point in time. To help avoid being categorically wrong, always think about what category an answer should have.

An especially common category error is to confuse a random variable with its distribution. We call this error sympathetic magic; this term comes from anthropology, where it is used for the belief that one can influence an object by manipulating a representation of that object. The following saying sheds light on the distinction between a random variable and its distribution:

The word is not the thing; the map is not the territory. - Alfred Korzybski

We can think of the distribution of a random variable as a map or blueprint describing the r.v. Just as different houses can share the same blueprint, different r.v.s can have the same distribution, even if the experiments they summarize, and the sample spaces they map from, are not the same.

Here are two examples of sympathetic magic:

- Given an r.v. X , trying to get the PMF of $2X$ by multiplying the PMF of X by 2. It does not make sense to multiply a PMF by 2, since the probabilities would no longer sum to 1. As we saw above, if X takes on values x_j with probabilities p_j , then $2X$ takes on values $2x_j$ with probabilities p_j . Therefore the PMF of $2X$ is a horizontal stretch of the PMF of X ; it is not a vertical stretch, as would result from multiplying the PMF by 2. Figure 3.11 shows the PMF of a discrete r.v. X with support $\{0, 1, 2, 3, 4\}$, along with the PMF of $2X$, which has support $\{0, 2, 4, 6, 8\}$. Note that X can take on odd values, but $2X$ is necessarily even.




FIGURE 3.11
PMF of X (above) and PMF of $2X$ (below).

- Claiming that because X and Y have the same distribution, X must always equal Y , i.e., $P(X = Y) = 1$. Just because two r.v.s have the same distribution does not mean they are always equal, or ever equal. We saw this in Example 3.2.5. As another example, consider flipping a fair coin once. Let X be the indicator of Heads and $Y = 1 - X$ be the indicator of Tails. Both X and Y have the $\text{Bern}(1/2)$ distribution, but the event $X = Y$ is impossible. The PMFs of X and Y are the same function, but X and Y are different mappings from the sample space to the real numbers.

If Z is the indicator of Heads in a second flip (independent of the first flip), then Z is also $\text{Bern}(1/2)$, but Z is not the same r.v. as X . Here

$$P(Z = X) = P(HH \text{ or } TT) = 1/2$$

4.3.8 3.8 Independence of r.v.s

Just as we had the notion of independence of events, we can define independence of random variables. Intuitively, if two r.v.s X and Y are independent, then knowing the value of X gives no information about the value of Y , and vice versa. The definition formalizes this idea.

Definition 3.8.1 (Independence of two r.v.s). Random variables X and Y are said to be independent if

$$P(X \leq x, Y \leq y) = P(X \leq x)P(Y \leq y),$$

for all $x, y \in \mathbb{R}$.

In the discrete case, this is equivalent to the condition

$$P(X = x, Y = y) = P(X = x)P(Y = y),$$

for all x, y with x in the support of X and y in the support of Y .

The definition for more than two r.v.s is analogous.

Definition 3.8.2 (Independence of many r.v.s). Random variables $X_1, \dots, X_n$ are independent if

$$P(X_1 \leq x_1, \dots, X_n \leq x_n) = P(X_1 \leq x_1) \dots P(X_n \leq x_n),$$

for all $x_1, \dots, x_n \in \mathbb{R}$. For infinitely many r.v.s, we say that they are independent if every finite subset of the r.v.s is independent.

Comparing this to the criteria for independence of n events, it may seem strange that the independence of $X_1, \dots, X_n$ requires just one equality, whereas for events we needed to verify pairwise independence for all $\binom{n}{2}$ pairs, three-way independence for all $\binom{n}{3}$ triplets, and so on. However, upon closer examination of the definition, we see that independence of r.v.s requires the equality to hold for all possible $x_1, \dots, x_n$ – infinitely many conditions! If we can find even a single list of values $x_1, \dots, x_n$ for which the above equality fails to hold, then $X_1, \dots, X_n$ are not independent.

3.8.3. If $X_1, \dots, X_n$ are independent, then they are pairwise independent, i.e., X_i is independent of X_j for $i \neq j$. The idea behind proving that X_i and X_j are independent is to let all the x_k other than x_i, x_j go to ∞ in the definition of independence, since we already know $X_k < \infty$ is true (though it takes some work to give a complete justification for the limit). But pairwise independence does not imply independence in general, as we saw in Chapter 2 for events.

Example 3.8.4. In a roll of two fair dice, if X is the number on the first die and Y is the number on the second die, then $X + Y$ is not independent of $X - Y$ since

$$0 = P(X + Y = 12, X - Y = 1) \neq P(X + Y = 12)P(X - Y = 1) = \frac{1}{36} \cdot \frac{5}{36}$$

Knowing the total is 12 tells us the difference must be 0, so the r.v.s provide information about each other.

If X and Y are independent then it is also true, e.g., that X^2 is independent of Y^4 , since if X^2 provided information about Y^4 , then X would give information about Y (using X^2 and Y^4 as intermediaries: X determines X^2 , which would give information about Y^4 , which in turn would give information about Y). More generally, we have the following result (for which we omit a formal proof).

Theorem 3.8.5 (Functions of independent r.v.s). If X and Y are independent r.v.s, then any function of X is independent of any function of Y .

Definition 3.8.6 (i.i.d.). We will often work with random variables that are independent and have the same distribution. We call such r.v.s independent and identically distributed, or i.i.d. for short.

3.8.7 (i. vs. i.d.). "Independent" and "identically distributed" are two often confused but completely different concepts. Random variables are independent if they provide no information about each other; they are identically distributed if they have the same PMF (or equivalently, the same CDF). Whether two r.v.s are independent has nothing to do with whether they have the same distribution. We can have r.v.s that are:

- independent and identically distributed. Let X be the result of a die roll, and let Y be the result of a second, independent die roll. Then X and Y are i.i.d.
- independent and not identically distributed. Let X be the result of a die roll, and let Y be the closing price of the Dow Jones (a stock market index) a month from now. Then X and Y provide no information about each other (one would fervently hope), and X and Y do not have the same distribution.

- dependent and identically distributed. Let X be the number of Heads in n independent fair coin tosses, and let Y be the number of Tails in those same n tosses. Then X and Y are both distributed $\text{Bin}(n, 1/2)$, but they are highly dependent: if we know X , then we know Y perfectly.
- dependent and not identically distributed. Let X be the indicator of whether the majority party retains control of the House of Representatives in the U.S. after the next election, and let Y be the average favorability rating of the majority party in polls taken within a month of the election. Then X and Y are dependent, and X and Y do not have the same distribution.

By taking a sum of i.i.d. Bernoulli r.v.s, we can write down the story of the Binomial distribution in an algebraic form.

Theorem 3.8.8. If $X \sim \text{Bin}(n, p)$, viewed as the number of successes in n independent Bernoulli trials with success probability p , then we can write $X = X_1 + \cdots + X_n$ where the X_i are i.i.d. $\text{Bern}(p)$.

Proof. Let $X_i = 1$ if the i th trial was a success, and 0 if the i th trial was a failure. It's as though we have a person assigned to each trial, and we ask each person to raise their hand if their trial was a success. If we count the number of raised hands (which is the same as adding up the X_i), we get the total number of successes.

An important fact about the Binomial distribution is that the sum of independent Binomial r.v.s with the same success probability is also Binomial.

Theorem 3.8.9. If $X \sim \text{Bin}(n, p)$, $Y \sim \text{Bin}(m, p)$, and X is independent of Y , then $X + Y \sim \text{Bin}(n + m, p)$.

Proof. We present three proofs, since each illustrates a useful technique.

1. LOTP: We can directly find the PMF of $X + Y$ by conditioning on X (or Y , whichever we prefer) and using the law of total probability:

$$\begin{aligned}
 P(X + Y = k) &= \sum_{j=0}^k P(X + Y = k \mid X = j)P(X = j) \\
 &= \sum_{j=0}^k P(Y = k - j)P(X = j) \\
 &= \sum_{j=0}^k \binom{m}{k-j} p^{k-j} q^{m-k+j} \binom{n}{j} p^j q^{n-j} \\
 &= p^k q^{n+m-k} \sum_{j=0}^k \binom{m}{k-j} \binom{n}{j} \\
 &= \binom{n+m}{k} p^k q^{n+m-k}.
 \end{aligned}$$

In the second line, we used the independence of X and Y to justify dropping the conditioning in

$$P(X + Y = k \mid X = j) = P(Y = k - j \mid X = j) = P(Y = k - j),$$

and in the last line, we used the fact that

$$\sum_{j=0}^k \binom{m}{k-j} \binom{n}{j} = \binom{n+m}{k}$$

by Vandermonde's identity. The resulting expression is the $\text{Bin}(n+m, p)$ PMF, so $X+Y \sim \text{Bin}(n+m, p)$.

2. Representation: A much simpler proof is to represent both X and Y as the sum of i.i.d. $\text{Bern}(p)$ r.v.s: $X = X_1 + \cdots + X_n$ and $Y = Y_1 + \cdots + Y_m$, where the X_i and Y_j are all i.i.d. $\text{Bern}(p)$. Then $X+Y$ is the sum of $n+m$ i.i.d. $\text{Bern}(p)$ r.v.s, so its distribution, by the previous theorem, is $\text{Bin}(n+m, p)$.

3. Story: By the Binomial story, X is the number of successes in n independent trials and Y is the number of successes in m additional independent trials, all with the same success probability, so $X+Y$ is the total number of successes in the $n+m$ trials, which is the story of the $\text{Bin}(n+m, p)$ distribution. $\square$

Of course, if we have a definition for independence of r.v.s, we should have an analogous definition for conditional independence of r.v.s.

Definition 3.8.10 (Conditional independence of r.v.s). Random variables X and Y are conditionally independent given an r.v. Z if for all $x, y \in \mathbb{R}$ and all z in the support of Z ,

$$P(X \leq x, Y \leq y \mid Z = z) = P(X \leq x \mid Z = z)P(Y \leq y \mid Z = z).$$

For discrete r.v.s, an equivalent definition is to require

$$P(X = x, Y = y \mid Z = z) = P(X = x \mid Z = z)P(Y = y \mid Z = z).$$

As we might expect from the name, this is the definition of independence, except that we condition on $Z = z$ everywhere, and require the equality to hold for all z in the support of Z .

Definition 3.8.11 (Conditional PMF). For any discrete r.v.s X and Z , the function $P(X = x \mid Z = z)$, when considered as a function of x for fixed z , is called the conditional PMF of X given $Z = z$.

Independence of r.v.s does not imply conditional independence, nor vice versa. First let us show why independence does not imply conditional independence.

Example 3.8.12 (Matching pennies). Consider the simple game called matching pennies. Each of two players, A and B, has a fair penny. They flip their pennies independently. If the pennies match, A wins; otherwise, B wins. Let X be 1 if A's penny lands Heads and -1 otherwise, and define Y similarly for B (the r.v.s X and Y are called random signs).

Let $Z = XY$, which is 1 if A wins and -1 if B wins. Then X and Y are unconditionally independent, but given $Z = 1$, we know that $X = Y$ (the pennies match). So X and Y are conditionally dependent given Z . $\square$

Example 3.8.13 (Two friends). Consider again the "I have only two friends who ever call me" scenario from Example 2.5.11, except now with r.v. notation. Let X be the indicator of Alice calling me next Friday, Y be the indicator of Bob calling me next Friday, and Z be the indicator of exactly one of them calling me next Friday. Then X and Y are independent (by assumption). But given $Z = 1$, we have that X and Y are completely dependent: given that $Z = 1$, we have $Y = 1 - X$. $\square$

Next let's see why conditional independence does not imply independence.

Example 3.8.14 (Mystery opponent). Suppose that you are going to play two games of tennis against one of two identical twins. Against one of the twins, you are evenly matched, and against the other you have a $3/4$ chance of winning. Suppose that you can't tell which twin you are playing against until after the two games. Let Z be the indicator of playing against

the twin with whom you're evenly matched, and let X and Y be the indicators of victory in the first and second games, respectively.

Conditional on $Z = 1$, X and Y are i.i.d. $\text{Bern}(1/2)$, and conditional on $Z = 0$, X and Y are i.i.d. $\text{Bern}(3/4)$. So X and Y are conditionally independent given Z . Unconditionally, X and Y are dependent because observing $X = 1$ makes it more likely that we are playing the twin who is worse. That is,

$$P(Y = 1 \mid X = 1) > P(Y = 1).$$

Past games give us information which helps us infer who our opponent is, which in turn helps us predict future games! Note that this example is isomorphic to the "random coin" scenario from Example 2.3.7. $\square$

4.3.9 3.9 Connections between Binomial and Hypergeometric

The Binomial and Hypergeometric distributions are connected in two important ways. As we will see in this section, we can get from the Binomial to the Hypergeometric by conditioning, and we can get from the Hypergeometric to the Binomial by taking a limit. We'll start with a motivating example.

Example 3.9.1 (Fisher exact test). A scientist wishes to study whether women or men are more likely to have a certain disease, or whether they are equally likely. A random sample of n women and m men is gathered, and each person is tested for the disease (assume for this problem that the test is completely accurate). The numbers of women and men in the sample who have the disease are X and Y respectively, with $X \sim \text{Bin}(n, p_1)$ and $Y \sim \text{Bin}(m, p_2)$, independently. Here p_1 and p_2 are unknown, and we are interested in testing whether $p_1 = p_2$ (this is known as a null hypothesis in statistics).

Consider a 2×2 table with rows corresponding to disease status and columns corresponding to gender. Each entry is the count of how many people have that disease status and gender, so $n + m$ is the sum of all 4 entries. Suppose that it is observed that $X + Y = r$.

The Fisher exact test is based on conditioning on both the row and column sums, so n, m, r are all treated as fixed, and then seeing if the observed value of X is "extreme" compared to this conditional distribution. Assuming the null hypothesis, find the conditional PMF of X given $X + Y = r$.

Solution:

First we'll build the 2×2 table, treating n, m , and r as fixed.

	Women	Men	Total
Disease	x	$r - x$	r
No disease	$n - x$	$m - r + x$	$n + m - r$
Total	n	m	$n + m$

Next, let's compute the conditional PMF $P(X = x \mid X + Y = r)$. By Bayes' rule,

$$\begin{aligned} P(X = x \mid X + Y = r) &= \frac{P(X + Y = r \mid X = x)P(X = x)}{P(X + Y = r)} \\ &= \frac{P(Y = r - x)P(X = x)}{P(X + Y = r)}. \end{aligned}$$

The step $P(X + Y = r \mid X = x) = P(Y = r - x)$ is justified by the independence of X and Y . Assuming the null hypothesis and letting $p = p_1 = p_2$, we have $X \sim \text{Bin}(n, p)$ and $Y \sim \text{Bin}(m, p)$, independently, so $X + Y \sim \text{Bin}(n + m, p)$. Thus,

$$\begin{aligned}
 P(X = x \mid X + Y = r) &= \frac{\binom{m}{r-x} p^{r-x} (1-p)^{m-r+x} \binom{n}{x} p^x (1-p)^{n-x}}{\binom{n+m}{r} p^r (1-p)^{n+m-r}} \\
 &= \frac{\binom{n}{x} \binom{m}{r-x}}{\binom{n+m}{r}}
 \end{aligned}$$

So the conditional distribution of X is Hypergeometric with parameters n, m, r .

To understand why the Hypergeometric appeared, seemingly out of nowhere, let's connect this problem to the elk story for the Hypergeometric. In the elk story, we are interested in the distribution of the number of tagged elk in the recaptured sample. By analogy, think of women as tagged elk and men as untagged elk. Instead of recapturing r elk at random from the forest, we infect $X + Y = r$ people with the disease; under the null hypothesis, the set of diseased people is equally likely to be any set of r people. Thus, conditional on $X + Y = r$, X represents the number of women among the r diseased individuals. This is exactly analogous to the number of tagged elk in the recaptured sample, which is distributed $\text{HGeom}(n, m, r)$.

An interesting fact, which turns out to be useful in statistics, is that the conditional distribution of X does not depend on p : unconditionally, $X \sim \text{Bin}(n, p)$, but p disappears from the parameters of the conditional distribution! This makes sense upon reflection, since once we know $X + Y = r$, we can work directly with the fact that we have a population with r diseased and $n + m - r$ healthy people, without worrying about the value of p that originally generated the population.

This motivating example serves as a proof of the following theorem.

Theorem 3.9.2. If $X \sim \text{Bin}(n, p)$, $Y \sim \text{Bin}(m, p)$, and X is independent of Y , then the conditional distribution of X given $X + Y = r$ is $\text{HGeom}(n, m, r)$.

In the other direction, the Binomial is a limiting case of the Hypergeometric.

Theorem 3.9.3. If $X \sim \text{HGeom}(w, b, n)$ and $N = w + b \rightarrow \infty$ such that $p = w/(w + b)$ remains fixed, then the PMF of X converges to the $\text{Bin}(n, p)$ PMF.

Proof. We take the stated limit of the $\text{HGeom}(w, b, n)$ PMF :

$$\begin{aligned}
 P(X = k) &= \frac{\binom{w}{k} \binom{b}{n-k}}{\binom{w+b}{n}} \\
 &= \binom{n}{k} \frac{\binom{w+b-n}{w-k}}{\binom{w+b}{w}} \text{ by Theorem 3.4.5} \\
 &= \binom{n}{k} \frac{w!}{(w-k)!} \frac{b!}{(b-n+k)!} \frac{(w+b-n)!}{(w+b)!} \\
 &= \binom{n}{k} \frac{w(w-1)\dots(w-k+1)b(b-1)\dots(b-n+k+1)}{(w+b)(w+b-1)\dots(w+b-n+1)} \\
 &= \binom{n}{k} \frac{p(p-\frac{1}{N})\dots(p-\frac{k-1}{N})q(q-\frac{1}{N})\dots(q-\frac{n-k-1}{N})}{(1-\frac{1}{N})(1-\frac{2}{N})\dots(1-\frac{n-1}{N})}
 \end{aligned}$$

As $N \rightarrow \infty$, the denominator goes to 1, and the numerator goes to $p^k q^{n-k}$. Thus

$$P(X = k) \rightarrow \binom{n}{k} p^k q^{n-k}$$

which is the $\text{Bin}(n, p)$ PMF.

The stories of the Binomial and Hypergeometric provide intuition for this result: given an urn with w white balls and b black balls, the Binomial distribution arises from sampling n balls from the urn with replacement, while the Hypergeometric arises from sampling without replacement. As the number of balls in the urn grows very large relative to the

number of balls that are drawn, sampling with replacement and sampling without replacement become essentially equivalent. In practical terms, this theorem tells us that if $N = w + b$ is large relative to n , we can approximate the $\text{HGeom}(w, b, n)$ PMF by the $\text{Bin}(n, w/(w + b))$ PMF.

The birthday problem implies that it is surprisingly likely that some ball will be sampled more than once if sampling with replacement; for example, if 1,200 out of 1,000,000 balls are drawn randomly with replacement, then there is about a 51% chance that some ball will be drawn more than once! But this becomes less and less likely as N grows, and even if it is likely that there will be a few coincidences, the approximation can still be reasonable if it is very likely that the vast majority of balls in the sample are sampled only once each.

4.3.10 3.10 Recap

A random variable (r.v.) is a function assigning a real number to every possible outcome of an experiment. The distribution of an r.v. X is a full specification of the probabilities for the events associated with X , such as $\{X = 3\}$ and $\{1 \leq X \leq 5\}$. The distribution of a discrete r.v. can be defined using a PMF, a CDF, or a story. The PMF of X is the function $P(X = x)$ for $x \in \mathbb{R}$. The CDF of X is the function $P(X \leq x)$ for $x \in \mathbb{R}$. A story for X describes an experiment that could give rise to a random variable with the same distribution as X .

For a PMF to be valid, it must be nonnegative and sum to 1. For a CDF to be valid, it must be increasing, right-continuous, converge to 0 as $x \rightarrow -\infty$, and converge to 1 as $x \rightarrow \infty$.

It is important to distinguish between a random variable and its distribution: the distribution is a blueprint for building the r.v., but different r.v.s can have the same distribution, just as different houses can be built from the same blueprint.

Four named discrete distributions are the Bernoulli, Binomial, Hypergeometric, and Discrete Uniform. Each of these is actually a family of distributions, indexed by parameters; to fully specify one of these distributions, we need to give both the name and the parameter values.

- A $\text{Bern}(p)$ r.v. is the indicator of success in a Bernoulli trial with probability of success p .
- A $\text{Bin}(n, p)$ r.v. is the number of successes in n independent Bernoulli trials, all with the same probability p of success.
- A $\text{HGeom}(w, b, n)$ r.v. is the number of white balls obtained in a sample of size n drawn without replacement from an urn of w white and b black balls.
- A $\text{DUnif}(C)$ r.v. is obtained by randomly choosing an element of the finite set C , with equal probabilities for each element.

A function of a random variable is still a random variable. If we know the PMF of X , we can find $P(g(X) = k)$, the PMF of $g(X)$, by translating the event $\{g(X) = k\}$ into an equivalent event involving X , then using the PMF of X .

Two random variables are independent if knowing the value of one r.v. gives no information about the value of the other. This is unrelated to whether the two r.v.s are identically distributed. In Chapter 7, we will learn how to deal with dependent random variables by considering them jointly rather than separately.

We have now seen four fundamental types of objects in probability: distributions, random variables, events, and numbers. Figure 3.12 shows connections between these four fundamental objects. A CDF can be used as a blueprint for generating an r.v., and then there are various events describing the behavior of the r.v., such as the events $X \leq x$ for all x . Knowing the probabilities of these events determines the CDF, taking us full circle. For a discrete r.v. we can also use the PMF as a blueprint, and go from distribution to r.v. to events and back again.




FIGURE 3.12

Four fundamental objects in probability: distributions (blueprints), random variables, events, and numbers. From a CDF we can generate an r.v. X . From X , we can generate many other r.v.s by taking functions of X . There are various events describing the behavior of X . Most notably, for any constant x the events $X \leq x$ and $X = x$ are of interest. Knowing the probabilities of these events for all x gives us the CDF and (in the discrete case) the PMF, taking us full circle.

4.3.11 3.11r

Distributions in R

All of the named distributions that we'll encounter in this book have been implemented in R. In this section we'll explain how to work with the Binomial and Hypergeometric distributions in R. We will also explain in general how to generate r.v.s from any discrete distribution with a finite support. Typing `help(distributions)` gives a handy list of built-in distributions; many others are available through R packages that can be loaded.

In general, for many named discrete distributions, three functions starting with `d`, `p`, and `r` will give the PMF, CDF, and random generation, respectively. Note that the function starting with `p` is not the PMF, but rather is the CDF.

Binomial Distribution

The Binomial distribution is associated with the following three R functions: `dbinom`, `pbinom`, and `rbinom`. For the Bernoulli distribution we can just use the Binomial functions with $n = 1$.

- `dbinom` is the Binomial PMF. It takes three inputs: the first is the value of x at which to evaluate the PMF, and the second and third are the parameters n and p . For example, `dbinom(3, 5, 0.2)` returns the probability $P(X = 3)$ where $X \sim \text{Bin}(5, 0.2)$. In other words,

$$\text{dbinom}(3, 5, 0.2) = \binom{5}{3} (0.2)^3 (0.8)^2 = 0.0512$$

- `pbinom` is the Binomial CDF. It takes three inputs: the first is the value of x at which to evaluate the CDF, and the second and third are the parameters. `pbinom(3, 5, 0.2)` is the probability $P(X \leq 3)$ where $X \sim \text{Bin}(5, 0.2)$. So

$$\text{pbinom}(3, 5, 0.2) = \sum_{k=0}^3 \binom{5}{k} (0.2)^k (0.8)^{5-k} = 0.9933.$$

- `rbinom` is a function for generating Binomial random variables. For `rbinom`, the first input is how many r.v.s we want to generate, and the second and third inputs are still the parameters. Thus the command `rbinom(7, 5, 0.2)` produces realizations of seven i.i.d. $\text{Bin}(5, 0.2)$ r.v.s. When we ran this command, we got

```
2100100
```

but you'll probably get something different when you try it!

We can also evaluate PMFs and CDFs at an entire vector of values. For example, recall that `0:n` is a quick way to list the integers from 0 to n . The command `dbinom(0:5, 5, 0.2)` returns 6 numbers, $P(X=0), P(X=1), \dots, P(X=5)$, where $X \sim \text{Bin}(5, 0.2)$.

Hypergeometric Distribution

The Hypergeometric distribution also has three functions: `dhyp`, `phyper`, and `rhyper`. As one might expect, `dhyp` is the Hypergeometric PMF, `phyper` is the Hypergeometric CDF, and `rhyper` generates Hypergeometric r.v.s. Since the Hypergeometric distribution has three parameters, each of these functions takes four inputs. For `dhyp` and `phyper`, the first input is the value at which we wish to evaluate the PMF or CDF, and the remaining inputs are the parameters of the distribution.

Thus `dhyp(k, w, b, n)` returns $P(X=k)$ where $X \sim \text{HGeom}(w, b, n)$, and `phyper(k, w, b, n)` returns $P(X \leq k)$. For `rhyper`, the first input is the number of r.v.s we want to generate, and the remaining inputs are the parameters; `rhyper(100, w, b, n)` generates 100 i.i.d. $\text{HGeom}(w, b, n)$ r.v.s.

Discrete Distributions with Finite Support

We can generate r.v.s from any discrete distribution with finite support using the `sample` command. When we first introduced the `sample` command, we said that it can be used in the form `sample(n, k)` or `sample(n, k, replace=TRUE)` to sample k times from the integers 1 through n , either without or with replacement. For example, to generate 5 independent $\text{DUnif}(1, 2, \dots, 100)$ r.v.s, we can use the command `sample(100, 5, replace = TRUE)`.

It turns out that `sample` is far more versatile. If we want to sample from the values $x_1, \dots, x_n$ with probabilities $p_1, \dots, p_n$, we simply create a vector `x` containing all the x_i and a vector `p` containing all the p_i , then feed them into `sample`. Suppose we want realizations of i.i.d. r.v.s $X_1, \dots, X_{100}$ whose PMF is

$$\begin{aligned} P(X_j = 0) &= 0.25 \\ P(X_j = 1) &= 0.5 \\ P(X_j = 5) &= 0.1 \\ P(X_j = 10) &= 0.15 \end{aligned}$$

and $P(X_j = x) = 0$ for all other values of x . First, we use the `c` function to create vectors with the support of the distribution and the corresponding probabilities.

```
x <- c(0, 1, 5, 10)
p <- c(0.25, 0.5, 0.1, 0.15)
```

Next, we use sample. Here's how to get 100 draws from the PMF above:

```
sample( x, 100, prob=p, replace=TRUE)
```

The inputs are the vector x to sample from, the sample size (100 in this case), the probabilities p to use when sampling from x (if this is omitted, the probabilities are assumed equal), and whether to sample with replacement.

4.3.12 3.12 Exercises

Exercises marked with (S) have detailed solutions at <http://stat110.net>.

PMFs and Cdfs

1. People are arriving at a party one at a time. While waiting for more people to arrive they entertain themselves by comparing their birthdays. Let X be the number of people needed to obtain a birthday match, i.e., before person X arrives no two people have the same birthday, but when person X arrives there is a match. Find the PMF of X .
2. (a) Independent Bernoulli trials are performed, with probability $1/2$ of success, until there has been at least one success. Find the PMF of the number of trials performed.
(b) Independent Bernoulli trials are performed, with probability $1/2$ of success, until there has been at least one success and at least one failure. Find the PMF of the number of trials performed.
3. Let X be an r.v. with CDF F , and $Y = \mu + \sigma X$, where μ and σ are real numbers with $\sigma > 0$. (Then Y is called a location-scale transformation of X ; we will encounter this concept many times in Chapter 5 and beyond.) Find the CDF of Y , in terms of F .
4. Let n be a positive integer and

$$F(x) = \frac{\lfloor x \rfloor}{n}$$

for $0 \leq x \leq n$, $F(x) = 0$ for $x < 0$, and $F(x) = 1$ for $x > n$, where $\lfloor x \rfloor$ is the greatest integer less than or equal to x . Show that F is a CDF, and find the PMF that it corresponds to.

5. (a) Show that $p(n) = \left(\frac{1}{2}\right)^{n+1}$ for $n = 0, 1, 2, \dots$ is a valid PMF for a discrete r.v.
(b) Find the CDF of a random variable with the PMF from (a).
6. (S) Benford's law states that in a very large variety of real-life data sets, the first digit approximately follows a particular distribution with about a 30% chance of a 1, an 18% chance of a 2, and in general

$$P(D = j) = \log_{10} \left(\frac{j+1}{j} \right), \text{ for } j \in \{1, 2, 3, \dots, 9\}$$

where D is the first digit of a randomly chosen element. Check that this is a valid PMF (using properties of logs, not with a calculator).

7. Bob is playing a video game that has 7 levels. He starts at level 1, and has probability p_1 of reaching level 2. In general, given that he reaches level j , he has probability p_j of reaching level $j+1$, for $1 \leq j \leq 6$. Let X be the highest level that he reaches. Find the PMF of X (in terms of $p_1, \dots, p_6$).
8. There are 100 prizes, with one worth \$1, one worth \$2, ..., and one worth \$100. There are 100 boxes, each of which contains one of the prizes. You get 5 prizes by picking random boxes one at a time, without replacement. Find the PMF of how much your most valuable prize is worth (as a simple expression in terms of binomial coefficients).

9. Let F_1 and F_2 be CDFs, $0 < p < 1$, and $F(x) = pF_1(x) + (1 - p)F_2(x)$ for all x .
(a) Show directly that F has the properties of a valid CDF (see Theorem 3.6.3). The distribution defined by F is called a mixture of the distributions defined by F_1 and F_2 .
(b) Consider creating an r.v. in the following way. Flip a coin with probability p of Heads. If the coin lands Heads, generate an r.v. according to F_1 ; if the coin lands Tails, generate an r.v. according to F_2 . Show that the r.v. obtained in this way has CDF F .
10. (a) Is there a discrete distribution with support $1, 2, 3, \dots$, such that the value of the PMF at n is proportional to $1/n$?
Hint: See the math appendix for a review of some facts about series.
(b) Is there a discrete distribution with support $1, 2, 3, \dots$, such that the value of the PMF at n is proportional to $1/n^2$?
11. (S) Let X be an r.v. whose possible values are $0, 1, 2, \dots$, with CDF F . In some countries, rather than using a CDF, the convention is to use the function G defined by $G(x) = P(X < x)$ to specify a distribution. Find a way to convert from F to G , i.e., if F is a known function, show how to obtain $G(x)$ for all real x .
12. (a) Give an example of r.v.s X and Y such that $F_X(x) \leq F_Y(x)$ for all x , where the inequality is strict for some x . Here F_X is the CDF of X and F_Y is the CDF of Y . For the example you gave, sketch the CDFs of both X and Y on the same axes. Then sketch their PMFs on a second set of axes.
(b) In Part (a), you found an example of two different CDFs where the first is less than or equal to the second everywhere. Is it possible to find two different PMFs where the first is less than or equal to the second everywhere? In other words, find discrete r.v.s X and Y such that $P(X = x) \leq P(Y = x)$ for all x , where the inequality is strict for some x , or show that it is impossible to find such r.v.s.
13. Let X, Y, Z be discrete r.v.s such that X and Y have the same conditional distribution given Z , i.e., for all a and z we have

$$P(X = a \mid Z = z) = P(Y = a \mid Z = z)$$

Show that X and Y have the same distribution (unconditionally, not just when given Z).

14. Let X be the number of purchases that Fred will make on the online site for a certain company (in some specified time period). Suppose that the PMF of X is $P(X = k) = e^{-\lambda} \lambda^k / k!$ for $k = 0, 1, 2, \dots$. This distribution is called the Poisson distribution with parameter λ , and it will be studied extensively in later chapters.
(a) Find $P(X \geq 1)$ and $P(X \geq 2)$ without summing infinite series.
(b) Suppose that the company only knows about people who have made at least one purchase on their site (a user sets up an account to make a purchase, but someone who has never made a purchase there doesn't appear in the customer database). If the company computes the number of purchases for everyone in their database, then these data are draws from the conditional distribution of the number of purchases, given that at least one purchase is made. Find the conditional PMF of X given $X \geq 1$. (This conditional distribution is called a truncated Poisson distribution.)

Named Distributions

15. Find the CDF of an r.v. $X \sim \text{DUnif}(1, 2, \dots, n)$.
16. Let $X \sim \text{DUnif}(C)$, and B be a nonempty subset of C . Find the conditional distribution of X , given that X is in B .
17. An airline overbooks a flight, selling more tickets for the flight than there are seats on the plane (figuring that it's likely that some people won't show up). The plane has 100

- seats, and 110 people have booked the flight. Each person will show up for the flight with probability 0.9, independently. Find the probability that there will be enough seats for everyone who shows up for the flight.
18. (S) (a) In the World Series of baseball, two teams (call them A and B) play a sequence of games against each other, and the first team to win four games wins the series. Let p be the probability that A wins an individual game, and assume that the games are independent. What is the probability that team A wins the series?
(b) Give a clear intuitive explanation of whether the answer to (a) depends on whether the teams always play 7 games (and whoever wins the majority wins the series), or the teams stop playing more games as soon as one team has won 4 games (as is actually the case in practice: once the match is decided, the two teams do not keep playing more games).
19. In a chess tournament, n games are being played, independently. Each game ends in a win for one player with probability 0.4 and ends in a draw (tie) with probability 0.6. Find the PMFs of the number of games ending in a draw, and of the number of players whose games end in draws.
20. Suppose that a lottery ticket has probability p of being a winning ticket, independently of other tickets. A gambler buys 3 tickets, hoping this will triple the chance of having at least one winning ticket.
(a) What is the distribution of how many of the 3 tickets are winning tickets?
(b) Show that the probability that at least 1 of the 3 tickets is winning is $3p - 3p^2 + p^3$, in two different ways: by using inclusion-exclusion, and by taking the complement of the desired event and then using the PMF of a certain named distribution.
(c) Show that the gambler's chances of having at least one winning ticket do not quite triple (compared with buying only one ticket), but that they do approximately triple if p is small.
21. (S) Let $X \sim \text{Bin}(n, p)$ and $Y \sim \text{Bin}(m, p)$, independent of X . Show that $X - Y$ is not Binomial.
22. There are two coins, one with probability p_1 of Heads and the other with probability p_2 of Heads. One of the coins is randomly chosen (with equal probabilities for the two coins). It is then flipped $n \geq 2$ times. Let X be the number of times it lands Heads.
(a) Find the PMF of X .
(b) What is the distribution of X if $p_1 = p_2$?
(c) Give an intuitive explanation of why X is not Binomial for $p_1 \neq p_2$ (its distribution is called a mixture of two Binomials). You can assume that n is large for your explanation, so that the frequentist interpretation of probability can be applied.
23. There are n people eligible to vote in a certain election. Voting requires registration. Decisions are made independently. Each of the n people will register with probability p_1 . Given that a person registers, they will vote with probability p_2 . Given that a person votes, they will vote for Kodos (who is one of the candidates) with probability p_3 . What is the distribution of the number of votes for Kodos (give the PMF, fully simplified, or the name of the distribution, including its parameters)?
24. Let X be the number of Heads in 10 fair coin tosses.
(a) Find the conditional PMF of X , given that the first two tosses both land Heads.
(b) Find the conditional PMF of X , given that at least two tosses land Heads.

25. (S) Alice flips a fair coin n times and Bob flips another fair coin $n + 1$ times, resulting in independent $X \sim \text{Bin}(n, \frac{1}{2})$ and $Y \sim \text{Bin}(n + 1, \frac{1}{2})$.
(a) Show that $P(X < Y) = P(n - X < n + 1 - Y)$.
(b) Compute $P(X < Y)$.

Hint: Use (a) and the fact that X and Y are integer-valued.

26. If $X \sim \text{HGeom}(w, b, n)$, what is the distribution of $n - X$? Give a short proof.
27. Recall de Montmort's matching problem from Chapter 1: in a deck of n cards labeled 1 through n , a match occurs when the number on the card matches the card's position in the deck. Let X be the number of matching cards. Is X Binomial? Is X Hypergeometric?
28. (S) There are n eggs, each of which hatches a chick with probability p (independently). Each of these chicks survives with probability r , independently. What is the distribution of the number of chicks that hatch? What is the distribution of the number of chicks that survive? (Give the PMFs; also give the names of the distributions and their parameters, if applicable.)
29. (S) A sequence of n independent experiments is performed. Each experiment is a success with probability p and a failure with probability $q = 1 - p$. Show that conditional on the number of successes, all valid possibilities for the list of outcomes of the experiment are equally likely.
30. A certain company has $n + m$ employees, consisting of n women and m men. The company is deciding which employees to promote.
(a) Suppose for this part that the company decides to promote t employees, where $1 \leq t \leq n + m$, by choosing t random employees (with equal probabilities for each set of t employees). What is the distribution of the number of women who get promoted?
(b) Now suppose that instead of having a predetermined number of promotions to give, the company decides independently for each employee, promoting the employee with probability p . Find the distributions of the number of women who are promoted, the number of women who are not promoted, and the number of employees who are promoted.
(c) In the set-up from (b), find the conditional distribution of the number of women who are promoted, given that exactly t employees are promoted.
31. Once upon a time, a famous statistician offered tea to a lady. The lady claimed that she could tell whether milk had been added to the cup before or after the tea. The statistician decided to run some experiments to test her claim.
(a) The lady is given 6 cups of tea, where it is known in advance that 3 will be milk-first and 3 will be tea-first, in a completely random order. The lady gets to taste each and then guess which 3 were milk-first. Assume for this part that she has no ability whatsoever to distinguish milk-first from tea-first cups of tea. Find the probability that at least 2 of her 3 guesses are correct.
(b) Now the lady is given one cup of tea, with probability $1/2$ of it being milk-first. She needs to say whether she thinks it was milk-first. Let p_1 be the lady's probability of being correct given that it was milk-first, and p_2 be her probability of being correct given that it was tea-first. She claims that the cup was milk-first. Find the posterior odds that the cup is milk-first, given this information.
32. In Evan's history class, 10 out of 100 key terms will be randomly selected to appear on the final exam; Evan must then choose 7 of those 10 to define. Since he knows the format of the exam in advance, Evan is trying to decide how many key terms he should study.
(a) Suppose that Evan decides to study s key terms, where s is an integer between 0 and 100. Let X be the number of key terms appearing on the exam that he has studied. What is the distribution of X ? Give the name and parameters, in terms of s .
(b) Using R or other software, calculate the probability that Evan knows at least 7 of the 10 key terms that appear on the exam, assuming that he studies $s = 75$ key terms.
33. A book has n typos. Two proofreaders, Prue and Frida, independently read the book. Prue

catches each typo with probability p_1 and misses it with probability $q_1 = 1 - p_1$, independently, and likewise for Frida, who has probabilities p_2 of catching and $q_2 = 1 - p_2$ of missing each typo. Let X_1 be the number of typos caught by Prue, X_2 be the number caught by Frida, and X be the number caught by at least one of the two proofreaders.

(a) Find the distribution of X .

(b) For this part only, assume that $p_1 = p_2$. Find the conditional distribution of X_1 given that $X_1 + X_2 = t$.

34. There are n students at a certain school, of whom $X \sim \text{Bin}(n, p)$ are Statistics majors. A simple random sample of size m is drawn ("simple random sample" means sampling without replacement, with all subsets of the given size equally likely).

(a) Find the PMF of the number of Statistics majors in the sample, using the law of total probability (don't forget to say what the support is). You can leave your answer as a sum (though with some algebra it can be simplified, by writing the binomial coefficients in terms of factorials and using the binomial theorem).

(b) Give a story proof derivation of the distribution of the number of Statistics majors in the sample; simplify fully.

Hint: Does it matter whether the students declare their majors before or after the random sample is drawn?

35. (S) Players A and B take turns in answering trivia questions, starting with player A answering the first question. Each time A answers a question, she has probability p_1 of getting it right. Each time B plays, he has probability p_2 of getting it right.

(a) If A answers m questions, what is the PMF of the number of questions she gets right?

(b) If A answers m times and B answers n times, what is the PMF of the total number of questions they get right (you can leave your answer as a sum)? Describe exactly when/whether this is a Binomial distribution.

(c) Suppose that the first player to answer correctly wins the game (with no predetermined maximum number of questions that can be asked). Find the probability that A wins the game.

36. There are n voters in an upcoming election in a certain country, where n is a large, even number. There are two candidates: Candidate A (from the Unite Party) and Candidate B (from the Untie Party). Let X be the number of people who vote for Candidate A. Suppose that each voter chooses randomly whom to vote for, independently and with equal probabilities.

(a) Find an exact expression for the probability of a tie in the election (so the candidates end up with the same number of votes).

(b) Use Stirling's approximation, which approximates the factorial function as

$$n! \approx \sqrt{2\pi n} \left(\frac{n}{e}\right)^n$$

to find a simple approximation to the probability of a tie. Your answer should be of the form $1/\sqrt{cn}$, with c a constant (which you should specify).

37. (S) A message is sent over a noisy channel. The message is a sequence $x_1, x_2, \dots, x_n$ of n bits ($x_i \in \{0, 1\}$). Since the channel is noisy, there is a chance that any bit might be corrupted, resulting in an error (a 0 becomes a 1 or vice versa). Assume that the error events are independent. Let p be the probability that an individual bit has an error ($0 < p < 1/2$). Let $y_1, y_2, \dots, y_n$ be the received message (so $y_i = x_i$ if there is no error in that bit, but $y_i = 1 - x_i$ if there is an error there).

To help detect errors, the n th bit is reserved for a parity check: x_n is defined to be 0 if $x_1 + x_2 + \dots + x_{n-1}$ is even, and 1 if $x_1 + x_2 + \dots + x_{n-1}$ is odd. When the message is received, the recipient checks whether y_n has the same parity as $y_1 + y_2 + \dots + y_{n-1}$. If the parity is wrong, the recipient knows that at least one error occurred; otherwise, the recipient assumes that there were no errors.

(a) For $n = 5, p = 0.1$, what is the probability that the received message has errors which go

undetected?

(b) For general n and p , write down an expression (as a sum) for the probability that the received message has errors which go undetected.

(c) Give a simplified expression, not involving a sum of a large number of terms, for the probability that the received message has errors which go undetected.

Hint for (c): Letting

$$a = \sum_{k \text{ even}, k \geq 0} \binom{n}{k} p^k (1-p)^{n-k} \text{ and } b = \sum_{k \text{ odd}, k \geq 1} \binom{n}{k} p^k (1-p)^{n-k},$$

the binomial theorem makes it possible to find simple expressions for $a + b$ and $a - b$, which then makes it possible to obtain a and b .

Independence of r.v.s

38. (a) Give an example of dependent r.v.s X and Y such that $P(X < Y) = 1$.
(b) Give an example of independent r.v.s X and Y such that $P(X < Y) = 1$.
39. Give an example of two discrete random variables X and Y on the same sample space such that X and Y have the same distribution, with support $\{1, 2, \dots, 10\}$, but the event $X = Y$ never occurs. If X and Y are independent, is it still possible to construct such an example?
40. Suppose X and Y are discrete r.v.s such that $P(X = Y) = 1$. This means that X and Y always take on the same value.
(a) Do X and Y have the same PMF?
(b) Is it possible for X and Y to be independent?
41. If X, Y, Z are r.v.s such that X and Y are independent and Y and Z are independent, does it follow that X and Z are independent?

Hint: Think about simple and extreme examples.

42. (S) Let X be a random day of the week, coded so that Monday is 1, Tuesday is 2, etc. (so X takes values 1, 2, ..., 7, with equal probabilities). Let Y be the next day after X (again represented as an integer between 1 and 7). Do X and Y have the same distribution? What is $P(X < Y)$?
43. (a) Is it possible to have two r.v.s X and Y such that X and Y have the same distribution but $P(X < Y) \geq p$, where:

- $p = 0.9$?
- $p = 0.99$?
- $p = 0.9999999999999999$?
- $p = 1$?

For each, give an example showing it is possible, or prove it is impossible.

Hint: Do the previous question first.

(b) Consider the same question as in Part (a), but now assume that X and Y are independent. Do your answers change?

44. For x and y binary digits (0 or 1), let $x \oplus y$ be 0 if $x = y$ and 1 if $x \neq y$ (this operation is called exclusive or (often abbreviated to XOR), or addition mod 2).

(a) Let $X \sim \text{Bern}(p)$ and $Y \sim \text{Bern}(1/2)$, independently. What is the distribution of $X \oplus Y$?

- (b) With notation as in (a), is $X \oplus Y$ independent of X ? Is $X \oplus Y$ independent of Y ? Be sure to consider both the case $p = 1/2$ and the case $p \neq 1/2$.
- (c) Let $X_1, \dots, X_n$ be i.i.d. $\text{Bern}(1/2)$. For each nonempty subset J of $\{1, 2, \dots, n\}$, let

$$Y_J = \bigoplus_{j \in J} X_j$$

where the notation means to "add" in the $\oplus$ sense all the elements of J ; the order in which this is done doesn't matter since $x \oplus y = y \oplus x$ and $(x \oplus y) \oplus z = x \oplus (y \oplus z)$. Show that $Y_J \sim \text{Bern}(1/2)$ and that these $2^n - 1$ r.v.s are pairwise independent, but not independent. For example, we can use this to simulate 1023 pairwise independent fair coin tosses using only 10 independent fair coin tosses.

Hint: Apply the previous parts with $p = 1/2$. Show that if J and K are two different nonempty subsets of $\{1, 2, \dots, n\}$, then we can write $Y_J = A \oplus B$, $Y_K = A \oplus C$, where A consists of the X_i with $i \in J \cap K$, B consists of the X_i with $i \in J \cap K^c$, and C consists of the X_i with $i \in J^c \cap K$. Then A, B, C are independent since they are based on disjoint sets of X_i . Also, at most one of these sets of X_i can be empty. If $J \cap K = \emptyset$, then $Y_J = B$, $Y_K = C$. Otherwise, compute $P(Y_J = y, Y_K = z)$ by conditioning on whether $A = 1$.

Mixed Practice

45. (S) A new treatment for a disease is being tested, to see whether it is better than the standard treatment. The existing treatment is effective on 50% of patients. It is believed initially that there is a $2/3$ chance that the new treatment is effective on 60% of patients, and a $1/3$ chance that the new treatment is effective on 50% of patients. In a pilot study, the new treatment is given to 20 random patients, and is effective for 15 of them.
- (a) Given this information, what is the probability that the new treatment is better than the standard treatment?
- (b) A second study is done later, giving the new treatment to 20 new random patients. Given the results of the first study, what is the PMF for how many of the new patients the new treatment is effective on? (Letting p be the answer to (a), your answer can be left in terms of p .)
46. Independent Bernoulli trials are performed, with success probability $1/2$ for each trial. An important question that often comes up in such settings is how many trials to perform. Many controversies have arisen in statistics over the issue of how to analyze data coming from an experiment where the number of trials can depend on the data collected so far.

For example, if we can follow the rule "keep performing trials until there are more than twice as many failures as successes, and then stop", then naively looking at the ratio of failures to successes (if and when the process stops) will give more than $2 : 1$ rather than the true theoretical $1 : 1$ ratio; this could be a very misleading result! However, it might never happen that there are more than twice as many failures as successes; in this problem, you will find the probability of that happening.

- (a) Two gamblers, A and B, make a series of bets, where each has probability $1/2$ of winning a bet, but A gets \$2 for each win and loses \$1 for each loss (a very favorable game for A!). Assume that the gamblers are allowed to borrow money, so they can and do gamble forever. Let p_k be the probability that A, starting with \$ k , will ever reach \$0, for each $k \geq 0$. Explain how this story relates to the original problem, and how the original problem can be solved if we can find p_k .
- (b) Find p_k .

Hint: As in the gambler's ruin, set up and solve a difference equation for p_k . We have $p_k \rightarrow 0$ as $k \rightarrow \infty$ (you don't need to prove this, but it should make sense since the game is so favorable to A, which will result in A's fortune going to ∞ ; a formal proof, not required here, could be done using the law of large numbers, an important theorem from Chapter 10). The solution can be written neatly in terms of the golden ratio.

(c) Find the probability of ever having more than twice as many failures as successes with independent $\text{Bern}(1/2)$ trials, as originally desired.

47. A copy machine is used to make n pages of copies per day. The machine has two trays in which paper gets loaded, and each page used is taken randomly and independently from one of the trays. At the beginning of the day, the trays are refilled so that they each have m pages.

(a) Let $\text{pbinom}(x, n, p)$ be the CDF of the $\text{Bin}(n, p)$ distribution, evaluated at x . In terms of pbinom , find a simple expression for the probability that both trays have enough paper on any particular day, when this probability is strictly between 0 and 1 (also specify the values of m for which the probability is 0 and the values for which it is 1).

Hint: Be careful about whether inequalities are strict, since the Binomial is discrete.

(b) Using a computer, find the smallest value of m for which there is at least a 95% chance that both trays have enough paper on a particular day, for $n = 10$, $n = 100$, $n = 1000$, and $n = 10000$.

Hint: If you use R, you may find the following commands useful:

`g <- function(m, n) [your answer from (a)]` defines a function g such that $g(m, n)$ is your answer from (a), `g(1 : 100, 100)` gives the vector $(g(1, 100), \dots, g(100, 100))$, which `(v > 0.95)` gives the indices of the components of vector $\mathbf{v}$ that exceed 0.95, and `min(w)` gives the minimum of a vector $\mathbf{w}$.

4.4 4 Expectation

4.4.1 4.1 Definition of Expectation

In the previous chapter, we introduced the distribution of a random variable, which gives us full information about the probability that the r.v. will fall into any particular set. For example, we can say how likely it is that the r.v. will exceed 1000, that it will equal 5, or that it will be in the interval $[0, 7]$. It can be unwieldy to manage so many probabilities though, so often we want just one number summarizing the "average" value of the r.v.

There are several senses in which the word "average" is used, but by far the most commonly used is the mean of an r.v., also known as its expected value. In addition, much of statistics is about understanding variability in the world, so it is often important to know how "spread out" the distribution is; we will formalize this with the concepts of variance and standard deviation. As we'll see, variance and standard deviation are defined in terms of expected values, so the uses of expected values go far beyond just computing averages.

Given a list of numbers $x_1, x_2, \dots, x_n$, the familiar way to average them is to add them up and divide by n . This is called the arithmetic mean, and is defined by

$$\bar{x} = \frac{1}{n} \sum_{j=1}^n x_j.$$

More generally, we can define a weighted mean of $x_1, \dots, x_n$ as

$$\text{weighted-mean}(x) = \sum_{j=1}^n x_j p_j,$$

where the weights $p_1, \dots, p_n$ are pre-specified nonnegative numbers that add up to 1 (so the unweighted mean $\bar{x}$ is obtained when $p_j = 1/n$ for all j).

The definition of expectation for a discrete r.v. is inspired by the weighted mean of a list of numbers, with weights given by probabilities.

Definition 4.1.1 (Expectation of a discrete r.v.). The expected value (also called the expectation or mean) of a discrete r.v. X whose distinct possible values are $x_1, x_2, \dots$ is defined by

$$E(X) = \sum_{j=1}^{\infty} x_j P(X = x_j).$$

If the support is finite, then this is replaced by a finite sum. We can also write

$$E(X) = \sum_x \underbrace{x}_{\text{value}} \underbrace{P(X = x)}_{\text{PMF at } x},$$

where the sum is over the support of X (in any case, $xP(X = x)$ is 0 for any x not in the support). The expectation is undefined if $\sum_{j=1}^{\infty} |x_j| P(X = x_j)$ diverges, since then the series for $E(X)$ diverges or its value depends on the order in which the x_j are listed.

In words, the expected value of X is a weighted average of the possible values that X can take on, weighted by their probabilities. Let's check that the definition makes sense in a few simple examples:

1. Let X be the result of rolling a fair 6-sided die, so X takes on the values 1, 2, 3, 4, 5, 6, with equal probabilities. Intuitively, we should be able to get the average by adding up these values and dividing by 6. Using the definition, the expected value is

$$E(X) = \frac{1}{6}(1 + 2 + \dots + 6) = 3.5$$

as we expected. Note that X never equals its mean in this example. This is similar to the fact that the average number of children per household in some country could be 1.8, but that doesn't mean that a typical household has 1.8 children!

2. Let $X \sim \text{Bern}(p)$ and $q = 1 - p$. Then

$$E(X) = 1p + 0q = p$$

which makes sense intuitively since it is between the two possible values of X , compromising between 0 and 1 based on how likely each is. This is illustrated in Figure 4.1 for a case with $p < 1/2$: two pebbles are being balanced on a seesaw. For the seesaw to balance, the fulcrum (shown as a triangle) must be at p , which in physics terms is the center of mass.

The frequentist interpretation would be to consider a large number of independent Bernoulli trials, each with probability p of success. Writing 1 for "success" and 0 for "failure", in the long run we would expect to have data consisting of a list of numbers where the proportion of 1's is very close to p . The average of a list of 0's and 1's is the proportion of 1's.

3. Let X have 3 distinct possible values, a_1, a_2, a_3 , with probabilities p_1, p_2, p_3 , respectively. Imagine running a simulation where n independent draws




FIGURE 4.1 Center of mass of two pebbles, depicting that $E(X) = p$ for $X \sim \text{Bern}(p)$. Here q and p denote the masses of the two pebbles.

from the distribution of X are generated. For n large, we would expect to have about $p_1 n a_1$'s, $p_2 n a_2$'s, and $p_3 n a_3$'s. (We will look at a more mathematical version of this example when we study the law of large numbers in Chapter 10.) If the simulation results are close to these expected results, then the arithmetic mean of the simulation results is approximately

$$\frac{p_1 n \cdot a_1 + p_2 n \cdot a_2 + p_3 n \cdot a_3}{n} = p_1 a_1 + p_2 a_2 + p_3 a_3 = E(X).$$

Note that $E(X)$ depends only on the distribution of X . This follows directly from the definition, but is worth recording since it is fundamental.

Proposition 4.1.2. If X and Y are discrete r.v.s with the same distribution, then $E(X) = E(Y)$ (if either side exists).

Proof. In the definition of $E(X)$, we only need to know the PMF of X .

The converse of the above proposition is false since the expected value is just a onenumber summary, not nearly enough to specify the entire distribution; it's a measure of where the "center" is but does not determine, for example, how spread out the distribution is or how likely the r.v. is to be positive. Figure 4.2 shows an example of two different PMFs with the same expected value (balancing point).

4.1.3 (Replacing an r.v. by its expectation). For any discrete r.v. X , the expected value $E(X)$ is a number (if it exists). A common mistake is to replace an r.v. by its expectation without justification, which is wrong both mathematically (X is a function, $E(X)$ is a constant) and statistically (it ignores the variability of X), except in the degenerate case where X is a constant.

Notation 4.1.4. We often abbreviate $E(X)$ to EX . Similarly, we often abbreviate $E(X^2)$ to EX^2 , and $E(X^n)$ to EX^n .

4.1.5. Paying attention to the order of operations is crucial when working with expectation. As stated above, EX^2 is the expectation of the random variable X^2 , not the square of the number EX . Unless the parentheses explicitly indicate otherwise,




FIGURE 4.2

The expected value does not determine the distribution: different PMFs can have the same balancing point.

for the expectation of an r.v. raised to a power, first we take the power and then we take the expectation. For example, $E(X - 1)^4$ is $E((X - 1)^4)$, not $(E(X - 1))^4$.

4.4.2 4.2 Linearity of Expectation

The most important property of expectation is linearity: the expected value of a sum of r.v.s is the sum of the individual expected values.

Theorem 4.2.1 (Linearity of expectation). For any r.v.s X, Y and any constant c ,

$$\begin{aligned} E(X + Y) &= E(X) + E(Y), \\ E(cX) &= cE(X). \end{aligned}$$

The second equation says that we can take out constant factors from an expectation; this is both intuitively reasonable and easily verified from the definition. The first equation, $E(X + Y) = E(X) + E(Y)$, also seems reasonable when X and Y are independent. What may be surprising is that it holds even if X and Y are dependent! To build intuition for this, consider the extreme case where X always equals Y . Then $X + Y = 2X$, and both sides of $E(X + Y) = E(X) + E(Y)$ are equal to $2E(X)$, so linearity still holds even in the most extreme case of dependence.

Linearity is true for all r.v.s, not just discrete r.v.s, but in this chapter we will prove it only for discrete r.v.s. Before proving linearity, it is worthwhile to recall some basic facts about averages. If we have a list of numbers, say $(1, 1, 1, 1, 1, 3, 3, 5)$, we can calculate their mean by adding all the values and dividing by the length of the list, so that each element of the list gets a weight of $\frac{1}{8}$:

$$\frac{1}{8}(1 + 1 + 1 + 1 + 1 + 3 + 3 + 5) = 2.$$

But another way to calculate the mean is to group together all the 1's, all the 3's, and all the 5's, and then take a weighted average, giving appropriate weights to 1's, 3's, and 5's:

$$\frac{5}{8} \cdot 1 + \frac{2}{8} \cdot 3 + \frac{1}{8} \cdot 5 = 2$$

This insight—that averages can be calculated in two ways, ungrouped or grouped—is all that is needed to prove linearity! Recall that X is a function which assigns a real number to every outcome s in the sample space. The r.v. X may assign the same value to multiple sample outcomes. When this happens, our definition of expectation groups all these outcomes together into a super-pebble whose weight, $P(X = x)$, is the total weight of the constituent pebbles. This grouping process is illustrated in Figure 4.3 for a hypothetical r.v. taking values in $\{0, 1, 2\}$. So our definition of expectation corresponds to the grouped way of taking averages.




FIGURE 4.3 Left: X assigns a number to each pebble in the sample space. Right: Grouping the pebbles by the value that X assigns to them, the 9 pebbles become 3 super-pebbles. The weight of a super-pebble is the sum of the weights of the constituent pebbles.

The advantage of this definition is that it allows us to work with the distribution of X directly, without returning to the sample space. The disadvantage comes when we have to prove theorems like this one, for if we have another r.v. Y on the same sample space, the super-pebbles created by Y are different from those created from X , with different weights $P(Y = y)$; this makes it difficult to combine $\sum_x xP(X = x)$ and $\sum_y yP(Y = y)$.

Fortunately, we know there's another equally valid way to calculate an average: we can take a weighted average of the values of individual pebbles. In other words, if $X(s)$ is the value that X assigns to pebble s , we can take the weighted average

$$E(X) = \sum_s X(s)P(\{s\}),$$

where $P(\{s\})$ is the weight of pebble s . This corresponds to the ungrouped way of taking averages. The advantage of this definition is that it breaks down the sample space into the smallest possible units, so we are now using the same weights $P(\{s\})$ for every random variable defined on this sample space. If Y is another random variable, then

$$E(Y) = \sum_s Y(s)P(\{s\}).$$

We can combine $\sum_s X(s)P(\{s\})$ and $\sum_s Y(s)P(\{s\})$, which gives $E(X) + E(Y) = \sum_s X(s)P(\{s\}) + \sum_s Y(s)P(\{s\}) = \sum_s (X + Y)(s)P(\{s\}) = E(X + Y)$.

Another intuition for linearity of expectation is via the concept of simulation. If we simulate many, many times from the distribution of X , the histogram of the simulated values will look very much like the true PMF of X . In particular, the arithmetic mean of the simulated values will be very close to the true value of $E(X)$ (the precise nature of this convergence is described by the law of large numbers, an important theorem that we will discuss in detail in Chapter 10).

Let X and Y be r.v.s summarizing a certain experiment. Suppose we perform the experiment n times, where n is a very large number, and we write down the values realized by X and Y each time. For each repetition of the experiment, we obtain an X value, a Y value, and (by adding them) an $X + Y$ value. In Figure 4.4, each row represents a repetition of the experiment. The left column contains the draws of X , the middle column contains the draws of Y , and the right column contains the draws of $X + Y$.

There are two ways to calculate the sum of all the numbers in the last column. The straightforward way is just to add all the numbers in that column. But an equally valid way is to add all the numbers in the first column, add all the numbers in the second column, and then add the two column sums.

Dividing by n everywhere, what we've argued is that the following procedures are equivalent:

- Taking the arithmetic mean of all the numbers in the last column. By the law of large numbers, this is very close to $E(X + Y)$.
- Taking the arithmetic mean of the first column and the arithmetic mean of the second column, then adding the two column means. By the law of large numbers, this is very close to $E(X) + E(Y)$.

Linearity of expectation thus emerges as a simple fact about arithmetic (we're just adding numbers in two different orders)! Notice that nowhere in our argument did we rely on whether X and Y were independent. In fact, in Figure 4.4, X and Y appear to be dependent: Y tends to be large when X is large, and Y tends to be small when X is small (in the language of Chapter 7, we say that X and Y are positively correlated). But this dependence is irrelevant: shuffling the draws of Y could completely alter the pattern of dependence between X and Y , but would have no effect on the column sums.

X	Y	$X + Y$
3	4	7
2	2	4
6	8	14
10	23	33
1	-3	-2
1	0	1
5	9	14
4	1	5
$\vdots$	$\vdots$	$\vdots$

FIGURE 4.4

Intuitive view of linearity of expectation. Each row represents a repetition of the experiment; the three columns are the realized values of X , Y , and $X + Y$, respectively. Adding all the numbers in the last column is equivalent to summing the first column and the second column separately, then adding the two column sums. So the mean of the last column is the sum of the first and second column means; this is linearity of expectation.

$$\frac{1}{n} \sum_{i=1}^n x_i + \frac{1}{n} \sum_{i=1}^n y_i = \frac{1}{n} \sum_{i=1}^n (x_i + y_i)$$

$$E(X) + E(Y) = E(X + Y)$$

Linearity is an extremely handy tool for calculating expected values, often allowing us to bypass the definition of expected value altogether. Let's use linearity to find the expectations of the Binomial and Hypergeometric distributions.

Example 4.2.2 (Binomial expectation). For $X \sim \text{Bin}(n, p)$, let's find $E(X)$ in two ways. By definition of expectation,

$$E(X) = \sum_{k=0}^n k P(X = k) = \sum_{k=0}^n k \binom{n}{k} p^k q^{n-k}.$$

From Example 1.5.2, we know $k \binom{n}{k} = n \binom{n-1}{k-1}$, so

$$\begin{aligned} \sum_{k=0}^n k \binom{n}{k} p^k q^{n-k} &= n \sum_{k=0}^n \binom{n-1}{k-1} p^k q^{n-k} \\ &= np \sum_{k=1}^n \binom{n-1}{k-1} p^{k-1} q^{n-k} \\ &= np \sum_{j=0}^{n-1} \binom{n-1}{j} p^j q^{n-1-j} \\ &= np. \end{aligned}$$

The sum in the penultimate line equals 1 because it is the sum of the $\text{Bin}(n-1, p)$ PMF (or by the binomial theorem). Therefore, $E(X) = np$.

This proof required us to remember combinatorial identities and manipulate binomial coefficients. Using linearity of expectation, we obtain a much shorter path to the same result. Let's write X as the sum of n independent $\text{Bern}(p)$ r.v.s:

$$X = I_1 + \cdots + I_n,$$

where each I_j has expectation $E(I_j) = 1p + 0q = p$. By linearity,

$$E(X) = E(I_1) + \cdots + E(I_n) = np.$$

Example 4.2.3 (Hypergeometric expectation). Let $X \sim \text{HGeom}(w, b, n)$, interpreted as the number of white balls in a sample of size n drawn without replacement from an urn with w white and b black balls. As in the Binomial case, we can write X as a sum of Bernoulli random variables,

$$X = I_1 + \cdots + I_n,$$

where I_j equals 1 if the j th ball in the sample is white and 0 otherwise. By symmetry, $I_j \sim \text{Bern}(p)$ with $p = w/(w + b)$, since unconditionally the j th ball drawn is equally likely to be any of the balls.

Unlike in the Binomial case, the I_j are not independent, since the sampling is without replacement: given that a ball in the sample is white, there is a lower chance that another ball in the sample is white. However, linearity still holds for dependent random variables! Thus,

$$E(X) = nw/(w + b).$$

As another example of the power of linearity, we can give a quick proof of the intuitive idea that "bigger r.v.s have bigger expectations".

Proposition 4.2.4 (Monotonicity of expectation). Let X and Y be r.v.s such that $X \geq Y$ with probability 1. Then $E(X) \geq E(Y)$, with equality holding if and only if $X = Y$ with probability 1.

Proof. This result holds for all r.v.s, but we will prove it only for discrete r.v.s since this chapter focuses on discrete r.v.s. The r.v. $Z = X - Y$ is nonnegative (with probability 1), so $E(Z) \geq 0$ since $E(Z)$ is defined as a sum of nonnegative terms. By linearity,

$$E(X) - E(Y) = E(X - Y) \geq 0$$

as desired. If $E(X) = E(Y)$, then by linearity we also have $E(Z) = 0$, which implies that $P(X = Y) = P(Z = 0) = 1$ since if even one term in the sum defining $E(Z)$ is positive, then the whole sum is positive.

4.4.3 4.3 Geometric and Negative Binomial

We now introduce two more famous discrete distributions, the Geometric and Negative Binomial, and calculate their expected values.

Story 4.3.1 (Geometric distribution). Consider a sequence of independent Bernoulli trials, each with the same success probability $p \in (0, 1)$, with trials performed until a success occurs. Let X be the number of failures before the first successful trial. Then X has the Geometric distribution with parameter p ; we denote this by $X \sim \text{Geom}(p)$.

For example, if we flip a fair coin until it lands Heads for the first time, then the number of Tails before the first occurrence of Heads is distributed as $\text{Geom}(1/2)$.

To get the Geometric PMF from the story, imagine the Bernoulli trials as a string of 0's (failures) ending in a single 1 (success). Each 0 has probability $q = 1 - p$ and the final 1 has probability p , so a string of k failures followed by one success has probability $q^k p$.

Theorem 4.3.2 (Geometric PMF). If $X \sim \text{Geom}(p)$, then the PMF of X is

$$P(X = k) = q^k p$$

for $k = 0, 1, 2, \dots$, where $q = 1 - p$.

This is a valid PMF because, summing a geometric series (see the math appendix for a review of geometric series), we have

$$\sum_{k=0}^{\infty} q^k p = p \sum_{k=0}^{\infty} q^k = p \cdot \frac{1}{1 - q} = 1$$

Just as the binomial theorem shows that the Binomial PMF is valid, a geometric series shows that the Geometric PMF is valid! A geometric series can also be used to obtain the Geometric CDF.

Theorem 4.3.3 (Geometric CDF). If $X \sim \text{Geom}(p)$, then the CDF of X is

$$F(x) = \begin{cases} 1 - q^{\lfloor x \rfloor + 1}, & \text{if } x \geq 0 \\ 0, & \text{if } x < 0 \end{cases}$$

where $q = 1 - p$ and $\lfloor x \rfloor$ is the greatest integer less than or equal to x .

Proof. Let F be the CDF of X . We will find $F(x)$ first for the case $x < 0$, then for the case that x is a nonnegative integer, and lastly for the case that x is a nonnegative real number. For $x < 0$, $F(x) = 0$ since X can't be negative. For n a nonnegative integer,

$$F(n) = \sum_{k=0}^n P(X = k) = p \sum_{k=0}^n q^k = p \cdot \frac{1 - q^{n+1}}{1 - q} = 1 - q^{n+1}.$$

We can also get the same result from the fact that the event $X \geq n + 1$ means that the first $n + 1$ trials were failures:

$$F(n) = 1 - P(X > n) = 1 - P(X \geq n + 1) = 1 - q^{n+1}$$

For real $x \geq 0$,

$$F(x) = P(X \leq x) = P(X \leq \lfloor x \rfloor),$$

since X always takes on integer values. For example,

$$P(X \leq 3.7) = P(X \leq 3) + P(3 < X \leq 3.7) = P(X \leq 3).$$

Therefore, F is as claimed.

Figure 4.3 displays the $\text{Geom}(0.5)$ PMF and CDF from 0 to 6. All Geometric PMFs have a similar shape; the greater the success probability p , the more quickly the PMF decays to 0.

4.3.4 (Conventions for the Geometric). There are differing conventions for the definition of the Geometric distribution; some sources define the Geometric as the total number of trials, including the success. In this book, the Geometric distribution excludes the success, and the First Success distribution includes the success.




FIGURE 4.5
Geom(0.5) PMF and CDF.

Definition 4.3.5 (First Success distribution). In a sequence of independent Bernoulli trials with success probability p , let Y be the number of trials until the first successful trial, including the success. Then Y has the First Success distribution with parameter p ; we denote this by $Y \sim \text{FS}(p)$.

It is easy to convert back and forth between the two but important to be careful about which convention is being used. If $Y \sim \text{FS}(p)$ then $Y - 1 \sim \text{Geom}(p)$, and we can convert between the PMFs of Y and $Y - 1$ by writing

$$P(Y = k) = P(Y - 1 = k - 1).$$

Conversely, if $X \sim \text{Geom}(p)$, then $X + 1 \sim \text{FS}(p)$.

Example 4.3.6 (Geometric expectation). Let $X \sim \text{Geom}(p)$. By definition,

$$E(X) = \sum_{k=0}^{\infty} kq^k p,$$

where $q = 1 - p$. This sum looks unpleasant; it's not a geometric series because of the extra k multiplying each term. But we notice that each term looks similar to kq^{k-1} , the derivative of q^k (with respect to q), so let's start there:

$$\sum_{k=0}^{\infty} q^k = \frac{1}{1 - q}$$

This geometric series converges since $0 < q < 1$. Differentiating both sides with respect to q , we get

$$\sum_{k=0}^{\infty} kq^{k-1} = \frac{1}{(1 - q)^2}$$

Finally, if we multiply both sides by pq , we recover the original sum we wanted to find:

$$E(X) = \sum_{k=0}^{\infty} kq^k p = pq \sum_{k=0}^{\infty} kq^{k-1} = pq \frac{1}{(1 - q)^2} = \frac{q}{p}.$$

In Example 9.1.8, we will give a story proof of the same result, based on first-step analysis: condition on the result of the first trial in the story interpretation of X . If the first trial is a success, we know $X = 0$ and if it's a failure, we have one wasted trial and then are back where we started. $\square$

Example 4.3.7 (First Success expectation). Since we can write $Y \sim \text{FS}(p)$ as $Y = X + 1$ where $X \sim \text{Geom}(p)$, we have

$$E(Y) = E(X + 1) = \frac{q}{p} + 1 = \frac{1}{p}$$

□

The Negative Binomial distribution generalizes the Geometric distribution: instead of waiting for just one success, we can wait for any predetermined number r of successes.

Story 4.3.8 (Negative Binomial distribution). In a sequence of independent Bernoulli trials with success probability p , if X is the number of failures before the r th success, then X is said to have the Negative Binomial distribution with parameters r and p , denoted $X \sim \text{NBin}(r, p)$.

□

Both the Binomial and the Negative Binomial distributions are based on independent Bernoulli trials; they differ in the stopping rule and in what they are counting. The Binomial counts the number of successes in a fixed number of trials; the Negative Binomial counts the number of failures until a fixed number of successes.

In light of these similarities, it comes as no surprise that the derivation of the Negative Binomial PMF bears a resemblance to the corresponding derivation for the Binomial.

Theorem 4.3.9 (Negative Binomial PMF). If $X \sim \text{NBin}(r, p)$, then the PMF of X is

$$P(X = n) = \binom{n+r-1}{r-1} p^r q^n$$

for $n = 0, 1, 2, \dots$, where $q = 1 - p$.

Proof. Imagine a string of 0's and 1's, with 1's representing successes. The probability of any specific string of n 0's and r 1's is $p^r q^n$. How many such strings are there? Because we stop as soon as we hit the r th success, the string must terminate in a 1. Among the other $n + r - 1$ positions, we choose $r - 1$ places for the remaining 1's to go. So the overall probability of exactly n failures before the r th success is

$$P(X = n) = \binom{n+r-1}{r-1} p^r q^n, \quad n = 0, 1, 2, \dots$$

□

Just as a Binomial r.v. can be represented as a sum of i.i.d. Bernoullis, a Negative Binomial r.v. can be represented as a sum of i.i.d. Geometrics.

Theorem 4.3.10. Let $X \sim \text{NBin}(r, p)$, viewed as the number of failures before the r th success in a sequence of independent Bernoulli trials with success probability p . Then we can write $X = X_1 + \dots + X_r$ where the X_i are i.i.d. $\text{Geom}(p)$.

Proof. Let X_1 be the number of failures until the first success, X_2 be the number of failures between the first success and the second success, and in general, X_i be the number of failures between the $(i - 1)$ st success and the i th success.

Then $X_1 \sim \text{Geom}(p)$ by the story of the Geometric distribution. After the first success, the number of additional failures until the next success is still Geometric! So $X_2 \sim \text{Geom}(p)$, and similarly for all the X_i . Furthermore, the X_i are independent because the trials are all independent of each other. Adding the X_i , we get the total number of failures before the r th success, which is X .

Using linearity, the expectation of the Negative Binomial now follows without any additional calculations.

Example 4.3.11 (Negative Binomial expectation). Let $X \sim \text{NBin}(r, p)$. By the previous theorem, we can write $X = X_1 + \dots + X_r$, where the X_i are i.i.d. $\text{Geom}(p)$. By linearity,

$$E(X) = E(X_1) + \dots + E(X_r) = r \cdot \frac{q}{p}$$

The next example is a famous problem in probability and an instructive application of the Geometric and First Success distributions. It is usually stated as a problem about collecting coupons, hence its name, but we'll use toys instead of coupons.

Example 4.3.12 (Coupon collector). Suppose there are n types of toys, which you are collecting one by one, with the goal of getting a complete set. When collecting toys, the toy types are random (as is sometimes the case, for example, with toys included in cereal boxes or included with kids' meals from a fast food restaurant). Assume that each time you collect a toy, it is equally likely to be any of the n types. What is the expected number of toys needed until you have a complete set?

Solution:

Let N be the number of toys needed; we want to find $E(N)$. Our strategy will be to break up N into a sum of simpler r.v.s so that we can apply linearity. So write

$$N = N_1 + N_2 + \cdots + N_n$$

where N_1 is the number of toys until the first toy type you haven't seen before (which is always 1, as the first toy is always a new type), N_2 is the additional number of toys until the second toy type you haven't seen before, and so forth. Figure 4.6 illustrates these definitions with $n = 3$ toy types.




FIGURE 4.6 Coupon collector, $n = 3$. Here N_1 is the time (number of toys collected) until the first new toy type, N_2 is the additional time until the second new type, and N_3 is the additional time until the third new type. The total number of toys for a complete set is $N_1 + N_2 + N_3$.

By the story of the FS distribution, $N_2 \sim \text{FS}((n-1)/n)$: after collecting the first toy type, there's a $1/n$ chance of getting the same toy you already had (failure) and an $(n-1)/n$ chance you'll get something new (success). Similarly, N_3 , the additional number of toys until the third new toy type, is distributed $\text{FS}((n-2)/n)$. In general,

$$N_j \sim \text{FS}((n-j+1)/n).$$

By linearity,

$$\begin{aligned} E(N) &= E(N_1) + E(N_2) + E(N_3) + \cdots + E(N_n) \\ &= 1 + \frac{n}{n-1} + \frac{n}{n-2} + \cdots + n \\ &= n \sum_{j=1}^n \frac{1}{j}. \end{aligned}$$

For large n , this is very close to $n(\log n + 0.577)$.

Before we leave this example, let's take a moment to connect it to our proof of Theorem 4.3.10, the representation of the Negative Binomial as a sum of i.i.d. Geometrics. In both problems, we are waiting for a specified number of successes, and we approach the problem by considering the intervals between successes. There are two major differences:

- In Theorem 4.3.10, we exclude the successes themselves, so the number of failures between two successes is Geometric. In the coupon collector problem, we include the successes because we want to count the total number of toys, so we have First Success r.v.s instead.
- In Theorem 4.3.10, the probability of success in each trial never changes, so the total number of failures is a sum of i.i.d. Geometrics. In the coupon collector problem, the probability of success decreases after each success, since it becomes harder and harder to find a new toy type you haven't seen before; so the N_j are not identically distributed, though they are independent. $\square$

4.3.13 (Expectation of a nonlinear function of an r.v.). Expectation is linear, but in general we do not have $E(g(X)) = g(E(X))$ for arbitrary functions g . We must be careful not to move the E around when g is not linear. The next example shows a situation in which $E(g(X))$ is very different from $g(E(X))$.

Example 4.3.14 (St. Petersburg paradox). Suppose a wealthy stranger offers to play the following game with you. You will flip a fair coin until it lands Heads for the first time, and you will receive \$2 if the game lasts for 1 round, \$4 if the game lasts for 2 rounds, \$8 if the game lasts for 3 rounds, and in general, $\$2^n$ if the game lasts for n rounds. What is the fair value of this game (the expected payoff)? How much would you be willing to pay to play this game once?

Solution:

Let X be your winnings from playing the game. By definition, $X = 2^N$ where N is the number of rounds that the game lasts. Then X is 2 with probability $1/2$, 4 with probability $1/4$, 8 with probability $1/8$, and so on, so

$$E(X) = \frac{1}{2} \cdot 2 + \frac{1}{4} \cdot 4 + \frac{1}{8} \cdot 8 + \cdots = \infty.$$

The expected winnings are infinite! On the other hand, the number of rounds N that the game lasts is the number of tosses until the first Heads, so $N \sim \text{FS}(1/2)$ and $E(N) = 2$. Thus $E(2^N) = \infty$ while $2^{E(N)} = 4$. Infinity certainly does not equal 4, illustrating the danger of confusing $E(g(X))$ with $g(E(X))$ when g is not linear.

This problem is often considered a paradox because although the game's expected payoff is infinite, most people would not be willing to pay very much to play the game (even if they could afford to lose the money). One explanation is to note that the amount of money in the real world is finite. Suppose that if the game lasts longer than 40 rounds, the wealthy stranger flees the country and you get nothing. Since $2^{40} \approx 1.1 \times 10^{12}$, this still gives you the potential to earn over a trillion dollars, and anyway it's incredibly unlikely that the game will last longer than 40 rounds. But in this setting, your expected value is

$$E(X) = \sum_{n=1}^{40} \frac{1}{2^n} \cdot 2^n + \sum_{n=41}^{\infty} \frac{1}{2^n} \cdot 0 = 40.$$

Is this drastic reduction because the wealthy stranger may flee the country? Let's suppose instead that the wealthy stranger caps your winnings at 2^{40} , so if the game lasts more than 40 rounds you will get this amount rather than walking away empty-handed. Now your expected value is

$$E(X) = \sum_{n=1}^{40} \frac{1}{2^n} \cdot 2^n + \sum_{n=41}^{\infty} \frac{1}{2^n} \cdot 2^{40} = 40 + 1 = 41$$

an increase of only \$1 from the previous scenario. The ∞ in the St. Petersburg paradox is driven by an infinite "tail" of extremely rare events where you get extremely large payoffs. Cutting off this tail at some point, which makes sense in the real world, dramatically reduces the expected value of the game.

4.4.4 4.4 Indicator r.v.s and the Fundamental Bridge

This section is devoted to indicator random variables, which we already encountered in the previous chapter but will treat in much greater detail here. In particular, we will show that indicator r.v.s are an extremely useful tool for calculating expected values.

Recall from the previous chapter that the indicator r.v. I_A (or $I(A)$) for an event A is defined to be 1 if A occurs and 0 otherwise. So I_A is a Bernoulli random variable, where success is defined as " A occurs" and failure is defined as " A does not occur". Some useful properties of indicator r.v.s are summarized below.

Theorem 4.4.1 (Indicator r.v. properties). Let A and B be events. Then the following properties hold.

1. $(I_A)^k = I_A$ for any positive integer k .
2. $I_{A^c} = 1 - I_A$.
3. $I_{A \cap B} = I_A I_B$.
4. $I_{A \cup B} = I_A + I_B - I_A I_B$.

Proof. Property 1 holds since $0^k = 0$ and $1^k = 1$ for any positive integer k . Property 2 holds since $1 - I_A$ is 1 if A does not occur and 0 if A occurs. Property 3 holds since $I_A I_B$ is 1 if both I_A and I_B are 1, and 0 otherwise. Property 4 holds since

$$I_{A \cup B} = 1 - I_{A^c \cap B^c} = 1 - I_{A^c} I_{B^c} = 1 - (1 - I_A)(1 - I_B) = I_A + I_B - I_A I_B$$

Indicator r.v.s provide a link between probability and expectation; we call this fact the fundamental bridge.

Theorem 4.4.2 (Fundamental bridge between probability and expectation). There is a one-to-one correspondence between events and indicator r.v.s, and the probability of an event A is the expected value of its indicator r.v. I_A :

$$P(A) = E(I_A).$$

Proof. For any event A , we have an indicator r.v. I_A . This is a one-to-one correspondence since A uniquely determines I_A and vice versa (to get from I_A back to A , we can use the fact that $A = \{s \in S : I_A(s) = 1\}$). Since $I_A \sim \text{Bern}(p)$ with $p = P(A)$, we have $E(I_A) = P(A)$.

The fundamental bridge connects events to their indicator r.v.s, and allows us to express any probability as an expectation. As an example, we give a short proof of inclusion-exclusion and a related inequality known as Boole's inequality or Bonferroni's inequality using indicator r.v.s.

Example 4.4.3 (Boole, Bonferroni, and inclusion-exclusion). Let $A_1, A_2, \dots, A_n$ be events. Note that

$$I(A_1 \cup \dots \cup A_n) \leq I(A_1) + \dots + I(A_n)$$

since if the left-hand side is 0 this is immediate, and if the left-hand side is 1 then at least one term on the right-hand side must be 1. Taking the expectation of both sides and using linearity and the fundamental bridge, we have

$$P(A_1 \cup \cdots \cup A_n) \leq P(A_1) + \cdots + P(A_n),$$

which is called Boole's inequality or Bonferroni's inequality. To prove inclusion-exclusion for $n = 2$, we can take the expectation of both sides in Property 4 of Theorem 4.4.1. For general n , we can use properties of indicator r.v.s as follows:

$$\begin{aligned} 1 - I(A_1 \cup \cdots \cup A_n) &= I(A_1^c \cap \cdots \cap A_n^c) \\ &= (1 - I(A_1)) \cdots (1 - I(A_n)) \\ &= 1 - \sum_i I(A_i) + \sum_{i < j} I(A_i) I(A_j) - \cdots + (-1)^n I(A_1) \cdots I(A_n). \end{aligned}$$

Taking the expectation of both sides, by the fundamental bridge we have proven the inclusion-exclusion theorem.

Conversely, the fundamental bridge is also extremely useful in many expected value problems. We can often express a complicated discrete r.v. whose distribution we don't know as a sum of indicator r.v.s, which are extremely simple. The fundamental bridge lets us find the expectation of the indicators; then, using linearity, we obtain the expectation of our original r.v. This strategy is extremely useful and versatile in fact, we already used it when deriving the expectations of the Binomial and Hypergeometric distributions earlier in this chapter!

Recognizing problems that are amenable to this strategy and then defining the indicator r.v.s takes practice, so it is important to study a lot of examples and solve a lot of problems. In applying the strategy to a random variable that counts the number of [noun]s, we should have an indicator for each potential [noun]. This [noun] could be a person, place, or thing; we will see examples of all three types.

We'll start by revisiting two problems from Chapter 1, de Montmort's matching problem and the birthday problem.

Example 4.4.4 (Matching continued). We have a well-shuffled deck of n cards, labeled 1 through n . A card is a match if the card's position in the deck matches the card's label. Let X be the number of matches; find $E(X)$.

Solution:

First let's check whether X could have any of the named distributions we have studied. The Binomial and Hypergeometric are the only two candidates since the value of X must be an integer between 0 and n . But neither of these distributions has the right support because X can't take on the value $n - 1$: if $n - 1$ cards are matches, then the n th card must be a match as well. So X does not follow a named distribution we have studied, but we can readily find its mean using indicator r.v.s: let's write $X = I_1 + I_2 + \cdots + I_n$, where

$$I_j = \begin{cases} 1 & \text{if the } j \text{ th card in the deck is a match,} \\ 0 & \text{otherwise.} \end{cases}$$

In other words, I_j is the indicator for A_j , the event that the j th card in the deck is a match. We can imagine that each I_j "raises its hand" to be counted if its card is a match; adding up the raised hands, we get the total number of matches, X .

By the fundamental bridge,

$$E(I_j) = P(A_j) = \frac{1}{n}$$

for all j . So by linearity,

$$E(X) = E(I_1) + \cdots + E(I_n) = n \cdot \frac{1}{n} = 1.$$

The expected number of matched cards is 1, regardless of n . Even though the I_j are dependent in a complicated way that makes the distribution of X neither Binomial nor Hypergeometric, linearity still holds.

Example 4.4.5 (Distinct birthdays, birthday matches). In a group of n people, under the usual assumptions about birthdays, what is the expected number of distinct birthdays among the n people, i.e., the expected number of days on which at least one of the people was born? What is the expected number of birthday matches, i.e., pairs of people with the same birthday?

Solution:

Let X be the number of distinct birthdays, and write $X = I_1 + \cdots + I_{365}$, where

$$I_j = \begin{cases} 1 & \text{if the } j \text{ th day is represented,} \\ 0 & \text{otherwise.} \end{cases}$$

We create an indicator for each day of the year because X counts the number of days of the year that are represented. By the fundamental bridge,

$$E(I_j) = P(j \text{ th day is represented}) = 1 - P(\text{no one born on day } j) = 1 - \left(\frac{364}{365}\right)^n$$

for all j . Then by linearity,

$$E(X) = 365 \left(1 - \left(\frac{364}{365}\right)^n\right).$$

Now let Y be the number of birthday matches. Label the people as $1, 2, \dots, n$, and order the $\binom{n}{2}$ pairs of people in some definite way. Then we can write

$$Y = J_1 + \cdots + J_{\binom{n}{2}}$$

where J_i is the indicator of the i th pair of people having the same birthday. We create an indicator for each pair of people since Y counts the number of pairs of people with the same birthday. The probability of any two people having the same birthday is $1/365$, so again by the fundamental bridge and linearity,

$$E(Y) = \frac{\binom{n}{2}}{365}$$

In addition to the fundamental bridge and linearity, the last two examples used a basic form of symmetry to simplify the calculations greatly: within each sum of indicator r.v.s, each indicator had the same expected value. For example, in the matching problem the probability of the j th card being a match does not depend on j , so we can just take n times the expected value of the first indicator r.v.

Other forms of symmetry can also be extremely helpful when available. The next two examples showcase a form of symmetry that stems from having equally likely permutations. Note how symmetry, linearity, and the fundamental bridge are used in tandem to make seemingly very hard problems manageable.

Example 4.4.6 (Putnam problem). A permutation $a_1, a_2, \dots, a_n$ of $1, 2, \dots, n$ has a local maximum at j if $a_j > a_{j-1}$ and $a_j > a_{j+1}$ (for $2 \leq j \leq n-1$; for $j = 1$, a local maximum at j means $a_1 > a_2$ while for $j = n$, it means $a_n > a_{n-1}$). For example, $4, 2, 5, 3, 6, 1$ has 3 local maxima, at positions 1, 3, and 5. The Putnam exam (a famous, hard math competition, on which the median score is often a 0) from 2006 posed the following question: for $n \geq 2$, what is the average number of local maxima of a random permutation of $1, 2, \dots, n$, with all $n!$ permutations equally likely?

Solution:

This problem can be solved quickly using indicator r.v.s, symmetry, and the fundamental bridge. Let $I_1, \dots, I_n$ be indicator r.v.s, where I_j is 1 if there is a local maximum at position j , and 0 otherwise. We are interested in the expected value of $\sum_{j=1}^n I_j$. For $1 < j < n$, $EI_j = 1/3$ since having a local maximum at j is equivalent to a_j being the largest of a_{j-1}, a_j, a_{j+1} , which has probability $1/3$ since all orders are equally likely. For $j = 1$ or $j = n$, we have $EI_j = 1/2$ since then there is only one neighbor. Thus, by linearity,

$$E\left(\sum_{j=1}^n I_j\right) = 2 \cdot \frac{1}{2} + (n-2) \cdot \frac{1}{3} = \frac{n+1}{3}.$$

□

The next example introduces the Negative Hypergeometric distribution, which completes the following table. The table shows the distributions for four sampling schemes: the sampling can be done with or without replacement, and the stopping rule can require a fixed number of draws or a fixed number of successes.

	With replacement	Without replacement
Fixed number of trials	Binomial	Hypergeometric
Fixed number of successes	Negative Binomial	Negative Hypergeometric

Example 4.4.7 (Negative Hypergeometric). An urn contains w white balls and b black balls, which are randomly drawn one by one without replacement, until r white balls have been obtained. The number of black balls drawn before drawing the r th white ball has a Negative Hypergeometric distribution with parameters w, b, r . We denote this distribution by $\text{NHGeom}(w, b, r)$. Of course, we assume that $r \leq w$. For example, if we shuffle a deck of cards and deal them one at a time, the number of cards dealt before uncovering the first ace is $\text{NHGeom}(4, 48, 1)$.

As another example, suppose a college offers g good courses and b bad courses (for some definition of "good" and "bad"), and a student wants to find 4 good courses to take. Not having any idea which of the courses are good, the student randomly tries out courses one at a time, stopping when they have obtained 4 good courses. Then the number of bad courses the student tries out is $\text{NHGeom}(g, b, 4)$.

We can obtain the PMF of $X \sim \text{NHGeom}(w, b, r)$ by noting that, in the urn context, $X = k$ means that the $(r+k)$ th ball chosen is white and exactly $r-1$ of the first $r+k-1$ balls chosen are white. This gives

$$P(X = k) = \frac{\binom{w}{r-1} \binom{b}{k}}{\binom{w+b}{r+k-1}} \cdot \frac{w-r+1}{w+b-r-k+1}$$

for $k = 0, 1, \dots, b$ (and 0 otherwise).

Alternatively, we can imagine that we continue drawing balls until the urn has been emptied out; this is valid since whether or not we continue to draw balls after obtaining the r th white

ball has no effect on X . Think of the $w + b$ balls as lined up in a random order, the order in which they will be drawn.

Then $X = k$ means that among the first $r + k - 1$ balls there are exactly $r - 1$ white balls, then there is a white ball, and then among the last $w + b - r - k$ balls there are exactly $w - r$ white balls. All $\binom{w+b}{w}$ possibilities for the locations of the white balls in the line are equally likely. So by the naive definition of probability, we have the following slightly simpler expression for the PMF:

$$P(X = k) = \frac{\binom{r+k-1}{r-1} \binom{w+b-r-k}{w-r}}{\binom{w+b}{w}},$$

for $k = 0, 1, \dots, b$ (and 0 otherwise).

Finding the expected value of a Negative Hypergeometric r.v. directly from the definition of expectation results in complicated sums. But the answer is very simple: for $X \sim \text{NHGeom}(w, b, r)$, we have $E(X) = rb/(w + 1)$.

Let's prove this using indicator r.v.s. As explained above, we can assume that we continue drawing balls until the urn is empty. First consider the case $r = 1$. Label the black balls as $1, 2, \dots, b$, and let I_j be the indicator of black ball j being drawn before any white balls have been drawn. Then $P(I_j = 1) = 1/(w + 1)$ since, listing out the order in which black ball j and the white balls are drawn (ignoring the other balls), all orders are equally likely by symmetry, and $I_j = 1$ is equivalent to black ball j being first in this list. So by linearity,

$$E\left(\sum_{j=1}^b I_j\right) = \sum_{j=1}^b E(I_j) = b/(w + 1).$$

Sanity check: This answer makes sense since it is increasing in b , decreasing in w , and correct in the extreme cases $b = 0$ (when no black balls will be drawn) and $w = 0$ (when all the black balls will be exhausted before drawing a nonexistent white ball). Moreover, note that $b/(w + 1)$ looks similar to, but is strictly smaller than, b/w , which is the expected value of a $\text{Geom}(w/(w + b))$ r.v. It makes sense that sampling without replacement should give a smaller expected waiting time than sampling with replacement. Similarly, if you are searching for something you lost, it makes more sense to choose locations to check without replacement, rather than wasting time looking over and over again in locations you already ruled out.

For general r , write $X = X_1 + X_2 + \dots + X_r$, where X_1 is the number of black balls before the first white ball, X_2 is the number of black balls after the first white ball but before the second white ball, etc. By essentially the same argument we used to handle the $r = 1$ case, we have $E(X_j) = b/(w + 1)$ for each j . So by linearity,

$$E(X) = rb/(w + 1)$$

Closely related to indicator r.v.s is an alternative expression for the expectation of a nonnegative integer-valued r.v. X . Rather than summing up values of X times values of the PMF of X , we can sum up probabilities of the form $P(X > n)$ (known as tail probabilities), over nonnegative integers n .

Theorem 4.4.8 (Expectation via survival function). Let X be a nonnegative integer-valued r.v. Let F be the CDF of X , and $G(x) = 1 - F(x) = P(X > x)$. The function G is called the survival function of X . Then

$$E(X) = \sum_{n=0}^{\infty} G(n).$$

That is, we can obtain the expectation of X by summing up the survival function (or, stated otherwise, summing up tail probabilities of the distribution).

Proof. For simplicity, we will prove the result only for the case that X is bounded, i.e., there is a nonnegative integer b such that X is always at most b . We can represent X as a sum of indicator r.v.s: $X = I_1 + I_2 + \cdots + I_b$, where $I_n = I(X \geq n)$. For example, if $X = 7$ occurs, then I_1 through I_7 equal 1 while the other indicators equal 0.

By linearity and the fundamental bridge, and the fact that $\{X \geq k\}$ is the same event as $\{X > k - 1\}$,

$$E(X) = \sum_{k=1}^b E(I_k) = \sum_{k=1}^b P(X \geq k) = \sum_{n=0}^{b-1} P(X > n) = \sum_{n=0}^{\infty} G(n).$$

As a quick example, we use the above result to give another derivation of the mean of a Geometric r.v.

Example 4.4.9 (Geometric expectation redux). Let $X \sim \text{Geom}(p)$, and $q = 1 - p$. Using the Geometric story, $\{X > n\}$ is the event that the first $n + 1$ trials are all failures. So by Theorem 4.4.8,

$$E(X) = \sum_{n=0}^{\infty} P(X > n) = \sum_{n=0}^{\infty} q^{n+1} = \frac{q}{1 - q} = \frac{q}{p}$$

confirming what we already knew about the mean of a Geometric. $\square$

4.4.5 4.5 Law of the Unconscious Statistician (lotus)

As we saw in the St. Petersburg paradox, $E(g(X))$ does not equal $g(E(X))$ in general if g is not linear. So how do we correctly calculate $E(g(X))$? Since $g(X)$ is an r.v., one way is to first find the distribution of $g(X)$ and then use the definition of expectation. Perhaps surprisingly, it turns out that it is possible to find $E(g(X))$ directly using the distribution of X , without first having to find the distribution of $g(X)$. This is done using the law of the unconscious statistician (LOTUS).

Theorem 4.5.1 (LOTUS). If X is a discrete r.v. and g is a function from $\mathbb{R}$ to $\mathbb{R}$, then

$$E(g(X)) = \sum_x g(x)P(X = x),$$

where the sum is taken over all possible values of X .

This means that we can get the expected value of $g(X)$ knowing only $P(X = x)$, the PMF of X ; we don't need to know the PMF of $g(X)$. The name comes from the fact that in going from $E(X)$ to $E(g(X))$ it is tempting just to change x to $g(x)$ in the definition, which can be done very easily and mechanically, perhaps in a state of unconsciousness. On second thought, it may sound too good to be true that finding the distribution of $g(X)$ is not needed for this calculation, but LOTUS says it is true.

Before proving LOTUS in general, let's see why it is true in some special cases. Let X have support $0, 1, 2, \dots$ with probabilities $p_0, p_1, p_2, \dots$, so the PMF is $P(X = n) = p_n$. Then X^3 has support $0^3, 1^3, 2^3, \dots$ with probabilities $p_0, p_1, p_2, \dots$, so

$$\begin{aligned} E(X) &= \sum_{n=0}^{\infty} np_n, \\ E(X^3) &= \sum_{n=0}^{\infty} n^3 p_n. \end{aligned}$$

As claimed by LOTUS, to edit the expression for $E(X)$ into an expression for $E(X^3)$, we can just change the n in front of the p_n to an n^3 . This was an easy example since the function $g(x) = x^3$ is one-to-one. But LOTUS holds much more generally. The key insight needed for the proof of LOTUS for general g is the same as the one we used for the proof of linearity: the expectation of $g(X)$ can be written in ungrouped form as

$$E(g(X)) = \sum_s g(X(s))P(\{s\}),$$

where the sum is over all the pebbles in the sample space, but we can also group the pebbles into super-pebbles according to the value that X assigns to them. Within the super-pebble $X = x$, $g(X)$ always takes on the value $g(x)$. Therefore,

$$\begin{aligned} E(g(X)) &= \sum_s g(X(s))P(\{s\}) \\ &= \sum_x \sum_{s: X(s)=x} g(X(s))P(\{s\}) \\ &= \sum_x g(x) \sum_{s: X(s)=x} P(\{s\}) \\ &= \sum_x g(x)P(X = x) \end{aligned}$$

In the last step, we used the fact that $\sum_{s: X(s)=x} P(\{s\})$ is the weight of the superpebble $X = x$.

4.4.6 4.6 Variance

One important application of LOTUS is for finding the variance of a random variable. Like expected value, variance is a single-number summary of the distribution of a random variable. While the expected value tells us the center of mass of a distribution, the variance tells us how spread out the distribution is.

Definition 4.6.1 (Variance and standard deviation). The variance of an r.v. X is

$$\text{Var}(X) = E(X - EX)^2$$

The square root of the variance is called the standard deviation (SD):

$$\text{SD}(X) = \sqrt{\text{Var}(X)}.$$

Recall that when we write $E(X - EX)^2$, we mean the expectation of the random variable $(X - EX)^2$, not $(E(X - EX))^2$ (which is 0 by linearity).

The variance of X measures how far X is from its mean on average, but instead of simply taking the average difference between X and its mean EX , we take the average squared difference. To see why, note that the average deviation from the mean, $E(X - EX)$, always equals 0 by linearity; positive and negative deviations cancel each other out. By squaring the deviations, we ensure that both positive and negative deviations contribute to the overall variability. However, because variance is an average squared distance, it has the wrong units: if X is in dollars, $\text{Var}(X)$ is in squared dollars. To get back to our original units, we take the square root; this gives us the standard deviation.

One might wonder why variance isn't defined as $E|X - EX|$, which would achieve the goal of counting both positive and negative deviations while maintaining the same units as X . This measure of variability isn't nearly as popular as $E(X - EX)^2$, for a variety of reasons. Most

notably, the absolute value function isn't differentiable at 0, whereas the squaring function is differentiable everywhere and is central in various fundamental mathematical results such as the Pythagorean theorem.

An equivalent expression for variance is $\text{Var}(X) = E(X^2) - (EX)^2$. This formula is often easier to work with when doing actual calculations. Since this is the variance formula we will use over and over again, we state it as its own theorem.

Theorem 4.6.2. For any r.v. X ,

$$\text{Var}(X) = E(X^2) - (EX)^2.$$

Proof. Let $\mu = EX$. Expanding $(X - \mu)^2$ and using linearity, the variance of X is

$$E(X - \mu)^2 = E(X^2 - 2\mu X + \mu^2) = E(X^2) - 2\mu EX + \mu^2 = E(X^2) - \mu^2.$$

Variance has the following properties. The first two are easily verified from the definition, the third will be addressed in a later chapter, and the last one is proven just after stating it.

- $\text{Var}(X + c) = \text{Var}(X)$ for any constant c . Intuitively, if we shift a distribution to the left or right, that should affect the center of mass of the distribution but not its spread.
- $\text{Var}(cX) = c^2 \text{Var}(X)$ for any constant c .
- If X and Y are independent, then $\text{Var}(X + Y) = \text{Var}(X) + \text{Var}(Y)$. We prove this and discuss it more in Chapter 7. This is not true in general if X and Y are dependent. For example, in the extreme case where X always equals Y , we have

$$\text{Var}(X + Y) = \text{Var}(2X) = 4 \text{Var}(X) > 2 \text{Var}(X) = \text{Var}(X) + \text{Var}(Y)$$

if $\text{Var}(X) > 0$ (which will be true unless X is a constant, as the next property shows).

- $\text{Var}(X) \geq 0$, with equality if and only if $P(X = a) = 1$ for some constant a . In other words, the only random variables that have zero variance are constants (which can be thought of as degenerate r.v.s); all other r.v.s have positive variance.

To prove the last property, note that $\text{Var}(X)$ is the expectation of the nonnegative r.v. $(X - EX)^2$, so $\text{Var}(X) \geq 0$. If $P(X = a) = 1$ for some constant a , then $E(X) = a$ and $E(X^2) = a^2$, so $\text{Var}(X) = 0$. Conversely, suppose that $\text{Var}(X) = 0$. Then $E(X - EX)^2 = 0$, which shows that $(X - EX)^2 = 0$ has probability 1, which in turn shows that X equals its mean with probability 1.

4.6.3 (Variance is not linear). Unlike expectation, variance is not linear. The constant comes out squared in $\text{Var}(cX) = c^2 \text{Var}(X)$, and the variance of the sum of r.v.s may not be the sum of their variances if they are dependent.

Example 4.6.4 (Geometric and Negative Binomial variance). In this example we'll use LOTUS to compute the variance of the Geometric distribution.

Let $X \sim \text{Geom}(p)$. We already know $E(X) = q/p$. By LOTUS,

$$E(X^2) = \sum_{k=0}^{\infty} k^2 P(X = k) = \sum_{k=0}^{\infty} k^2 pq^k = \sum_{k=1}^{\infty} k^2 pq^k$$

We'll find this using a tactic similar to how we found the expectation, starting from the geometric series

$$\sum_{k=0}^{\infty} q^k = \frac{1}{1-q}$$

and taking derivatives. After differentiating once with respect to q , we have

$$\sum_{k=1}^{\infty} kq^{k-1} = \frac{1}{(1-q)^2}$$

We start the sum from $k = 1$ since the $k = 0$ term is 0 anyway. If we differentiate again, we'll get $k(k-1)$ instead of k^2 as we want, so let's replenish our supply of q 's by multiplying both sides by q . This gives

$$\sum_{k=1}^{\infty} kq^k = \frac{q}{(1-q)^2}$$

Now we are ready to take another derivative:

$$\sum_{k=1}^{\infty} k^2 q^{k-1} = \frac{1+q}{(1-q)^3}$$

so

$$E(X^2) = \sum_{k=1}^{\infty} k^2 p q^k = p q \frac{1+q}{(1-q)^3} = \frac{q(1+q)}{p^2}$$

Finally,

$$\text{Var}(X) = E(X^2) - (EX)^2 = \frac{q(1+q)}{p^2} - \left(\frac{q}{p}\right)^2 = \frac{q}{p^2}.$$

This is also the variance of the First Success distribution, since shifting by a constant does not affect the variance.

Since an $\text{NBin}(r, p)$ r.v. can be represented as a sum of r i.i.d. $\text{Geom}(p)$ r.v.s by Theorem 4.3.10, and since variance is additive for independent random variables, it follows that the variance of the $\text{NBin}(r, p)$ distribution is $r \cdot \frac{q}{p^2}$.

LOTUS is an all-purpose tool for computing $E(g(X))$ for any g , but as it usually leads to complicated sums, it should be used as a last resort. For variance calculations, our trusty indicator r.v.s can sometimes be used in place of LOTUS, as in the next example.

Example 4.6.5 (Binomial variance). Let's find the variance of $X \sim \text{Bin}(n, p)$ using indicator r.v.s to avoid tedious sums. Represent $X = I_1 + I_2 + \cdots + I_n$, where I_j is the indicator of the j th trial being a success. Each I_j has variance

$$\text{Var}(I_j) = E(I_j^2) - (E(I_j))^2 = p - p^2 = p(1-p).$$

(Recall that $I_j^2 = I_j$, so $E(I_j^2) = E(I_j) = p$.)

Since the I_j are independent, we can add their variances to get the variance of their sum:

$$\text{Var}(X) = \text{Var}(I_1) + \cdots + \text{Var}(I_n) = np(1-p).$$

Alternatively, we can find $E(X^2)$ by first finding $E\binom{X}{2}$. The latter sounds more complicated, but actually it is simpler since $\binom{X}{2}$ is the number of pairs of successful trials. Creating an indicator r.v. for each pair of trials, we have

$$E\binom{X}{2} = \binom{n}{2}p^2.$$

Thus,

$$n(n-1)p^2 = E(X(X-1)) = E(X^2) - E(X) = E(X^2) - np$$

which again gives

$$\text{Var}(X) = E(X^2) - (EX)^2 = (n(n-1)p^2 + np) - (np)^2 = np(1-p).$$

Exercise 48 uses this strategy to find the variance of the Hypergeometric.

4.4.7 4.7 Poisson

The last discrete distribution that we'll introduce in this chapter is the Poisson, which is an extremely popular distribution for modeling discrete data. We'll introduce its PMF, mean, and variance, and then discuss its story in more detail.

Definition 4.7.1 (Poisson distribution). An r.v. X has the Poisson distribution with parameter λ , where $\lambda > 0$, if the PMF of X is

$$P(X = k) = \frac{e^{-\lambda}\lambda^k}{k!}, \quad k = 0, 1, 2, \dots$$

We write this as $X \sim \text{Pois}(\lambda)$.

This is a valid PMF because of the Taylor series $\sum_{k=0}^{\infty} \frac{\lambda^k}{k!} = e^\lambda$.

Example 4.7.2 (Poisson expectation and variance). Let $X \sim \text{Pois}(\lambda)$. We will show that the mean and variance are both equal to λ . For the mean, we have

$$\begin{aligned} E(X) &= e^{-\lambda} \sum_{k=0}^{\infty} k \frac{\lambda^k}{k!} \\ &= e^{-\lambda} \sum_{k=1}^{\infty} k \frac{\lambda^k}{k!} \\ &= \lambda e^{-\lambda} \sum_{k=1}^{\infty} \frac{\lambda^{k-1}}{(k-1)!} \\ &= \lambda e^{-\lambda} e^\lambda = \lambda \end{aligned}$$

First we dropped the $k = 0$ term because it was 0. Then we took a λ out of the sum so that what was left inside was just the Taylor series for e^λ .

To get the variance, we first find $E(X^2)$. By LOTUS,

$$E(X^2) = \sum_{k=0}^{\infty} k^2 P(X = k) = e^{-\lambda} \sum_{k=0}^{\infty} k^2 \frac{\lambda^k}{k!}.$$

From here, the derivation is very similar to that of the variance of the Geometric. Differentiate the familiar series

$$\sum_{k=0}^{\infty} \frac{\lambda^k}{k!} = e^\lambda$$

with respect to λ and replenish:

$$\sum_{k=1}^{\infty} k \frac{\lambda^{k-1}}{k!} = e^{\lambda}$$

$$\sum_{k=1}^{\infty} k \frac{\lambda^k}{k!} = \lambda e^{\lambda}$$

Rinse and repeat:

$$\sum_{k=1}^{\infty} k^2 \frac{\lambda^{k-1}}{k!} = e^{\lambda} + \lambda e^{\lambda} = e^{\lambda}(1 + \lambda)$$

$$\sum_{k=1}^{\infty} k^2 \frac{\lambda^k}{k!} = e^{\lambda} \lambda (1 + \lambda)$$

Finally,

$$E(X^2) = e^{-\lambda} \sum_{k=0}^{\infty} k^2 \frac{\lambda^k}{k!} = e^{-\lambda} e^{\lambda} \lambda (1 + \lambda) = \lambda(1 + \lambda),$$

so

$$\text{Var}(X) = E(X^2) - (EX)^2 = \lambda(1 + \lambda) - \lambda^2 = \lambda.$$

Thus, the mean and variance of a $\text{Pois}(\lambda)$ r.v. are both equal to λ . $\square$

Figure 4.7 shows the PMF and CDF of the $\text{Pois}(2)$ and $\text{Pois}(5)$ distributions from $k = 0$ to $k = 10$. It appears that the mean of the $\text{Pois}(2)$ is around 2 and the mean of the $\text{Pois}(5)$ is around 5, consistent with our findings above. The PMF of the $\text{Pois}(2)$ is highly skewed, but as λ grows larger, the skewness is reduced and the PMF becomes more bell-shaped.




FIGURE 4.7

Top: $\text{Pois}(2)$ PMF and CDF. Bottom: $\text{Pois}(5)$ PMF and CDF.

The Poisson distribution is often used in situations where we are counting the number of successes in a particular region or interval of time, and there are a large number of

trials, each with a small probability of success. For example, the following random variables could follow a distribution that is approximately Poisson.

- The number of emails you receive in an hour. There are a lot of people who could potentially email you in that hour, but it is unlikely that any specific person will actually email you in that hour. Alternatively, imagine subdividing the hour into milliseconds. There are 3.6×10^6 seconds in an hour, but in any specific millisecond it is unlikely that you will get an email.
- The number of chips in a chocolate chip cookie. Imagine subdividing the cookie into small cubes; the probability of getting a chocolate chip in a single cube is small, but the number of cubes is large.
- The number of earthquakes in a year in some region of the world. At any given time and location, the probability of an earthquake is small, but there are a large number of possible times and locations for earthquakes to occur over the course of the year.

The parameter λ is interpreted as the rate of occurrence of these rare events; in the examples above, λ could be 20 (emails per hour), 10 (chips per cookie), and 2 (earthquakes per year). The Poisson paradigm says that in applications similar to the ones above, we can approximate the distribution of the number of events that occur by a Poisson distribution.

Approximation 4.7.3 (Poisson paradigm). Let $A_1, \dots, A_n$ be events with $p_j = P(A_j)$, where n is large, the p_j are small, and the A_j are independent or weakly dependent. Let

$$X = \sum_{j=1}^n I(A_j)$$

count how many of the A_j occur. Then X is approximately distributed as $\text{Pois}(\lambda)$, with $\lambda = \sum_{j=1}^n p_j$.

Proving that the above approximation is good is difficult, and would require first giving precise definitions of weak dependence (there are various ways to measure dependence of r.v.s) and of good approximations (there are various ways to measure how good an approximation is). A remarkable theorem is that if the A_j are independent, $N \sim \text{Pois}(\lambda)$, and B is any set of nonnegative integers, then

$$|P(X \in B) - P(N \in B)| \leq \min\left(1, \frac{1}{\lambda}\right) \sum_{j=1}^n p_j^2$$

This gives an upper bound on how much error is incurred from using a Poisson approximation. It also makes more precise how small the p_j should be: we want $\sum_{j=1}^n p_j^2$ to be very small, or at least very small compared to λ . The result can be shown using an advanced technique known as the Stein-Chen method.

The Poisson paradigm is also called the law of rare events. The interpretation of "rare" is that the p_j are small, not that λ is small. For example, in the email example, the low probability of getting an email from a specific person in a particular hour is offset by the large number of people who could send you an email in that hour.

In the examples we gave above, the number of events that occur isn't exactly Poisson because a Poisson random variable has no upper bound, whereas how many of $A_1, \dots, A_n$ occur is at most n , and there is a limit to how many chocolate chips can be crammed into a cookie. But the Poisson distribution often gives good approximations. Note that the conditions for the Poisson paradigm to hold are fairly flexible: the n trials can have different success probabilities, and the trials don't have to be independent, though they should not be very dependent. So there

are a wide variety of situations that can be cast in terms of the Poisson paradigm. This makes the Poisson a popular model, or at least a starting point, for data whose values are nonnegative integers (called count data in statistics).

Example 4.7.4 (Balls in boxes). There are k distinguishable balls and n distinguishable boxes. The balls are randomly placed in the boxes, with all n^k possibilities equally likely. Problems in this setting are called occupancy problems, and are at the core of many widely used algorithms in computer science.

- (a) Find the expected number of empty boxes (fully simplified, not as a sum).
 (b) Find the probability that at least one box is empty. Express your answer as a sum of at most n terms.
 (c) Now let $n = 1000, k = 5806$. The expected number of empty boxes is then approximately 3. Find a good approximation as a decimal for the probability that at least one box is empty. The handy fact $e^3 \approx 20$ may help.

Solution:

- (a) Let I_j be the indicator r.v. for the j th box being empty. Then

$$E(I_j) = P(I_j = 1) = \left(1 - \frac{1}{n}\right)^k$$

By linearity,

$$E\left(\sum_{j=1}^n I_j\right) = \sum_{j=1}^n E(I_j) = n \left(1 - \frac{1}{n}\right)^k.$$

- (b) The probability is 1 for $k < n$. In general, let A_j be the event that box j is empty. By inclusion-exclusion,

$$\begin{aligned} P(A_1 \cup A_2 \cup \cdots \cup A_n) &= \sum_{j=1}^n (-1)^{j+1} \binom{n}{j} P(A_1 \cap A_2 \cap \cdots \cap A_j) \\ &= \sum_{j=1}^{n-1} (-1)^{j+1} \binom{n}{j} \left(1 - \frac{j}{n}\right)^k. \end{aligned}$$

- (c) The number X of empty boxes is approximately $\text{Pois}(3)$, since there are a lot of boxes but each is very unlikely to be empty; the probability that a specific box is empty is $\left(1 - \frac{1}{n}\right)^k = \frac{1}{n}$. $E(X) \approx 0.003$. So

$$P(X \geq 1) = 1 - P(X = 0) \approx 1 - e^{-3} \approx 1 - \frac{1}{20} = 0.95.$$

Poisson approximation greatly simplifies obtaining a good approximate solution to the birthday problem discussed in Chapter 1, and makes it possible to obtain good approximations to various variations which would be hard to solve exactly.

Example 4.7.5 (Birthday problem continued). If we have m people and make the usual assumptions about birthdays, then each pair of people has probability $p = 1/365$ of having the same birthday, and there are $\binom{m}{2}$ pairs. By the Poisson paradigm the distribution of the number X of birthday matches is approximately $\text{Pois}(\lambda)$, where $\lambda = \binom{m}{2} \frac{1}{365}$. Then the probability of at least one match is

$$P(X \geq 1) = 1 - P(X = 0) \approx 1 - e^{-\lambda}$$

For $m = 23$, $\lambda = 253/365$ and $1 - e^{-\lambda} \approx 0.500002$, which agrees with our finding from Chapter 1 that we need 23 people to have a 50 – 50 chance of a matching birthday.

Note that even though $m = 23$ is fairly small, the relevant quantity in this problem is actually $\binom{m}{2}$, which is the total number of "trials" for a successful birthday match, so the Poisson approximation still performs well.

Example 4.7.6 (Near-birthday problem). What if we want to find the number of people required in order to have a 50 – 50 chance that two people would have birthdays within one day of each other (i.e., on the same day or one day apart)? Unlike the original birthday problem, this is difficult to obtain an exact answer for, but the Poisson paradigm still applies. The probability that any two people have birthdays within one day of each other is $3/365$ (choose a birthday for the first person, and then the second person needs to be born on that day, the day before, or the day after). Again there are $\binom{m}{2}$ possible pairs, so the number of within-one-day matches is approximately $\text{Pois}(\lambda)$ where $\lambda = \binom{m}{2} \frac{3}{365}$. Then a calculation similar to the one above tells us that we need $m = 14$ or more. This was a quick approximation, but it turns out that $m = 14$ is the exact answer!

Example 4.7.7 (Birth-minute and birth-hour). There are 1600 sophomores at a certain college. Throughout this example, make the usual assumptions as in the birthday problem.

(a) Find a Poisson approximation for the probability that there are two sophomores who were born not only on the same day of the year, but also at the same hour and the same minute (e.g., both sophomores were born at 8:20 pm on March 31, not necessarily in the same year).

(b) With assumptions as in (a), what is the probability that there are four sophomores who were born not only on the same day, but also at the same hour (e.g., all were born between 2 pm and 3 pm on March 31, not necessarily in the same year)?

Give two different Poisson approximations for this value, one based on creating an indicator r.v. for each quadruplet of sophomores, and the other based on creating an indicator r.v. for each possible day-hour. Which do you think is more accurate?

Solution:

(a) This is the birthday problem, with $c = 365 \cdot 24 \cdot 60 = 525600$ categories rather than 365 categories. ¹ Let $n = 1600$. Creating an indicator r.v. for each pair of sophomores, by linearity the expected number of pairs born on the same day-hourminute is

$$\lambda_1 = \binom{n}{2} \frac{1}{c}.$$

By Poisson approximation, the probability of at least one match is approximately

$$1 - \exp(-\lambda_1) \approx 0.9122$$

This approximation is very accurate: typing `pbirthday(1600, classes = 365 * 24 * 60)` in R yields 0.9125 .

(b) Now there are $b = 365 \cdot 24 = 8760$ categories. Let's explore two different methods of Poisson approximation.

Method 1: Create an indicator for each set of 4 sophomores. By linearity, the expected number of sets of 4 sophomores born on the same day-hour is

$$\lambda_2 = \binom{n}{4} \frac{1}{b^3}$$

Poisson approximation gives that the desired probability is approximately

$$1 - \exp(-\lambda_2) \approx 0.333$$

Method 2: Create an indicator for each possible day-hour. Let I_j be the indicator for at least 4 people having been born on the j th day-hour of the year (ordered chronologically), for $1 \leq j \leq b$. Let $p = 1/b$ and $q = 1 - p$. Then

$$\begin{aligned} E(I_j) &= P(I_j = 1) \\ &= 1 - P(\text{at most 3 people born on the } j \text{ th day-hour}) \\ &= 1 - q^n - npq^{n-1} - \binom{n}{2}p^2q^{n-2} - \binom{n}{3}p^3q^{n-3} \end{aligned}$$

The expected number of day-hours on which at least 4 sophomores were born is

$$\lambda_3 = b \cdot E(I_1)$$

with $E(I_1)$ as above. We then have the Poisson approximation

$$1 - \exp(-\lambda_3) \approx 0.295$$

The command `pbirthday(1600, classes = 8760, coincident = 4)` in *R* gives that the correct answer is 0.296. So Method 2 is more accurate than Method 1.

An intuitive explanation for why Method 1 is less accurate is that there is a more substantial dependence in the indicators in that method. For example, being given that sophomores 1, 2, 3, 4 share the same birth day-hour greatly increases the chance that sophomores 1, 2, 3, 5 share the same birth day-hour. In contrast, knowing that at least 4 sophomores were born on a specific day-hour provides very little information about whether at least 4 were born on a different specific day-hour.

4.4.8 4.8 Connections between Poisson and Binomial

The Poisson and Binomial distributions are closely connected, and their relationship is exactly parallel to the relationship between the Binomial and Hypergeometric distributions that we examined in the previous chapter: we can get from the Poisson to the Binomial by conditioning, and we can get from the Binomial to the Poisson by taking a limit.

Our results will rely on the fact that the sum of independent Poissons is Poisson, just as the sum of independent Binomials is Binomial. We'll prove this result using the law of total probability for now; in Chapter 6 we'll learn a faster method that uses a tool called the moment generating function. Chapter 13 gives further insight into these results.

Theorem 4.8.1 (Sum of independent Poissons). If $X \sim \text{Pois}(\lambda_1)$, $Y \sim \text{Pois}(\lambda_2)$, and X is independent of Y , then $X + Y \sim \text{Pois}(\lambda_1 + \lambda_2)$.

Proof. To get the PMF of $X + Y$, condition on X and use the law of total probability:

⁰¹ The song "Seasons of Love" from *Rent* gives a musical interpretation of this fact.

$$\begin{aligned}
P(X + Y = k) &= \sum_{j=0}^k P(X + Y = k \mid X = j)P(X = j) \\
&= \sum_{j=0}^k P(Y = k - j)P(X = j) \\
&= \sum_{j=0}^k \frac{e^{-\lambda_2} \lambda_2^{k-j}}{(k-j)!} \frac{e^{-\lambda_1} \lambda_1^j}{j!} \\
&= \frac{e^{-(\lambda_1 + \lambda_2)}}{k!} \sum_{j=0}^k \binom{k}{j} \lambda_1^j \lambda_2^{k-j} \\
&= \frac{e^{-(\lambda_1 + \lambda_2)} (\lambda_1 + \lambda_2)^k}{k!}.
\end{aligned}$$

The last step used the binomial theorem. Since we've arrived at the $\text{Pois}(\lambda_1 + \lambda_2)$ PMF, we have $X + Y \sim \text{Pois}(\lambda_1 + \lambda_2)$.

The story of the Poisson distribution provides intuition for this result. If there are two different types of events occurring at rates λ_1 and λ_2 , independently, then the overall event rate is $\lambda_1 + \lambda_2$.

Theorem 4.8.2 (Poisson given a sum of Poissons). If $X \sim \text{Pois}(\lambda_1)$, $Y \sim \text{Pois}(\lambda_2)$, and X is independent of Y , then the conditional distribution of X given $X + Y = n$ is $\text{Bin}(n, \lambda_1 / (\lambda_1 + \lambda_2))$.

Proof. Exactly as in the corresponding proof for the Binomial and Hypergeometric, we use Bayes' rule to compute the conditional PMF $P(X = k \mid X + Y = n)$:

$$\begin{aligned}
P(X = k \mid X + Y = n) &= \frac{P(X + Y = n \mid X = k)P(X = k)}{P(X + Y = n)} \\
&= \frac{P(Y = n - k)P(X = k)}{P(X + Y = n)}.
\end{aligned}$$

Now we plug in the PMFs of X , Y , and $X + Y$; the last of these is distributed $\text{Pois}(\lambda_1 + \lambda_2)$ by the previous theorem. This gives

$$\begin{aligned}
P(X = k \mid X + Y = n) &= \frac{\left(\frac{e^{-\lambda_2} \lambda_2^{n-k}}{(n-k)!} \right) \left(\frac{e^{-\lambda_1} \lambda_1^k}{k!} \right)}{\frac{e^{-(\lambda_1 + \lambda_2)} (\lambda_1 + \lambda_2)^n}{n!}} \\
&= \binom{n}{k} \frac{\lambda_1^k \lambda_2^{n-k}}{(\lambda_1 + \lambda_2)^n} \\
&= \binom{n}{k} \left(\frac{\lambda_1}{\lambda_1 + \lambda_2} \right)^k \left(\frac{\lambda_2}{\lambda_1 + \lambda_2} \right)^{n-k},
\end{aligned}$$

which is the $\text{Bin}(n, \lambda_1 / (\lambda_1 + \lambda_2))$ PMF, as desired.

Conversely, if we take the limit of the $\text{Bin}(n, p)$ distribution as $n \rightarrow \infty$ and $p \rightarrow 0$ with np fixed, we arrive at a Poisson distribution. This provides the basis for the Poisson approximation to the Binomial distribution.

Theorem 4.8.3 (Poisson approximation to Binomial). If $X \sim \text{Bin}(n, p)$ and we let $n \rightarrow \infty$ and $p \rightarrow 0$ such that $\lambda = np$ remains fixed, then the PMF of X converges to the $\text{Pois}(\lambda)$ PMF. More generally, the same conclusion holds if $n \rightarrow \infty$ and $p \rightarrow 0$ in such a way that np converges to a constant λ .

This is a special case of the Poisson paradigm, where the A_j are independent with the same probabilities, so that $\sum_{j=1}^n I(A_j)$ has a Binomial distribution. In this special case, we can prove that the Poisson approximation makes sense just by taking a limit of the Binomial PMF.

Proof. We will prove this for the case that $\lambda = np$ is fixed while $n \rightarrow \infty$ and $p \rightarrow 0$, by showing that the $\text{Bin}(n, p)$ PMF converges to the $\text{Pois}(\lambda)$ PMF. For $0 \leq k \leq n$,

$$\begin{aligned} P(X = k) &= \binom{n}{k} p^k (1-p)^{n-k} \\ &= \frac{n(n-1)\dots(n-k+1)}{k!} \left(\frac{\lambda}{n}\right)^k \left(1 - \frac{\lambda}{n}\right)^n \left(1 - \frac{\lambda}{n}\right)^{-k} \\ &= \frac{\lambda^k}{k!} \frac{n(n-1)\dots(n-k+1)}{n^k} \left(1 - \frac{\lambda}{n}\right)^n \left(1 - \frac{\lambda}{n}\right)^{-k}. \end{aligned}$$

Letting $n \rightarrow \infty$ with k fixed,

$$\begin{aligned} \frac{n(n-1)\dots(n-k+1)}{n^k} &\rightarrow 1, \\ \left(1 - \frac{\lambda}{n}\right)^n &\rightarrow e^{-\lambda}, \\ \left(1 - \frac{\lambda}{n}\right)^{-k} &\rightarrow 1, \end{aligned}$$

where the $e^{-\lambda}$ comes from the compound interest formula from Section A.2.5 of the math appendix. So

$$P(X = k) \rightarrow \frac{e^{-\lambda} \lambda^k}{k!},$$

which is the $\text{Pois}(\lambda)$ PMF.

This theorem implies that if n is large, p is small, and np is moderate, we can approximate the $\text{Bin}(n, p)$ PMF by the $\text{Pois}(np)$ PMF. The main thing that matters here is that p should be small; in fact, the result mentioned after the statement of the Poisson paradigm says in this case that the error in approximating $P(X \in B) \approx P(N \in B)$ for $X \sim \text{Bin}(n, p)$, $N \sim \text{Pois}(np)$ is at most $\min(p, np^2)$.

Example 4.8.4 (Visitors to a website). The owner of a certain website is studying the distribution of the number of visitors to the site. Every day, a million people independently decide whether to visit the site, with probability $p = 2 \times 10^{-6}$ of visiting. Give a good approximation for the probability of getting at least three visitors on a particular day.

Solution:

Let $X \sim \text{Bin}(n, p)$ be the number of visitors, where $n = 10^6$. It is easy to run into computational difficulties or numerical errors in exact calculations with this distribution since n is so large and p is so small. But since n is large, p is small, and $np = 2$ is moderate, $\text{Pois}(2)$ is a good approximation. This gives

$$P(X \geq 3) = 1 - P(X < 3) \approx 1 - e^{-2} - e^{-2} \cdot 2 - e^{-2} \cdot \frac{2^2}{2!} = 1 - 5e^{-2} \approx 0.3233$$

which turns out to be extremely accurate.

4.4.9 4.9 *Using Probability and Expectation to Prove Existence

An amazing and beautiful fact is that we can use probability and expectation to prove the existence of objects with properties we care about. This technique is called the probabilistic method, and it is based on two simple but surprisingly powerful ideas. Suppose I want to show that there exists an object in a collection with a certain property. This desire seems at first to have nothing to do with probability; I could simply examine each object in the collection one by one until finding an object with the desired property.

The probabilistic method rejects such painstaking inspection in favor of random selection: our strategy is to pick an object at random from the collection and show that there is a positive probability of the random object having the desired property. Note that we are not required to compute the exact probability, but merely to show it is greater than 0. If we can show that the probability of the property holding is positive, then we know that there must exist an object with the property—even if we don't know how to explicitly construct such an object.

Similarly, suppose each object has a score, and I want to show that there exists an object with a "good" score - that is, a score exceeding a particular threshold. Again, we proceed by choosing a random object and considering its score, X . We know there is an object in the collection whose score is at least $E(X)$ —it's impossible for every object to be below average! If $E(X)$ is already a good score, then there must also be an object in the collection with a good score. Thus we can show the existence of an object with a good score by showing that the average score is already good.

Let's state the two key ideas formally.

- The possibility principle: Let A be the event that a randomly chosen object in a collection has a certain property. If $P(A) > 0$, then there exists an object with the property.
- The good score principle: Let X be the score of a randomly chosen object. If $E(X) \geq c$, then there is an object with a score of at least c .

To see why the possibility principle is true, consider its contrapositive: if there is no object with the desired property, then the probability of a randomly chosen object having the property is 0. Similarly, the contrapositive of the good score principle is "if all objects have a score below c , then the average score is below c ", which is true since a weighted average of numbers less than c is a number less than c .

The probabilistic method doesn't tell us how to find an object with the desired property; it only assures us that one exists.

Example 4.9.1. A group of 100 people are assigned to 15 committees of size 20, such that each person serves on 3 committees. Show that there exist 2 committees that have at least 3 people in common.

Solution:

A direct approach is inadvisable here: one would have to list all possible committee assignments and compute, for each one, the number of people in common in every pair of committees. The probabilistic method lets us bypass brute-force calculations. To prove the existence of two committees with an overlap of at least three people, we'll calculate the average overlap of two randomly chosen committees in an arbitrary committee assignment. So choose two committees at random, and let X be the number of people on both committees. We can represent $X = I_1 + I_2 + \cdots + I_{100}$, where $I_j = 1$ if the j th person is on both committees and 0 otherwise. By symmetry, all of the indicators have the same expected value, so $E(X) = 100E(I_1)$, and we just need to find $E(I_1)$.

By the fundamental bridge, $E(I_1)$ is the probability that person 1 (whom we'll name Bob) is on both committees (which we'll call A and B). There are a variety of ways to calculate this probability; one way is to think of Bob's committees as 3 tagged elk in a population of 15. Then A and B are a sample of 2 elk, made without replacement. Using the HGeom(3, 12, 2) PMF, the probability that both of these elk are tagged (i.e., the probability that both committees contain Bob) is $\binom{3}{2}\binom{12}{0}/\binom{15}{2} = 1/35$. Therefore,

$$E(X) = 100/35 = 20/7,$$

which is just shy of the desired "good score" of 3. But hope is not lost! The good score principle says there exist two committees with an overlap of at least $20/7$, but since the overlap between two committees must be an integer, an overlap of at least $20/7$ implies an overlap of at least 3. Thus, there exist two committees with at least 3 people in common.

4.9.1 *communicating over a Noisy Channel

Another major application of the probabilistic method is in information theory, the subject which studies (among other things) how to achieve reliable communication across a noisy channel. Consider the problem of trying to send a message when there is noise. This problem is encountered by millions of people every day, such as when talking on the phone (you may be misheard). Suppose that the message you want to send is represented as a binary vector $x \in \{0, 1\}^k$, and that you want to use a code to improve the chance that your message will get through successfully.

Definition 4.9.2 (Codes and rates). Given positive integers k and n , a code is a function c that assigns to each input message $x \in \{0, 1\}^k$ a codeword $c(x) \in \{0, 1\}^n$. The rate of this code is k/n (the number of input bits per output bit). After $c(x)$ is sent, a decoder takes the received message, which may be a corrupted version of $c(x)$, and attempts to recover the correct x .

For example, an obvious code would be to repeat yourself a bunch of times, sending x a bunch of times in a row, say m (with m odd); this is called a repetition code. The receiver could then decode by going with the majority, e.g., decoding the first bit of x as a 1 if that bit was received more times as a 1 than as a 0. But this code may be very inefficient; to get the probability of failure very small, you may need to repeat yourself many times, resulting in a very low rate $1/m$ of communication.

Claude Shannon, the founder of information theory, showed something amazing: even in a very noisy channel, there is a code allowing for very reliable communication at a rate that does not go to 0 as we require the probability of failure to be lower and lower. His proof was even more amazing: he studied the performance of a completely random code. Richard Hamming, who worked with Shannon at Bell Labs, described Shannon's approach as follows.

Courage is another attribute of those who do great things. Shannon is a good example. For some time he would come to work at about 10:00 am, play chess until about 2:00 pm and go home.

The important point is how he played chess. When attacked he seldom, if ever, defended his position, rather he attacked back. Such a method of playing soon produces a very interrelated board. He would then pause a bit, think and advance his queen saying, "I ain't [scared] of nothin'." It took me a while to realize that of course that is why he was able to prove the existence of good coding methods. Who but Shannon would think to average over all random codes and expect to find that the average was close to ideal? I learned from him to say the same to myself when stuck, and on some occasions his approach enabled me to get significant results. [15]

We will prove a version of Shannon's result, for the case of a channel where each transmitted bit gets flipped (from 0 to 1 or from 1 to 0) with probability p , independently. First we need

two definitions. A natural measure of distance between binary vectors, named after Hamming, is as follows.

Definition 4.9.3 (Hamming distance). For two binary vectors v and w of the same length, the Hamming distance $d(v, w)$ is the number of positions in which they differ. We can write this as

$$d(v, w) = \sum_i |v_i - w_i|$$

The following function arises very frequently in information theory.

Definition 4.9.4 (Binary entropy function). For $0 < p < 1$, the binary entropy function H is given by

$$H(p) = -p \log_2 p - (1 - p) \log_2 (1 - p)$$

We also define $H(0) = H(1) = 0$.

The interpretation of $H(p)$ in information theory is that it is a measure of how much information we get from observing a $\text{Bern}(p)$ r.v.; $H(1/2) = 1$ says that a fair coin flip provides 1 bit of information, while $H(1) = 0$ says that with a coin that always lands Heads, there's no information gained from being told the result of the flip, since we already know the result.

Now consider a channel where each transmitted bit gets flipped with probability p , independently. Intuitively, it may seem that smaller p is always better, but note that $p = 1/2$ is actually the worst-case scenario. In that case, technically known as a useless channel, it is impossible to send information over the channel: the output will be independent of the input! Analogously, in deciding whether to watch a movie, would you rather hear a review from someone you always disagree with or someone you agree with half the time? We now prove that for $0 < p < 1/2$, it is possible to communicate very reliably with rate very close to $1 - H(p)$.

Theorem 4.9.5 (Shannon). Consider a channel where each transmitted bit gets flipped with probability p , independently. Let $0 < p < 1/2$ and $\epsilon > 0$. There exists a code with rate at least $1 - H(p) - \epsilon$ that can be decoded with probability of error less than ϵ .

Proof. We can assume that $1 - H(p) - \epsilon > 0$, since otherwise there is no constraint on the rate. Let n be a large positive integer (chosen according to conditions given below), and

$$k = \lceil n(1 - H(p) - \epsilon) \rceil + 1.$$

The ceiling function is there since k must be an integer. Choose $p' \in (p, 1/2)$ such that $|H(p') - H(p)| < \epsilon/2$ (this can be done since H is continuous). We will now study the performance of a random code C . To generate a random code C , we need to generate a random encoded message $C(x)$ for all possible input messages x .

For each $x \in \{0, 1\}^k$, choose $C(x)$ to be a uniformly random vector in $\{0, 1\}^n$ (making these choices independently). So we can think of $C(x)$ as a vector consisting of n i.i.d. $\text{Bern}(1/2)$ r.v.s. The rate k/n exceeds $1 - H(p) - \epsilon$ by definition, but let's see how well we can decode the received message!

Let $x \in \{0, 1\}^k$ be the input message, $C(x)$ be the encoded message, and $Y \in \{0, 1\}^n$ be the received message. For now, treat x as deterministic. But $C(x)$ is random since the codewords are chosen randomly, and Y is random since $C(x)$ is random and due to the random noise in the channel. Intuitively, we hope that $C(x)$ will be close to Y (in Hamming distance) and $C(z)$ will be far from Y for all $z \neq x$, in which case it will be clear how to decode Y and the decoding will succeed. To make this precise, decode Y as follows:

If there exists a unique $z \in \{0, 1\}^k$ such that $d(C(z), Y) \leq np'$, decode Y to that z ; otherwise, declare decoder failure.

We will show that for n large enough, the probability of the decoder failing to recover the correct x is less than ϵ . There are two things that could go wrong:

- (a) $d(C(x), Y) > np'$, or
- (b) There could be some impostor $z \neq x$ with $d(C(z), Y) \leq np'$.

Note that $d(C(x), Y)$ is an r.v., so $d(C(x), Y) > np'$ is an event. To handle (a), represent

$$d(C(x), Y) = B_1 + \cdots + B_n \sim \text{Bin}(n, p)$$

where B_i is the indicator of the i th bit being flipped. The law of large numbers (see Chapter 10) says that as n grows, the r.v. $d(C(x), Y)/n$ will get very close to p (its expected value), and so will be very unlikely to exceed p' :

$$P(d(C(x), Y) > np') = P\left(\frac{B_1 + \cdots + B_n}{n} > p'\right) \rightarrow 0 \text{ as } n \rightarrow \infty.$$

So by choosing n large enough, we can make

$$P(d(C(x), Y) > np') < \epsilon/4.$$

To handle (b), note that $d(C(z), Y) \sim \text{Bin}(n, 1/2)$ for $z \neq x$, since the n bits in $C(z)$ are i.i.d. $\text{Bern}(1/2)$, independent of Y (to show this in more detail, condition on Y using LOTP). Let $B \sim \text{Bin}(n, 1/2)$. By Boole's inequality,

$$P(d(C(z), Y) \leq np' \text{ for some } z \neq x) \leq (2^k - 1) P(B \leq np').$$

To simplify notation, suppose that np' is an integer. A crude way to upper bound a sum of m terms is to use m times the largest term, and a crude way to upper bound a binomial coefficient $\binom{n}{j}$ is to use $r^{-j}(1-r)^{-(n-j)}$ for any $r \in (0, 1)$. Combining these two crudities,

$$P(B \leq np') = \frac{1}{2^n} \sum_{j=0}^{np'} \binom{n}{j} \leq \frac{np' + 1}{2^n} \binom{n}{np'} \leq (np' + 1) 2^{nH(p')-n},$$

using the fact that $(p')^{-np'} (q')^{-nq'} = 2^{nH(p')}$ for $q' = 1 - p'$. Thus,

$$2^k P(B \leq np') \leq (np' + 1) 2^{n(1-H(p)-\epsilon)+2+n(H(p)+\epsilon/2)-n} = 4(np' + 1) 2^{-n\epsilon/2} \rightarrow 0,$$

so we can choose n to make $P(d(C(z), Y) \leq np' \text{ for some } z \neq x) < \epsilon/4$.

Assume that k and n have been chosen in accordance with the above, and let $F(c, x)$ be the event of failure when code c is used with input message x . Putting together the above results, we have shown that for a random C and any fixed x ,

$$P(F(C, x)) < \epsilon/2.$$

It follows that for each x , there is a code c with $P(F(c, x)) < \epsilon/2$, but this is not good enough: we want one code that works well for all x ! Let X be a uniformly random input message in $\{0, 1\}^k$, independent of C . By LOTP, we have

$$P(F(C, X)) = \sum_x P(F(C, x))P(X = x) < \epsilon/2.$$

Again using LOTP, but this time conditioning on C , we have

$$\sum_c P(F(c, X))P(C = c) = P(F(C, X)) < \epsilon/2.$$

Therefore, there exists a code c such that $P(F(c, X)) < \epsilon/2$, i.e., a code c such that the probability of failure for a random input message X is less than $\epsilon/2$. Lastly, we will improve c , obtaining a code that works well for all x , not just a random x . We do this by expurgating the worst 50% of the x 's. That is, remove as legal input messages the 2^{k-1} values of x with the highest failure probabilities for code c . For all remaining x , we have $P(F(c, x)) < \epsilon$, since otherwise more than half of the $x \in \{0, 1\}^k$ would have more than double the average failure probability (see Markov's inequality in Chapter 10 for more about this kind of argument). By relabeling the remaining x using vectors in $\{0, 1\}^{k-1}$, we obtain a code $c' : \{0, 1\}^{k-1} \rightarrow \{0, 1\}^n$ with rate $(k-1)/n \geq 1 - H(p) - \epsilon$ and probability less than ϵ of failure for all input messages in $\{0, 1\}^{k-1}$.

There is also a converse to the above theorem, showing that if we require the rate to be at least $1 - H(p) + \epsilon$, it is impossible to find codes that make the probability of error arbitrarily small. This is why $1 - H(p)$ is called the capacity of the channel. Shannon also obtained analogous results for much more general channels. These results give theoretical bounds on what can be achieved, without saying explicitly which codes to use. Decades of subsequent work have been devoted to developing specific codes that work well in practice, by coming close to the Shannon bound and allowing for efficient encoding and decoding.

4.4.10 4.10 Recap

The expectation of a discrete r.v. X is

$$E(X) = \sum_x xP(X = x)$$

An equivalent "ungrouped" way of calculating expectation is

$$E(X) = \sum_s X(s)P(\{s\}),$$

where the sum is taken over pebbles in the sample space. Expectation is a single number summarizing the center of mass of a distribution. A single-number summary of the spread of a distribution is the variance, defined by

$$\text{Var}(X) = E(X - EX)^2 = E(X^2) - (EX)^2.$$

The square root of the variance is called the standard deviation.

Expectation is linear:

$$E(cX) = cE(X) \text{ and } E(X + Y) = E(X) + E(Y),$$

regardless of whether X and Y are independent or not. Variance is not linear:

$$\text{Var}(cX) = c^2 \text{Var}(X),$$

and

$$\text{Var}(X + Y) \neq \text{Var}(X) + \text{Var}(Y)$$

in general (an important exception is when X and Y are independent).

A very important strategy for calculating the expectation of a discrete r.v. X is to express it as a sum of indicator r.v.s, and then apply linearity and the fundamental bridge. This technique is especially powerful because the indicator r.v.s need not be independent; linearity holds even for dependent r.v.s. The strategy can be summarized in the following three steps.

1. Represent the r.v. X as a sum of indicator r.v.s. To decide how to define the indicators, think about what X is counting. For example, if X is the number of local maxima, as in the Putnam problem, then we should create an indicator for each local maximum that could occur.
2. Use the fundamental bridge to calculate the expected value of each indicator. When applicable, symmetry may be very helpful at this stage.
3. By linearity of expectation, $E(X)$ can be obtained by adding up the expectations of the indicators.

Another tool for computing expectations is LOTUS, which says we can calculate the expectation of $g(X)$ using only the PMF of X , via

$$E(g(X)) = \sum_x g(x)P(X = x)$$

If g is non-linear, it is a grave mistake to attempt to calculate $E(g(X))$ by swapping the E and the g .

Four new discrete distributions to add to our list are the Geometric, Negative Binomial, Negative Hypergeometric, and Poisson distributions. A $\text{Geom}(p)$ r.v. is the number of failures before the first success in a sequence of independent Bernoulli trials with probability p of success, and an $\text{NBin}(r, p)$ r.v. is the number of failures before r successes. The Negative Hypergeometric is similar to the Negative Binomial except, in terms of drawing balls from an urn, the Negative Hypergeometric samples without replacement and the Negative Binomial samples with replacement. (We also introduced the First Success distribution, which is just a Geometric shifted so that the success is included.)

A Poisson r.v. is often used as an approximation for the number of successes that occur when there are many independent or weakly dependent trials, where each trial has a small probability of success. In the Binomial story, all the trials have the same probability p of success, but in the Poisson approximation, different trials can have different (but small) probabilities p_j of success.

The Poisson, Binomial, and Hypergeometric distributions are mutually connected via the operations of conditioning and taking limits, as illustrated in Figure 4.8. In the rest of this book, we'll continue to introduce new named distributions and add them to this family tree, until everything is connected!




FIGURE 4.8
Relationships between the Poisson, Binomial, and Hypergeometric.

Figure 4.9 expands upon the corresponding figure from the previous chapter, further exploring the connections between the four fundamental objects we have considered: distributions, random variables, events, and numbers.




FIGURE 4.9

Four fundamental objects in probability: distributions, random variables, events, and numbers. From an r.v. X , we can generate many other r.v.s by taking functions of X , and we can use LOTUS to find their expected values. The mean, variance, and standard deviation of X express the average and spread of the distribution of X (in particular, they only depend on F , not directly on X itself).

4.4.11 4.11r

Geometric, Negative Binomial, and Poisson

The three functions for the Geometric distribution in R are `dgeom`, `pgeom`, and `rgeom`, corresponding to the PMF, CDF, and random generation. For `dgeom` and `pgeom`, we need to supply the following as inputs: (1) the value at which to evaluate the PMF or CDF, and (2) the parameter p . For `rgeom`, we need to input (1) the number of random variables to generate and (2) the parameter p .

For example, to calculate $P(X = 3)$ and $P(X \leq 3)$ where $X \sim \text{Geom}(0.5)$, we use `dgeom(3, 0.5)` and `pgeom(3, 0.5)`, respectively. To generate 100 i.i.d. $\text{Geom}(0.8)$ r.v.s, we use `rgeom(100, 0.8)`. If instead we want 100 i.i.d. $\text{FS}(0.8)$ r.v.s, we just need to add 1 to include the success: `rgeom(100, 0.8) + 1`.

For the Negative Binomial distribution, we have `dnbinom`, `pnbinom`, and `rnbinom`.

These take three inputs. For example, to calculate the $\text{NBin}(5, 0.5)$ PMF at 3, we type `dnbinom(3, 5, 0.5)`.

Finally, for the Poisson distribution, the three functions are `dpois`, `ppois`, and `rpois`. These take two inputs. For example, to find the $\text{Pois}(10)$ CDF at 2, we type `ppois(2, 10)`.

Matching Simulation

Continuing with Example 4.4.4, let's use simulation to calculate the expected number of matches in a deck of cards. As in Chapter 1, we let n be the number of cards in the deck and perform the experiment 10^4 times using `replicate`.

```
n <- 100
```

```
r <- replicate(10^4, sum(sample(n) == (1:n)))
```

Now r contains the number of matches from each of the 10^4 simulations. But instead of looking at the probability of at least one match, as in Chapter 1, we now want to find the expected number of matches. We can approximate this by the average of all the simulation results, that

is, the arithmetic mean of the elements of r . This is accomplished with the mean function:
`mean(r)`

The command `mean( r )` is equivalent to `sum( r )/length( r )`. The result we get is very close to 1, confirming the calculation we did in Example 4.4.4 using indicator r.v.s. You can verify that no matter what value of n you choose, `mean( r )` will be very close to 1.

Distinct Birthdays Simulation

Let's calculate the expected number of distinct birthdays in a group of k people by simulation. We'll let $k = 20$, but you can choose whatever value of k you like.

```
k <- 20
r <- replicate (10^4, { bdays <- sample (365, k, replace = TRUE );
                        length(unique(bdays)) })
```

In the second line, `replicate` repeats the expression in the curly braces 10^4 times, so we just need to understand what is inside the curly braces. First, we sample k times with replacement from the numbers 1 through 365 and call these the birthdays of the k people, `bdays`. Then, `unique(bdays)` removes duplicates in the vector `bdays`, and `length(unique(bdays))` calculates the length of the vector after duplicates have been removed. The two commands need to be separated by a semicolon.

Now r contains the number of distinct birthdays that we observed in each of the 10^4 simulations. The average number of distinct birthdays across the 10^4 simulations is `mean(r)`. We can compare the simulated value to the theoretical value that we found in Example 4.4.5 using indicator r.v.s:

```
mean (r)
365 * (1 - (364/365)^k)
```

When we ran the code, both the simulated and theoretical values gave us approximately 19.5.

4.4.12 4.12 Exercises

Exercises marked with (S) have detailed solutions at <http://stat110.net>.

Expectations and Variances

1. Bobo, the amoeba from Chapter 2, currently lives alone in a pond. After one minute Bobo will either die, split into two amoebas, or stay the same, with equal probability. Find the expectation and variance for the number of amoebas in the pond after one minute.
2. In the Gregorian calendar, each year has either 365 days (a normal year) or 366 days (a leap year). A year is randomly chosen, with probability $3/4$ of being a normal year and $1/4$ of being a leap year. Find the mean and variance of the number of days in the chosen year.
3. (a) A fair die is rolled. Find the expected value of the roll.
(b) Four fair dice are rolled. Find the expected total of the rolls.
4. A fair die is rolled some number of times. You can choose whether to stop after 1, 2, or 3 rolls, and your decision can be based on the values that have appeared so far. You receive the value shown on the last roll of the die, in dollars. What is your optimal strategy (to maximize your expected winnings)? Find the expected winnings for this strategy.
Hint: Start by considering a simpler version of this problem, where there are at most 2 rolls. For what values of the first roll should you continue for a second roll?

5. Find the mean and variance of a Discrete Uniform r.v. on $1, 2, \dots, n$.

Hint: See the math appendix for some useful facts about sums.

6. Two teams are going to play a best-of-7 match (the match will end as soon as either team has won 4 games). Each game ends in a win for one team and a loss for the other team. Assume that each team is equally likely to win each game, and that the games played are independent. Find the mean and variance of the number of games played.

7. A certain small town, whose population consists of 100 families, has 30 families with 1 child, 50 families with 2 children, and 20 families with 3 children. The birth rank of one of these children is 1 if the child is the firstborn, 2 if the child is the secondborn, and 3 if the child is the thirdborn.

(a) A random family is chosen (with equal probabilities), and then a random child within that family is chosen (with equal probabilities). Find the PMF, mean, and variance of the child's birth rank.

(b) A random child is chosen in the town (with equal probabilities). Find the PMF, mean, and variance of the child's birth rank.

8. A certain country has four regions: North, East, South, and West. The populations of these regions are 3 million, 4 million, 5 million, and 8 million, respectively. There are 4 cities in the North, 3 in the East, 2 in the South, and there is only 1 city in the West. Each person in the country lives in exactly one of these cities.

(a) What is the average size of a city in the country? (This is the arithmetic mean of the populations of the cities, and is also the expected value of the population of a city chosen uniformly at random.)

Hint: Give the cities names (labels).

(b) Show that without further information it is impossible to find the variance of the population of a city chosen uniformly at random. That is, the variance depends on how the people within each region are allocated between the cities in that region.

(c) A region of the country is chosen uniformly at random, and then a city within that region is chosen uniformly at random. What is the expected population size of this randomly chosen city?

Hint: First find the selection probability for each city.

(d) Explain intuitively why the answer to (c) is larger than the answer to (a).

9. Consider the following simplified scenario based on Who Wants to Be a Millionaire?, a game show in which the contestant answers multiple-choice questions that have 4 choices per question. The contestant (Fred) has answered 9 questions correctly already, and is now being shown the 10th question. He has no idea what the right answers are to the 10th or 11th questions are. He has one "lifeline" available, which he can apply on any question, and which narrows the number of choices from 4 down to 2. Fred has the following options available.

(a) Walk away with \$16,000.

(b) Apply his lifeline to the 10th question, and then answer it. If he gets it wrong, he will leave with \$1,000. If he gets it right, he moves on to the 11th question. He then leaves with \$32,000 if he gets the 11th question wrong, and \$64,000 if he gets the 11th question right.

(c) Same as the previous option, except not using his lifeline on the 10th question, and instead applying it to the 11th question (if he gets the 10th question right).

Find the expected value of each of these options. Which option has the highest expected value? Which option has the lowest variance?

10. Consider the St. Petersburg paradox (Example 4.3.14), except that you receive $\$n$ rather than $\$2^n$ if the game lasts for n rounds. What is the fair value of this game? What if the payoff is $\$n^2$?

11. Martin has just heard about the following exciting gambling strategy: bet \$1 that a fair coin will land Heads. If it does, stop. If it lands Tails, double the bet for the next toss, now betting

\$2 on Heads. If it does, stop. Otherwise, double the bet for the next toss to \$4. Continue in this way, doubling the bet each time and then stopping right after winning a bet. Assume that each individual bet is fair, i.e., has an expected net winnings of 0. The idea is that

$$1 + 2 + 2^2 + 2^3 + \cdots + 2^n = 2^{n+1} - 1$$

so the gambler will be \$1 ahead after winning a bet, and then can walk away with a profit. Martin decides to try out this strategy. However, he only has \$31, so he may end up walking away bankrupt rather than continuing to double his bet. On average, how much money will Martin win?

12. Let X be a discrete r.v. with support $-n, -n+1, \dots, 0, \dots, n-1, n$ for some positive integer n . Suppose that the PMF of X satisfies the symmetry property $P(X = -k) = P(X = k)$ for all integers k . Find $E(X)$.

13. (S) Are there discrete random variables X and Y such that $E(X) > 100E(Y)$ but Y is greater than X with probability at least 0.99?

14. Let X have PMF

$$P(X = k) = cp^k/k \text{ for } k = 1, 2, \dots,$$

where p is a parameter with $0 < p < 1$ and c is a normalizing constant. We have $c = -1/\log(1-p)$, as seen from the Taylor series

$$-\log(1-p) = p + \frac{p^2}{2} + \frac{p^3}{3} + \dots$$

This distribution is called the Logarithmic distribution (because of the log in the above Taylor series), and has often been used in ecology. Find the mean and variance of X .

15. Player A chooses a random integer between 1 and 100, with probability p_j of choosing j (for $j = 1, 2, \dots, 100$). Player B guesses the number that player A picked, and receives from player A that amount in dollars if the guess is correct (and 0 otherwise).

(a) Suppose for this part that player B knows the values of p_j . What is player B's optimal strategy (to maximize expected earnings)?

(b) Show that if both players choose their numbers so that the probability of picking j is proportional to $1/j$, then neither player has an incentive to change strategies, assuming the opponent's strategy is fixed. (In game theory terminology, this says that we have found a Nash equilibrium.)

(c) Find the expected earnings of player B when following the strategy from (b). Express your answer both as a sum of simple terms and as a numerical approximation. Does the value depend on what strategy player A uses?

16. The dean of Blotchville University boasts that the average class size there is 20. But the reality experienced by the majority of students there is quite different: they find themselves in huge courses, held in huge lecture halls, with hardly enough seats or Haribo gummi bears for everyone. The purpose of this problem is to shed light on the situation. For simplicity, suppose that every student at Blotchville University takes only one course per semester.

(a) Suppose that there are 16 seminar courses, which have 10 students each, and 2 large lecture courses, which have 100 students each. Find the dean's-eye-view average class size (the simple average of the class sizes) and the student's-eye-view average class size (the average class size experienced by students, as it would be reflected by surveying students and asking them how big their classes are). Explain the discrepancy intuitively.

(b) Give a short proof that for any set of class sizes (not just those given above), the dean's-eye-view average class size will be strictly less than the student's-eye-view average class size, unless all classes have exactly the same size.

Hint: Relate this to the fact that variances are nonnegative.

17. The sociologist Elizabeth Wrigley-Field posed the following puzzle [29]:

American fertility fluctuated dramatically in the decades surrounding the Second World War. Parents created the smallest families during the Great Depression, and the largest families during the postwar Baby Boom. Yet children born during the Great Depression came from larger families than those born during the Baby Boom. How can this be?

(a) For a particular era, let n_k be the number of American families with exactly k children, for each $k \geq 0$. (Assume for simplicity that American history has cleanly been separated into eras, where each era has a well-defined set of families, and each family has a well-defined set of children; we are ignoring the fact that a particular family's size may change over time, that children grow up, etc.) For each $j \geq 0$, let

$$m_j = \sum_{k=0}^{\infty} k^j n_k$$

For a family selected randomly in that era (with all families equally likely), find the expected number of children in the family. Express your answer only in terms of the m_j 's.

(b) For a child selected randomly in that era (with all children equally likely), find the expected number of children in the child's family, only in terms of the m_j 's.

(c) Give an intuitive explanation in words for which of the answers to (a) and (b) is larger, or whether they are equal. Explain how this relates to the Wrigley-Field puzzle.

Named Distributions

18. (S) A fair coin is tossed repeatedly, until it has landed Heads at least once and has landed Tails at least once. Find the expected number of tosses.

19. (S) A coin is tossed repeatedly until it lands Heads for the first time. Let X be the number of tosses that are required (including the toss that landed Heads), and let p be the probability of Heads, so that $X \sim \text{FS}(p)$. Find the CDF of X , and for $p = 1/2$ sketch its graph.

20. Let $X \sim \text{Bin}(100, 0.9)$. For each of the following parts, construct an example showing that it is possible, or explain clearly why it is impossible. In this problem, Y is a random variable on the same probability space as X ; note that X and Y are not necessarily independent.

(a) Is it possible to have $Y \sim \text{Pois}(0.01)$ with $P(X \geq Y) = 1$?

(b) Is it possible to have $Y \sim \text{Bin}(100, 0.5)$ with $P(X \geq Y) = 1$?

(c) Is it possible to have $Y \sim \text{Bin}(100, 0.5)$ with $P(X \leq Y) = 1$?

21. (S) Let $X \sim \text{Bin}(n, \frac{1}{2})$ and $Y \sim \text{Bin}(n+1, \frac{1}{2})$, independently.

(a) Let $V = \min(X, Y)$ be the smaller of X and Y , and let $W = \max(X, Y)$ be the larger of X and Y . So if X crystallizes to x and Y crystallizes to y , then V crystallizes to $\min(x, y)$ and W crystallizes to $\max(x, y)$. Find $E(V) + E(W)$.

(b) Show that $E|X - Y| = E(W) - E(V)$, with notation as in (a).

(c) Compute $\text{Var}(n - X)$ in two different ways.

22. (S) Raindrops are falling at an average rate of 20 drops per square inch per minute. What would be a reasonable distribution to use for the number of raindrops hitting a particular region measuring 5 inches² in t minutes? Why? Using your chosen distribution, compute the probability that the region has no rain drops in a given 3-second time interval.

23. (S) Alice and Bob have just met, and wonder whether they have a mutual friend. Each has 50 friends, out of 1000 other people who live in their town. They think that it's unlikely that they have a friend in common, saying "each of us is only friends with 5% of the people here, so it would be very unlikely that our two 5% 's overlap."

Assume that Alice's 50 friends are a random sample of the 1000 people (equally likely to be any 50 of the 1000), and similarly for Bob. Also assume that knowing who Alice's friends are gives no information about who Bob's friends are.

- (a) Compute the expected number of mutual friends Alice and Bob have.
(b) Let X be the number of mutual friends they have. Find the PMF of X .
(c) Is the distribution of X one of the important distributions we have looked at? If so, which?
24. Let $X \sim \text{Bin}(n, p)$ and $Y \sim \text{NBin}(r, p)$. Using a story about a sequence of Bernoulli trials, prove that $P(X < r) = P(Y > n - r)$.

25. (S) Calvin and Hobbes play a match consisting of a series of games, where Calvin has probability p of winning each game (independently). They play with a "win by two" rule: the first player to win two games more than his opponent wins the match. Find the expected number of games played.

Hint: Consider the first two games as a pair, then the next two as a pair, etc.

26. Nick and Penny are independently performing independent Bernoulli trials. For concreteness, assume that Nick is flipping a nickel with probability p_1 of Heads and Penny is flipping a penny with probability p_2 of Heads. Let $X_1, X_2, \dots$ be Nick's results and $Y_1, Y_2, \dots$ be Penny's results, with $X_i \sim \text{Bern}(p_1)$ and $Y_j \sim \text{Bern}(p_2)$.

- (a) Find the distribution and expected value of the first time at which they are simultaneously successful, i.e., the smallest n such that $X_n = Y_n = 1$.

Hint: Define a new sequence of Bernoulli trials and use the story of the Geometric.

- (b) Find the expected time until at least one has a success (including the success).

Hint: Define a new sequence of Bernoulli trials and use the story of the Geometric.

- (c) For $p_1 = p_2$, find the probability that their first successes are simultaneous, and use this to find the probability that Nick's first success precedes Penny's.

27. (S) Let X and Y be $\text{Pois}(\lambda)$ r.v.s, and $T = X + Y$. Suppose that X and Y are not independent, and in fact $X = Y$. Prove or disprove the claim that $T \sim \text{Pois}(2\lambda)$ in this scenario.

28. William is on a treasure hunt. There are t pieces of treasure, each of which is hidden in one of n locations. William searches these locations one by one, without replacement, until he has found all the treasure. (Assume that no location contains more than one piece of treasure, and that William will find the treasure piece when he searches a location that does have treasure.) Let X be the number of locations that William searches during his treasure hunt. Find the distribution of X , and find $E(X)$.

29. Let $X \sim \text{Geom}(p)$, and define the function f by $f(x) = P(X = x)$, for all real x . Find $E(f(X))$. (The notation $f(X)$ means first evaluate $f(x)$ in terms of p and x , and then plug in X for x ; it is not correct to say " $f(X) = P(X = X) = 1$ ".)

30. (a) Use LOTUS to show that for $X \sim \text{Pois}(\lambda)$ and any function g ,

$$E(Xg(X)) = \lambda E(g(X+1)),$$

assuming that both sides exist. This is called the Stein-Chen identity for the Poisson.

- (b) Find the third moment $E(X^3)$ for $X \sim \text{Pois}(\lambda)$ by using the identity from (a) and a bit of algebra to reduce the calculation to the fact that X has mean λ and variance λ .

31. In many problems about modeling count data, it is found that values of zero in the data are far more common than can be explained well using a Poisson model (we can make $P(X = 0)$ large for $X \sim \text{Pois}(\lambda)$ by making λ small, but that also constrains the mean and variance of X

to be small since both are λ). The Zero-Inflated Poisson distribution is a modification of the Poisson to address this issue, making it easier to handle frequent zero values gracefully.

A Zero-Inflated Poisson r.v. X with parameters p and λ can be generated as follows. First flip a coin with probability of p of Heads. Given that the coin lands Heads, $X = 0$. Given that the coin lands Tails, X is distributed $\text{Pois}(\lambda)$. Note that if $X = 0$ occurs, there are two possible explanations: the coin could have landed Heads (in which case the zero is called a structural zero), or the coin could have landed Tails but the Poisson r.v. turned out to be zero anyway.

For example, if X is the number of chicken sandwiches consumed by a random person in a week, then $X = 0$ for vegetarians (this is a structural zero), but a chicken-eater could still have $X = 0$ occur by chance (since they might happen not to eat any chicken sandwiches that week).

- (a) Find the PMF of a Zero-Inflated Poisson r.v. X .
- (b) Explain why X has the same distribution as $(1 - I)Y$, where $I \sim \text{Bern}(p)$ is independent of $Y \sim \text{Pois}(\lambda)$.
- (c) Find the mean of X in two different ways: directly using the PMF of X , and using the representation from (b). For the latter, you can use the fact (which we prove in Chapter 7) that if r.v.s Z and W are independent, then $E(ZW) = E(Z)E(W)$.
- (d) Find the variance of X .

32. (S) A discrete distribution has the memoryless property if for X a random variable with that distribution, $P(X \geq j + k \mid X \geq j) = P(X \geq k)$ for all nonnegative integers j, k .

(a) If X has a memoryless distribution with CDF F and PMF $p_i = P(X = i)$, find an expression for $P(X \geq j + k)$ in terms of $F(j), F(k), p_j, p_k$.

(b) Name a discrete distribution which has the memoryless property. Justify your answer with a clear interpretation in words or with a computation.

33. Find values of w, b, r such that the Negative Hypergeometric distribution with parameters w, b, r reduces to a Discrete Uniform on $\{0, 1, \dots, n\}$. Justify your answer both in terms of the story of the Negative Hypergeometric and in terms of its PMF.

Indicator r.v.s

34. (S) Randomly, k distinguishable balls are placed into n distinguishable boxes, with all possibilities equally likely. Find the expected number of empty boxes.
35. (S) A group of 50 people are comparing their birthdays (as usual, assume their birthdays are independent, are not February 29, etc.). Find the expected number of pairs of people with the same birthday, and the expected number of days in the year on which at least two of these people were born.
36. (S) A group of $n \geq 4$ people are comparing their birthdays (as usual, assume their birthdays are independent, are not February 29, etc.). Let I_{ij} be the indicator r.v. of i and j having the same birthday (for $i < j$). Is I_{12} independent of I_{34} ? Is I_{12} independent of I_{13} ? Are the I_{ij} independent?
37. (S) A total of 20 bags of Haribo gummi bears are randomly distributed to 20 students. Each bag is obtained by a random student, and the outcomes of who gets which bag are independent. Find the average number of bags of gummi bears that the first three students get in total, and find the average number of students who get at least one bag.
38. Each of $n \geq 2$ people puts their name on a slip of paper (no two have the same name). The slips of paper are shuffled in a hat, and then each person draws one (uniformly at

- random at each stage, without replacement). Find the average number of people who draw their own names.
39. Two researchers independently select simple random samples from a population of size N , with sample sizes m and n (for each researcher, the sampling is done without replacement, with all samples of the prescribed size equally likely). Find the expected size of the overlap of the two samples.
40. In a sequence of n independent fair coin tosses, what is the expected number of occurrences of the pattern HTH (consecutively)? Note that overlap is allowed, e.g., $HTHTH$ contains two overlapping occurrences of the pattern.
41. You have a well-shuffled 52-card deck. On average, how many pairs of adjacent cards are there such that both cards are red?
42. Suppose there are n types of toys, which you are collecting one by one. Each time you collect a toy, it is equally likely to be any of the n types. What is the expected number of distinct toy types that you have after you have collected t toys? (Assume that you will definitely collect t toys, whether or not you obtain a complete set before then.)
43. A building has n floors, labeled $1, 2, \dots, n$. At the first floor, k people enter the elevator, which is going up and is empty before they enter. Independently, each decides which of floors $2, 3, \dots, n$ to go to and presses that button (unless someone has already pressed it).
(a) Assume for this part only that the probabilities for floors $2, 3, \dots, n$ are equal. Find the expected number of stops the elevator makes on floors $2, 3, \dots, n$.
(b) Generalize (a) to the case that floors $2, 3, \dots, n$ have probabilities $p_2, \dots, p_n$ (respectively); you can leave your answer as a finite sum.
44. (S) There are 100 shoelaces in a box. At each stage, you pick two random ends and tie them together. Either this results in a longer shoelace (if the two ends came from different pieces), or it results in a loop (if the two ends came from the same piece). What are the expected number of steps until everything is in loops, and the expected number of loops after everything is in loops? (This is a famous interview problem; leave the latter answer as a sum.)
Hint: For each step, create an indicator r.v. for whether a loop was created then, and note that the number of free ends goes down by 2 after each step.
45. Show that for any events $A_1, \dots, A_n$,

$$P(A_1 \cap A_2 \cdots \cap A_n) \geq \sum_{j=1}^n P(A_j) - n + 1$$

Hint: First prove a similar-looking statement about indicator r.v.s, by interpreting what the events $I(A_1 \cap A_2 \cdots \cap A_n) = 1$ and $I(A_1 \cap A_2 \cdots \cap A_n) = 0$ mean.

46. You have a well-shuffled 52-card deck. You turn the cards face up one by one, without replacement. What is the expected number of non-aces that appear before the first ace? What is the expected number between the first ace and the second ace?
47. You are being tested for psychic powers. Suppose that you do not have psychic powers. A standard deck of cards is shuffled, and the cards are dealt face down one by one. Just after each card is dealt, you name any card (as your prediction). Let X be the number of cards you predict correctly. (See Diaconis [5] for much more about the statistics of testing for psychic powers.)

(a) Suppose that you get no feedback about your predictions. Show that no matter what strategy you follow, the expected value of X stays the same; find this value. (On the other hand, the variance may be very different for different strategies. For example, saying "Ace of Spades" every time gives variance 0.)

Hint: Indicator r.v.s.

(b) Now suppose that you get partial feedback: after each prediction, you are told immediately whether or not it is right (but without the card being revealed). Suppose you use the following strategy: keep saying a specific card's name (e.g., "Ace of Spades") until you hear that you are correct. Then keep saying a different card's name (e.g., "Two of Spades") until you hear that you are correct (if ever). Continue in this way, naming the same card over and over again until you are correct and then switching to a new card, until the deck runs out. Find the expected value of X , and show that it is very close to $e - 1$.

Hint: Indicator r.v.s.

(c) Now suppose that you get complete feedback: just after each prediction, the card is revealed. Call a strategy "stupid" if it allows, e.g., saying "Ace of Spades" as a guess after the Ace of Spades has already been revealed. Show that any non-stupid strategy gives the same expected value for X ; find this value.

Hint: Indicator r.v.s.

48. (S) Let X be Hypergeometric with parameters w, b, n .

(a) Find $E\binom{X}{2}$ by thinking, without any complicated calculations.

(b) Use (a) to find the variance of X . You should get

$$\text{Var}(X) = \frac{N-n}{N-1}npq,$$

where $N = w + b, p = w/N, q = 1 - p$.

49. There are n prizes, with values \$1, \$2, ..., \$ n . You get to choose k random prizes, without replacement. What is the expected total value of the prizes you get?

Hint: Express the total value in the form $a_1I_1 + \cdots + a_nI_n$, where the a_j are constants and the I_j are indicator r.v.s. Or find the expected value of the j th prize received directly.

50. Ten random chords of a circle are chosen, independently. To generate each of these chords, two independent uniformly random points are chosen on the circle (intuitively, "uniformly" means that the choice is completely random, with no favoritism toward certain angles; formally, it means that the probability of any arc is proportional to the length of that arc). On average, how many pairs of chords intersect?

Hint: Consider two random chords. An equivalent way to generate them is to pick four independent uniformly random points on the circle, and then pair them up randomly.

51. (S) A hash table is being used to store the phone numbers of k people, storing each person's phone number in a uniformly random location, represented by an integer between 1 and n (see Exercise 27 from Chapter 1 for a description of hash tables). Find the expected number of locations with no phone numbers stored, the expected number with exactly one phone number, and the expected number with more than one phone number (should these quantities add up to n ?).

52. A coin with probability p of Heads is flipped n times. The sequence of outcomes can be divided into runs (blocks of H 's or blocks of T 's), e.g., $HHHTTHTTTTH$ becomes

HHH	TT	H	TTT	H

, which has 5 runs. Find the expected number of runs.

Hint: Start by finding the expected number of tosses (other than the first) where the outcome is different from the previous one.

53. A coin with probability p of Heads is flipped 4 times. Let X be the number of occurrences of HH (for example, $THHT$ has 1 occurrence and $HHHH$ has 3 occurrences). Find $E(X)$

and $\text{Var}(X)$.

54. A population has N people, with ID numbers from 1 to N . Let y_j be the value of some numerical variable for person j , and

$$\bar{y} = \frac{1}{N} \sum_{j=1}^N y_j$$

be the population average of the quantity. For example, if y_j is the height of person j then $\bar{y}$ is the average height in the population, and if y_j is 1 if person j holds a certain belief and 0 otherwise, then $\bar{y}$ is the proportion of people in the population who hold that belief. In this problem, $y_1, y_2, \dots, y_n$ are thought of as constants rather than random variables.

A researcher is interested in learning about $\bar{y}$, but it is not feasible to measure y_j for all j . Instead, the researcher gathers a random sample of size n , by choosing people one at a time, with equal probabilities at each stage and without replacement. Let W_j be the value of the numerical variable (e.g., height) for the j th person in the sample. Even though $y_1, \dots, y_n$ are constants, W_j is a random variable because of the random sampling. A natural way to estimate the unknown quantity $\bar{y}$ is using

$$\bar{W} = \frac{1}{n} \sum_{j=1}^n W_j$$

Show that $E(\bar{W}) = \bar{y}$ in two different ways:

- (a) by directly evaluating $E(W_j)$ using symmetry;
- (b) by showing that $\bar{W}$ can be expressed as a sum over the population by writing

$$\bar{W} = \frac{1}{n} \sum_{j=1}^N I_j y_j$$

where I_j is the indicator of person j being included in the sample, and then using linearity and the fundamental bridge.

55. (S) Consider the following algorithm, known as bubble sort, for sorting a list of n distinct numbers into increasing order. Initially they are in a random order, with all orders equally likely. The algorithm compares the numbers in positions 1 and 2, and swaps them if needed, then it compares the new numbers in positions 2 and 3, and swaps them if needed, etc., until it has gone through the whole list. Call this one "sweep" through the list. After the first sweep, the largest number is at the end, so the second sweep (if needed) only needs to work with the first $n - 1$ positions. Similarly, the third sweep (if needed) only needs to work with the first $n - 2$ positions, etc. Sweeps are performed until $n - 1$ sweeps have been completed or there is a swapless sweep.

For example, if the initial list is 53241 (omitting commas), then the following 4 sweeps are performed to sort the list, with a total of 10 comparisons:

$$\begin{aligned} 53241 &\rightarrow 35241 \rightarrow 32541 \rightarrow 32451 \rightarrow 32415 \\ 32415 &\rightarrow 23415 \rightarrow 23415 \rightarrow 23145 \\ 23145 &\rightarrow 23145 \rightarrow 21345 \\ 21345 &\rightarrow 12345 \end{aligned}$$

- (a) An inversion is a pair of numbers that are out of order (e.g., 12345 has no inversions, while 53241 has 8 inversions). Find the expected number of inversions in the original list.
- (b) Show that the expected number of comparisons is between $\frac{1}{2} \binom{n}{2}$ and $\binom{n}{2}$.

Hint: For one bound, think about how many comparisons are made if $n - 1$ sweeps are done; for the other bound, use Part (a).

56. A certain basketball player practices shooting free throws over and over again. The shots are independent, with probability p of success.

(a) In n shots, what is the expected number of streaks of 7 consecutive successful shots? (Note that, for example, 9 in a row counts as 3 streaks.)

(b) Now suppose that the player keeps shooting until making 7 shots in a row for the first time. Let X be the number of shots taken. Show that $E(X) \leq 7/p^7$.

Hint: Consider the first 7 trials as a block, then the next 7 as a block, etc.

57. (S) An urn contains red, green, and blue balls. Balls are chosen randomly with replacement (each time, the color is noted and then the ball is put back). Let r, g, b be the probabilities of drawing a red, green, blue ball, respectively ($r + g + b = 1$).

(a) Find the expected number of balls chosen before obtaining the first red ball, not including the red ball itself.

(b) Find the expected number of different colors of balls obtained before getting the first red ball.

(c) Find the probability that at least 2 of n balls drawn are red, given that at least 1 is red.

58. (S) Job candidates $C_1, C_2, \dots$ are interviewed one by one, and the interviewer compares them and keeps an updated list of rankings (if n candidates have been interviewed so far, this is a list of the n candidates, from best to worst). Assume that there is no limit on the number of candidates available, that for any n the candidates $C_1, C_2, \dots, C_n$ are equally likely to arrive in any order, and that there are no ties in the rankings given by the interview.

Let X be the index of the first candidate to come along who ranks as better than the very first candidate C_1 (so C_X is better than C_1 , but the candidates after 1 but prior to X (if any) are worse than C_1). For example, if C_2 and C_3 are worse than C_1 but C_4 is better than C_1 , then $X = 4$. All $4!$ orderings of the first 4 candidates are equally likely, so it could have happened that the first candidate was the best out of the first 4 candidates, in which case $X > 4$.

What is $E(X)$ (which is a measure of how long, on average, the interviewer needs to wait to find someone better than the very first candidate)?

Hint: Find $P(X > n)$ by interpreting what $X > n$ says about how C_1 compares with other candidates, and then apply the result of Theorem 4.4.8.

59. People are arriving at a party one at a time. While waiting for more people to arrive they entertain themselves by comparing their birthdays. Let X be the number of people needed to obtain a birthday match, i.e., before person X arrives there are no two people with the same birthday, but when person X arrives there is a match.

Assume for this problem that there are 365 days in a year, all equally likely. By the result of the birthday problem from Chapter 1, for 23 people there is a 50.7% chance of a birthday match (and for 22 people there is a less than 50% chance). But this has to do with the median of X (defined below); we also want to know the mean of X , and in this problem we will find it, and see how it compares with 23.

(a) A median of a random variable Y is a value m for which $P(Y \leq m) \geq 1/2$ and $P(Y \geq m) \geq 1/2$ (this is also called a median of the distribution of Y ; note that the notion is completely determined by the CDF of Y). Every distribution has a median, but for some distributions it is not unique. Show that 23 is the unique median of X .

(b) Show that $X = I_1 + I_2 + \dots + I_{366}$, where I_j is the indicator r.v. for the event $X \geq j$. Then find $E(X)$ in terms of p_j 's defined by $p_1 = p_2 = 1$ and for $3 \leq j \leq 366$,

$$p_j = \left(1 - \frac{1}{365}\right) \left(1 - \frac{2}{365}\right) \dots \left(1 - \frac{j-2}{365}\right).$$

(c) Compute $E(X)$ numerically. In R, the pithy command `cumprod(1-(0:364)/365)` pro-

duces the vector $(p_2, \dots, p_{366})$.

(d) Find the variance of X , both in terms of the p_j 's and numerically.

Hint: What is I_i^2 , and what is $I_i I_j$ for $i < j$? Use this to simplify the expansion

$$X^2 = I_1^2 + \dots + I_{366}^2 + 2 \sum_{j=2}^{366} \sum_{i=1}^{j-1} I_i I_j$$

Note: In addition to being an entertaining game for parties, the birthday problem has many applications in computer science, such as in a method called the birthday attack in cryptography. It can be shown that if there are n days in a year and n is large, then

$$E(X) \approx \sqrt{\frac{\pi n}{2}} + \frac{2}{3}$$

60. Elk dwell in a certain forest. There are N elk, of which a simple random sample of size n is captured and tagged (so all $\binom{N}{n}$ sets of n elk are equally likely). The captured elk are returned to the population, and then a new sample is drawn. This is an important method that is widely used in ecology, known as capture-recapture. If the new sample is also a simple random sample, with some fixed size, then the number of tagged elk in the new sample is Hypergeometric.

For this problem, assume that instead of having a fixed sample size, elk are sampled one by one without replacement until m tagged elk have been recaptured, where m is specified in advance (of course, assume that $1 \leq m \leq n \leq N$). An advantage of this sampling method is that it can be used to avoid ending up with a very small number of tagged elk (maybe even zero), which would be problematic in many applications of capture-recapture. A disadvantage is not knowing how large the sample will be.

- (a) Find the PMFs of the number of untagged elk in the new sample (call this X) and of the total number of elk in the new sample (call this Y).
- (b) Find the expected sample size EY using symmetry, linearity, and indicator r.v.s.
- (c) Suppose that m, n, N are such that EY is an integer. If the sampling is done with a fixed sample size equal to EY rather than sampling until exactly m tagged elk are obtained, find the expected number of tagged elk in the sample. Is it less than m , equal to m , or greater than m (for $n < N$)?

Lotus

61. (S) For $X \sim \text{Pois}(\lambda)$, find $E(X!)$ (the average factorial of X), if it is finite.
62. For $X \sim \text{Pois}(\lambda)$, find $E(2^X)$, if it is finite.
63. For $X \sim \text{Geom}(p)$, find $E(2^X)$ (if it is finite) and $E(2^{-X})$ (if it is finite). For each, make sure to clearly state what the values of p are for which it is finite.
64. (S) Let $X \sim \text{Geom}(p)$ and let t be a constant. Find $E(e^{tX})$, as a function of t (this is known as the moment generating function; we will see in Chapter 6 how this function is useful).
65. (S) The number of fish in a certain lake is a $\text{Pois}(\lambda)$ random variable. Worried that there might be no fish at all, a statistician adds one fish to the lake. Let Y be the resulting number of fish (so Y is 1 plus a $\text{Pois}(\lambda)$ random variable).
- (a) Find $E(Y^2)$.
- (b) Find $E(1/Y)$.

66. (S) Let X be a $\text{Pois}(\lambda)$ random variable, where λ is fixed but unknown. Let $\theta = e^{-3\lambda}$, and suppose that we are interested in estimating θ based on the data. Since X is what we observe, our estimator is a function of X , call it $g(X)$. The bias of the estimator $g(X)$ is defined to be $E(g(X)) - \theta$, i.e., how far off the estimate is on average; the estimator is unbiased if its bias is 0.
- (a) For estimating λ , the r.v. X itself is an unbiased estimator. Compute the bias of the estimator $T = e^{-3X}$. Is it unbiased for estimating θ ?
- (b) Show that $g(X) = (-2)^X$ is an unbiased estimator for θ . (In fact, it turns out to be the only unbiased estimator for θ .)
- (c) Explain intuitively why $g(X)$ is a silly choice for estimating θ , despite (b), and show how to improve it by finding an estimator $h(X)$ for θ that is always at least as good as $g(X)$ and sometimes strictly better than $g(X)$. That is,

$$|h(X) - \theta| \leq |g(X) - \theta|,$$

with the inequality sometimes strict.

Poisson Approximation

67. (S) Law school courses often have assigned seating to facilitate the Socratic method. Suppose that there are 100 first-year law students, and each takes the same two courses: Torts and Contracts. Both are held in the same lecture hall (which has 100 seats), and the seating is uniformly random and independent for the two courses.
- (a) Find the probability that no one has the same seat for both courses (exactly; you should leave your answer as a sum).
- (b) Find a simple but accurate approximation to the probability that no one has the same seat for both courses.
- (c) Find a simple but accurate approximation to the probability that at least two students have the same seat for both courses.
68. (S) A group of n people play "Secret Santa" as follows: each puts their name on a slip of paper in a hat, picks a name randomly from the hat (without replacement), and then buys a gift for that person. Unfortunately, they overlook the possibility of drawing one's own name, so some may have to buy gifts for themselves (on the bright side, some may like self-selected gifts better). Assume $n \geq 2$.
- (a) Find the expected value of the number X of people who pick their own names.
- (b) Find the expected number of pairs of people, A and B , such that A picks B 's name and B picks A 's name (where $A \neq B$ and order doesn't matter).
- (c) What is the approximate distribution of X if n is large (specify the parameter value or values)? What does $P(X = 0)$ converge to as $n \rightarrow \infty$?
69. A survey is being conducted in a city with a million (10^6) people. A sample of size 1000 is collected by choosing people in the city at random, with replacement and with equal probabilities for everyone in the city. Find a simple, accurate approximation to the probability that at least one person will get chosen more than once (in contrast, Exercise 26 from Chapter 1 asks for an exact answer).
- Hint: Indicator r.v.s are useful here, but creating 1 indicator for each of the million people is not recommended since it leads to a messy calculation. Feel free to use the fact that $999 \approx 1000$.
70. (S) Ten million people enter a certain lottery. For each person, the chance of winning is one in ten million, independently.

- (a) Find a simple, good approximation for the PMF of the number of people who win the lottery.
- (b) Congratulations! You won the lottery. However, there may be other winners. Assume now that the number of winners other than you is $W \sim \text{Pois}(1)$, and that if there is more than one winner, then the prize is awarded to one randomly chosen winner. Given this information, find the probability that you win the prize (simplify).
71. In a group of 90 people, find a simple, good approximation for the probability that there is at least one pair of people such that they share a birthday and their biological mothers share a birthday. Assume that no one among the 90 people is the biological mother of another one of the 90 people, nor do two of the 90 people have the same biological mother. Express your answer as a fully simplified fraction in the form a/b , where a and b are positive integers and $b \leq 100$.
Make the usual assumptions as in the birthday problem. To simplify the calculation, you can use the approximations $365 \approx 360$ and $89 \approx 90$, and the fact that $e^x \approx 1 + x$ for $x \approx 0$.
72. Use Poisson approximations to investigate the following types of coincidences. The usual assumptions of the birthday problem apply.
- (a) How many people are needed to have a 50% chance that at least one of them has the same birthday as you?
- (b) How many people are needed to have a 50% chance that there is at least one pair of people who not only were born on the same day of the year, but also were born at the same hour (e.g., two people born between 2 pm and 3 pm are considered to have been born at the same hour)?
- (c) Considering that only $1/24$ of pairs of people born on the same day were born at the same hour, why isn't the answer to (b) approximately $24 \cdot 23$?
- (d) With 100 people, there is a 64% chance that there is at least one set of 3 people with the same birthday (according to R , using `pbirthday(100, classes = 365, coincident = 3)` to compute it). Provide two different Poisson approximations for this value, one based on creating an indicator r.v. for each triplet of people, and the other based on creating an indicator r.v. for each day of the year. Which is more accurate?
73. A chess tournament has 100 players. In the first round, they are randomly paired to determine who plays whom (so 50 games are played). In the second round, they are again randomly paired, independently of the first round. In both rounds, all possible pairings are equally likely. Let X be the number of people who play against the same opponent twice.
- (a) Find the expected value of X .
- (b) Explain why X is not approximately Poisson.
- (c) Find good approximations to $P(X = 0)$ and $P(X = 2)$, by thinking about games in the second round such that the same pair played each other in the first round.

***existence**

74. (S) Each of 111 people names their 5 favorite movies out of a list of 11 movies.
- (a) Alice and Bob are 2 of the 111 people. Assume for this part only that Alice's 5 favorite movies out of the 11 are random, with all sets of 5 equally likely, and likewise for Bob, independently. Find the expected number of movies in common to Alice's and Bob's lists of favorite movies.
- (b) Show that there are 2 movies such that at least 21 of the people name both of these movies as favorites.

75. (S) The circumference of a circle is colored with red and blue ink such that $2/3$ of the circumference is red and $1/3$ is blue. Prove that no matter how complicated the coloring scheme is, there is a way to inscribe a square in the circle such that at least three of the four corners of the square touch red ink.
76. (S) A hundred students have taken an exam consisting of 8 problems, and for each problem at least 65 of the students got the right answer. Show that there exist two students who collectively got everything right, in the sense that for each problem, at least one of the two got it right.
77. (S) Ten points in the plane are designated. You have ten circular coins (of the same radius). Show that you can position the coins in the plane (without stacking them) so that all ten points are covered.

Hint: Consider a honeycomb tiling of the plane (this is a way to divide the plane into hexagons). You can use the fact from geometry that if a circle is inscribed in a hexagon then the ratio of the area of the circle to the area of the hexagon is $\frac{\pi}{2\sqrt{3}} > 0.9$.

78. (S) Let S be a set of binary strings $a_1 \dots a_n$ of length n (where juxtaposition means concatenation). We call S k -complete if for any indices $1 \leq i_1 < \dots < i_k \leq n$ and any binary string $b_1 \dots b_k$ of length k , there is a string $s_1 \dots s_n$ in S such that $s_{i_1} s_{i_2} \dots s_{i_k} = b_1 b_2 \dots b_k$. For example, for $n = 3$, the set $S = \{001, 010, 011, 100, 101, 110\}$ is 2-complete since all 4 patterns of 0's and 1's of length 2 can be found in any 2 positions. Show that if $\binom{n}{k} 2^k (1 - 2^{-k})^m < 1$, then there exists a k -complete set of size at most m .

Mixed Practice

79. A hacker is trying to break into a password-protected website by randomly trying to guess the password. Let m be the number of possible passwords.
- (a) Suppose for this part that the hacker makes random guesses (with equal probability), with replacement. Find the average number of guesses it will take until the hacker guesses the correct password (including the successful guess).
- (b) Now suppose that the hacker guesses randomly, without replacement. Find the average number of guesses it will take until the hacker guesses the correct password (including the successful guess).
- Hint: Use symmetry.
- (c) Show that the answer to (a) is greater than the answer to (b) (except in the degenerate case $m = 1$), and explain why this makes sense intuitively.
- (d) Now suppose that the website locks out any user after n incorrect password attempts, so the hacker can guess at most n times. Find the PMF of the number of guesses that the hacker makes, both for the case of sampling with replacement and for the case of sampling without replacement.
80. A fair 20-sided die is rolled repeatedly, until a gambler decides to stop. The gambler receives the amount shown on the die when the gambler stops. The gambler decides in advance to roll the die until a value of m or greater is obtained, and then stop (where m is a fixed integer with $1 \leq m \leq 20$).
- (a) What is the expected number of rolls (simplify)?
- (b) What is the expected square root of the number of rolls (as a sum)?
81. (S) A group of 360 people is going to be split into 120 teams of 3 (where the order of teams and the order within a team don't matter).
- (a) How many ways are there to do this?

- (b) The group consists of 180 married couples. A random split into teams of 3 is chosen, with all possible splits equally likely. Find the expected number of teams containing married couples.
82. (S) The gambler de Méré asked Pascal whether it is more likely to get at least one six in 4 rolls of a die, or to get at least one double-six in 24 rolls of a pair of dice. Continuing this pattern, suppose that a group of n fair dice is rolled $4 \cdot 6^{n-1}$ times.
- (a) Find the expected number of times that "all sixes" is achieved (i.e., how often among the $4 \cdot 6^{n-1}$ rolls it happens that all n dice land 6 simultaneously).
- (b) Give a simple but accurate approximation of the probability of having at least one occurrence of "all sixes", for n large (in terms of e but not n).
- (c) de Méré finds it tedious to re-roll so many dice. So after one normal roll of the n dice, in going from one roll to the next, with probability $6/7$ he leaves the dice in the same configuration and with probability $1/7$ he re-rolls. For example, if $n = 3$ and the 7th roll is $(3, 1, 4)$, then $6/7$ of the time the 8th roll remains $(3, 1, 4)$ and $1/7$ of the time the 8th roll is a new random outcome. Does the expected number of times that "all sixes" is achieved stay the same, increase, or decrease (compared with (a))? Give a short but clear explanation.
83. (S) Five people have just won a \$100 prize, and are deciding how to divide the \$100 up between them. Assume that whole dollars are used, not cents. Also, for example, giving \$50 to the first person and \$10 to the second is different from vice versa.
- (a) How many ways are there to divide up the \$100, such that each gets at least \$10?
- (b) Assume that the \$100 is randomly divided up, with all of the possible allocations counted in (a) equally likely. Find the expected amount of money that the first person receives.
- (c) Let A_j be the event that the j th person receives more than the first person (for $2 \leq j \leq 5$), when the \$100 is randomly allocated as in (b). Are A_2 and A_3 independent?
84. (S) Joe's iPod has 500 different songs, consisting of 50 albums of 10 songs each. He listens to 11 random songs on his iPod, with all songs equally likely and chosen independently (so repetitions may occur).
- (a) What is the PMF of how many of the 11 songs are from his favorite album?
- (b) What is the probability that there are 2 (or more) songs from the same album among the 11 songs he listens to?
- (c) A pair of songs is a match if they are from the same album. If, say, the 1st, 3rd, and 7th songs are all from the same album, this counts as 3 matches. Among the 11 songs he listens to, how many matches are there on average?
85. (S) Each day that the Mass Cash lottery is run in Massachusetts, 5 of the integers from 1 to 35 are chosen (randomly and without replacement).
- (a) When playing this lottery, find the probability of guessing exactly 3 numbers right, given that you guess at least 1 of the numbers right.
- (b) Find an exact expression for the expected number of days needed so that all of the $\binom{35}{5}$ possible lottery outcomes will have occurred.
- (c) Approximate the probability that after 50 days of the lottery, every number from 1 to 35 has been picked at least once.
86. A certain country has three political parties, denoted by A, B, and C. Each adult in the country is a member of exactly one of the three parties. There are n adults in the country, consisting of n_A members of party A, n_B members of party B, and n_C members of party C, where n_A, n_B, n_C are positive integers with $n_A + n_B + n_C = n$.

A simple random sample of size m is chosen from the adults in the country (the sampling is done without replacement, and all possible samples of size m are equally likely). Let X be the number of members of party A in the sample, Y be the number of members of party B in the sample, and Z be the number of members of party C in the sample.

(a) Find $P(X = x, Y = y, Z = z)$, for x, y, z nonnegative integers with $x + y + z = m$.

(b) Find $E(X)$.

(c) Find $\text{Var}(X)$, and briefly explain why your answer makes sense in the extreme cases $m = 1$ and $m = n$.

87. The U.S. Senate consists of 100 senators, with 2 from each of the 50 states. There are d Democrats in the Senate. A committee of size c is formed, by picking a random set of senators such that all sets of size c are equally likely.

(a) Find the expected number of Democrats on the committee.

(b) Find the expected number of states represented on the committee (by at least one senator).

(c) Find the expected number of states such that both of the state's senators are on the committee.

(d) Each state has a junior senator and a senior senator (based on which of them has served longer). A committee of size 20 is formed randomly, with all sets of 20 senators equally likely. Find the distribution of the number of junior senators on the committee, and the expected number of junior senators on the committee.

(e) For the committee from (d), find the expected number of states such that both senators from that state are on the committee.

88. A certain college has g good courses and b bad courses, where g and b are positive integers. Alice, who is hoping to find a good course, randomly shops courses one at a time (without replacement) until she finds a good course.

(a) Find the expected number of bad courses that Alice shops before finding a good course (as a simple expression in terms of g and b).

(b) Should the answer to (a) be less than, equal to, or greater than b/g ? Explain this using properties of the Geometric distribution.

89. A DNA sequence can be represented as a sequence of letters, where the alphabet has 4 letters: A, C, T, G. Suppose such a sequence is generated randomly, where the letters are independent and the probabilities of A, C, T, G are p_A, p_C, p_T, p_G , respectively.

(a) In a DNA sequence of length 115, what is the variance of the number of occurrences of the letter C?

(b) In a DNA sequence of length 115, what is the expected number of occurrences of the expression CATCAT? Note that, for example, the expression CATCATCAT counts as 2 occurrences.

(c) In a DNA sequence of length 6, what is the probability that the expression CAT occurs at least once?

90. Alice is conducting a survey in a town with population size 1000. She selects a simple random sample of size 100 (i.e., sampling without replacement, such that all samples of size 100 are equally likely). Bob is also conducting a survey in this town. Bob selects a simple random sample of size 20, independent of Alice's sample. Let A be the set of people in Alice's sample and B be the set of people in Bob's sample.

(a) Find the expected number of people in $A \cap B$.

(b) Find the expected number of people in $A \cup B$.

(c) The 1000 people consist of 500 married couples. Find the expected number of couples such that both members of the couple are in Bob's sample.

91. The Wilcoxon rank sum test is a widely used procedure for assessing whether two groups of observations come from the same distribution. Let group 1 consist of i.i.d. $X_1, \dots, X_m$ with CDF F and group 2 consist of i.i.d. $Y_1, \dots, Y_n$ with CDF G , with all of these r.v.s independent. Assume that the probability of 2 of the observations being equal is 0 (this will be true if the

distributions are continuous).

After the $m + n$ observations are obtained, they are listed in increasing order, and each is assigned a rank between 1 and $m + n$: the smallest has rank 1, the second smallest has rank 2, etc. Let R_j be the rank of X_j among all the observations for $1 \leq j \leq m$, and let $R = \sum_{j=1}^m R_j$ be the sum of the ranks for group 1.

Intuitively, the Wilcoxon rank sum test is based on the idea that a very large value of R is evidence that observations from group 1 are usually larger than observations from group 2 (and vice versa if R is very small). But how large is "very large" and how small is "very small"? Answering this precisely requires studying the distribution of the test statistic R .

(a) The null hypothesis in this setting is that $F = G$. Show that if the null hypothesis is true, then $E(R) = m(m + n + 1)/2$.

(b) The power of a test is an important measure of how good the test is about saying to reject the null hypothesis if the null hypothesis is false. To study the power of the Wilcoxon rank sum test, we need to study the distribution of R in general. So for this part, we do not assume $F = G$. Let $p = P(X_1 > Y_1)$. Find $E(R)$ in terms of m, n, p .

Hint: Write R_j in terms of indicator r.v.s for X_j being greater than various other r.v.s.

92. The legendary Caltech physicist Richard Feynman and two editors of The Feynman Lectures on Physics (Michael Gottlieb and Ralph Leighton) posed the following problem about how to decide what to order at a restaurant. You plan to eat m meals at a certain restaurant, where you have never eaten before. Each time, you will order one dish.

The restaurant has n dishes on the menu, with $n \geq m$. Assume that if you had tried all the dishes, you would have a definite ranking of them from 1 (your least favorite) to n (your favorite). If you knew which your favorite was, you would be happy to order it always (you never get tired of it).

Before you've eaten at the restaurant, this ranking is completely unknown to you. After you've tried some dishes, you can rank those dishes amongst themselves, but don't know how they compare with the dishes you haven't yet tried. There is thus an exploration/exploitation tradeoff: should you try new dishes, or should you order your favorite among the dishes you have tried before?

A natural strategy is to have two phases in your series of visits to the restaurant: an exploration phase, where you try different dishes each time, and an exploitation phase, where you always order the best dish you obtained in the exploration phase. Let k be the length of the exploration phase (so $m - k$ is the length of the exploitation phase).

Your goal is to maximize the expected sum of the ranks of the dishes you eat there (the rank of a dish is the "true" rank from 1 to n that you would give that dish if you could try all the dishes). Show that the optimal choice is

$$k = \sqrt{2(m+1)} - 1,$$

or this rounded up or down to an integer if needed. Do this in the following steps:

- (a) Let X be the rank of the best dish that you find in the exploration phase. Find the expected sum of the ranks of all the dishes you eat (including both phases), in terms of k, n , and $E(X)$.
- (b) Find the PMF of X , as a simple expression in terms of binomial coefficients.
- (c) Show that

$$E(X) = \frac{k(n+1)}{k+1}$$

Hint: Use Example 1.5.2 (about the team captain) and Exercise 20 from Chapter 1 (about the hockey stick identity).

- (d) Use calculus to find the optimal value of k .

4.5 5 Continuous Random Variables

So far we have been working with discrete random variables, whose possible values can be written down as a list. In this chapter we will discuss continuous r.v.s, which can take on any real value in an interval (possibly of infinite length, such as $(0, \infty)$ or the entire real line). First we'll look at properties of continuous r.v.s in general. Then we'll introduce three famous continuous distributions—the Uniform, Normal, and Exponential—which, in addition to having important stories in their own right, serve as building blocks for many other useful continuous distributions.

4.5.1 5.1 Probability Density Functions




FIGURE 5.1

Discrete vs. continuous r.v.s. Left: The CDF of a discrete r.v. has jumps at each point in the support. Right: The CDF of a continuous r.v. increases smoothly.

Recall that for a discrete r.v., the CDF jumps at every point in the support, and is flat everywhere else. In contrast, for a continuous r.v. the CDF increases smoothly; see Figure 5.1 for a comparison of discrete vs. continuous CDFs.

Definition 5.1.1 (Continuous r.v.). An r.v. has a continuous distribution if its

CDF is differentiable. We also allow there to be endpoints (or finitely many points) where the CDF is continuous but not differentiable, as long as the CDF is differentiable everywhere else. A continuous random variable is a random variable with a continuous distribution.

For discrete r.v.s, the CDF is awkward to work with because of its jumpiness, and its derivative is almost useless since it's undefined at the jumps and 0 everywhere else. But for continuous r.v.s, the CDF is often convenient to work with, and its derivative is a very useful function, called the probability density function.

Definition 5.1.2 (Probability density function). For a continuous r.v. X with CDF F , the probability density function (PDF) of X is the derivative f of the CDF, given by $f(x) = F'(x)$. The support of X , and of its distribution, is the set of all x where $f(x) > 0$.

An important way in which continuous r.v.s differ from discrete r.v.s is that for a continuous r.v. X , $P(X = x) = 0$ for all x . This is because $P(X = x)$ is the height of a jump in the CDF at x , but the CDF of X has no jumps! Since the PMF of a continuous r.v. would just be 0 everywhere, we work with a PDF instead.

The PDF is analogous to the PMF in many ways, but there is a key difference: for a PDF f , the quantity $f(x)$ is not a probability, and in fact it is possible to have $f(x) > 1$ for some values of x . To obtain a probability, we need to integrate the PDF. The fundamental theorem of calculus tells us how to get from the PDF back to the CDF.

Proposition 5.1.3 (PDF to CDF). Let X be a continuous r.v. with PDF f . Then the CDF of X is given by

$$F(x) = \int_{-\infty}^x f(t)dt$$

Proof. By definition of PDF, F is an antiderivative of f . So by the fundamental theorem of calculus,

$$\int_{-\infty}^x f(t)dt = F(x) - F(-\infty) = F(x)$$

The above result is analogous to how we obtained the value of a discrete CDF at x by summing the PMF over all values less than or equal to x ; here we integrate the PDF over all values up to x , so the CDF is the accumulated area under the PDF. Since we can freely convert between the PDF and the CDF using the inverse operations of integration and differentiation, both the PDF and CDF carry complete information about the distribution of a continuous r.v.

Since the PDF determines the distribution, we should be able to use it to find the probability of X falling into an interval (a, b) . A handy fact is that we can include or exclude the endpoints as we wish without altering the probability, since the endpoints have probability 0 :

$$P(a < X < b) = P(a < X \leq b) = P(a \leq X < b) = P(a \leq X \leq b).$$

5.1.4 (Including or excluding endpoints). We can be carefree about including or excluding endpoints as above for continuous r.v.s, but we must not be careless about this for discrete r.v.s.

By definition of CDF and the fundamental theorem of calculus,

$$P(a < X \leq b) = F(b) - F(a) = \int_a^b f(x)dx$$

Therefore, to find the probability of X falling in the interval $(a, b]$ (or (a, b) , $[a, b)$, or $[a, b]$) using the PDF, we simply integrate the PDF from a to b . In general, for an arbitrary region $A \subseteq \mathbb{R}$,

$$P(X \in A) = \int_A f(x)dx$$

In summary:

To get a desired probability, integrate the PDF over the appropriate range.

Just as a valid PMF must be nonnegative and sum to 1 , a valid PDF must be nonnegative and integrate to 1 .

Theorem 5.1.5 (Valid PDFs). The PDF f of a continuous r.v. must satisfy the following two criteria:

- Nonnegative: $f(x) \geq 0$;
- Integrates to 1: $\int_{-\infty}^{\infty} f(x)dx = 1$.

Proof. The first criterion is true because probability is nonnegative; if $f(x_0)$ were negative, then we could integrate over a tiny region around x_0 and get a negative probability. Alternatively, note that the PDF at x_0 is the slope of the CDF at x_0 , so $f(x_0) < 0$ would imply that the CDF is decreasing at x_0 , which is not allowed. The second criterion is true since $\int_{-\infty}^{\infty} f(x)dx$ is the probability of X falling somewhere on the real line, which is 1 .

Conversely, any such function f is the PDF of some r.v. This is because if f satisfies these properties, we can integrate it as in Proposition 5.1.3 to get a function F satisfying the

properties of a CDF. Then a version of Universality of the Uniform, the main concept in Section 5.3, can be used to create an r.v. with CDF F .

Now let's look at some specific examples of PDFs. The two distributions in the following examples are named the Logistic and Rayleigh distributions, but we won't discuss their stories here; their appearance is intended mainly as a way of getting comfortable with PDFs.

Example 5.1.6 (Logistic). The Logistic distribution has CDF

$$F(x) = \frac{e^x}{1 + e^x}, \quad x \in \mathbb{R}$$

To get the PDF, we differentiate the CDF, which gives

$$f(x) = \frac{e^x}{(1 + e^x)^2}, \quad x \in \mathbb{R}$$

Let $X \sim \text{Logistic}$. To find $P(-2 < X < 2)$, integrate the PDF from -2 to 2 :

$$P(-2 < X < 2) = \int_{-2}^2 \frac{e^x}{(1 + e^x)^2} dx = F(2) - F(-2) \approx 0.76.$$

The integral was easy to evaluate since we already knew that F was an antiderivative for f , and we had a nice expression for F . Otherwise, we could have made the substitution $u = 1 + e^x$, so $du = e^x dx$, giving

$$\int_{-2}^2 \frac{e^x}{(1 + e^x)^2} dx = \int_{1+e^{-2}}^{1+e^2} \frac{1}{u^2} du = \left(-\frac{1}{u} \right) \Big|_{1+e^{-2}}^{1+e^2} \approx 0.76.$$

Figure 5.2 shows the Logistic PDF (left) and CDF (right). On the PDF, the probability $P(-2 < X < 2)$ is represented by the shaded area; on the CDF, it is represented by the height of the curly brace. You can check that the properties of a valid PDF and CDF are satisfied. $\square$




FIGURE 5.2

Logistic PDF and CDF. The probability $P(-2 < X < 2)$ is indicated by the shaded area under the PDF and the height of the curly brace on the CDF.

Example 5.1.7 (Rayleigh). The Rayleigh distribution has CDF

$$F(x) = 1 - e^{-x^2/2}, \quad x > 0.$$

To get the PDF, we differentiate the CDF, which gives

$$f(x) = xe^{-x^2/2}, \quad x > 0.$$

For $x \leq 0$, both the CDF and the PDF are equal to 0 .

Let $X \sim \text{Rayleigh}$. To find $P(X > 2)$, we need to integrate the PDF from 2 to ∞ . We can do

that by making the substitution $u = -x^2/2$, but since we already have the CDF in a nice form we know the integral is $F(\infty) - F(2) = 1 - F(2)$:

$$P(X > 2) = \int_2^\infty x e^{-x^2/2} dx = 1 - F(2) \approx 0.14$$

The Rayleigh PDF and CDF are plotted in Figure 5.3. Again, probability is represented by a shaded area on the PDF and a vertical height on the CDF. $\square$




FIGURE 5.3

Rayleigh PDF and CDF. The probability $P(X > 2)$ is indicated by the shaded area under the PDF and the height of the curly brace on the CDF.

Although the height of a PDF at x does not represent a probability, it is closely related to the probability of falling into a tiny interval around x , as the following intuition explains.

Intuition 5.1.8 (Units). Let F be the CDF and f be the PDF of a continuous r.v. X . As mentioned earlier, $f(x)$ is not a probability; for example, we could have $f(3) > 1$, and we know $P(X = 3) = 0$. But thinking about the probability of X being very close to 3 gives us a way to interpret $f(3)$. Specifically, the probability of X being in a tiny interval of length ϵ , centered at 3, will essentially be $f(3)\epsilon$. This is because

$$P(3 - \epsilon/2 < X < 3 + \epsilon/2) = \int_{3-\epsilon/2}^{3+\epsilon/2} f(x) dx \approx f(3)\epsilon$$

if the interval is so tiny that f is approximately the constant $f(3)$ on that interval. In general, we can think of $f(x)dx$ as the probability of X being in an infinitesimally small interval containing x , of length dx .

In practice, X often has units in some system of measurement, such as units of distance, time, area, or mass. Thinking about the units is not only important in applied problems, but also it often helps in checking that answers make sense.

Suppose for concreteness that X is a length, measured in centimeters (cm). Then $f(x) = dF(x)/dx$ is the probability per cm at x , which explains why $f(x)$ is a probability density. Probability is a dimensionless quantity (a number without physical units), so the units of $f(x)$ are cm^{-1} . Therefore, to be able to get a probability again, we need to multiply $f(x)$ by a length. When we do an integral such as $\int_0^5 f(x) dx$, this is achieved by the often-forgotten dx . $\square$

Intuition 5.1.9 (Simulation). For another intuitive way to think about PDFs, consider the following graphical way to simulate a draw X from a continuous distribution, based on looking at the graph of the PDF. To generate X , choose a uniformly random point under the PDF curve; this means that the probability of any region under the curve is the area of that region. Then let X be the x -coordinate of the random point. This is illustrated in Figure 5.4.




FIGURE 5.4

A complicated PDF (no numbers are shown on the vertical axis since the scale is whatever it needs to be to make the area under the curve 1). To generate an r.v. X with this PDF, choose a uniformly random point (x, y) under the curve and let $X = x$. This method works since, for example, X will be in the interval $[3, 6]$ if and only if the randomly chosen point (x, y) is in the shaded region.

Then X has the desired distribution since, by construction, $P(a \leq X \leq b)$ is the area under the PDF curve between the lines $x = a$ and $x = b$. Thinking about this method helps build intuition for PDFs, by giving us a feel for random variables sampled according to a particular PDF curve. $\square$

The definition of expectation for continuous r.v.s is analogous to the definition for discrete r.v.s: replace the sum with an integral and the PMF with the PDF.

Definition 5.1.10 (Expectation of a continuous r.v.). The expected value (also called the expectation or mean) of a continuous r.v. X with PDF f is

$$E(X) = \int_{-\infty}^{\infty} xf(x)dx$$

As in the discrete case, the expectation of a continuous r.v. may or may not exist. When discussing expectations, it would be very tedious to have to add "(if it exists)" after every mention of an expectation not yet shown to exist, so we will often leave this implicit.

The integral is taken over the entire real line, but if the support of X is not the entire real line we can just integrate over the support. The units in this definition make sense: if X is measured in centimeters, then so is $E(X)$, since $xf(x)dx$ has units of $\text{cm} \cdot \text{cm}^{-1} \cdot \text{cm} = \text{cm}$.

With this definition, the expected value retains its interpretation as a center of mass. As shown in Figure 5.5, using the Rayleigh PDF for illustrative purposes, the expected value is the balancing point of the PDF, just as it was the balancing point of the PMF in the discrete case.




FIGURE 5.5

The expected value of a continuous r.v. is the balancing point of the PDF .

Linearity of expectation holds for continuous r.v.s, as it did for discrete r.v.s (we will show this later in Example 7.2.4). LOTUS also holds for continuous r.v.s, replacing the sum with an integral and the PMF with the PDF:

Theorem 5.1.11 (LOTUS, continuous). If X is a continuous r.v. with PDF f and g is a function from $\mathbb{R}$ to $\mathbb{R}$, then

$$E(g(X)) = \int_{-\infty}^{\infty} g(x)f(x)dx$$

We now have all the tools we need to tackle the named distributions of this chapter, starting with the Uniform distribution.

4.5.2 5.2 Uniform

Intuitively, a Uniform r.v. on the interval (a, b) is a completely random number between a and b . We formalize the notion of "completely random" on an interval by specifying that the PDF should be constant over the interval.

Definition 5.2.1 (Uniform distribution). A continuous r.v. U is said to have the Uniform distribution on the interval (a, b) if its PDF is

$$f(x) = \begin{cases} \frac{1}{b-a} & \text{if } a < x < b \\ 0 & \text{otherwise} \end{cases}$$

We denote this by $U \sim \text{Unif}(a, b)$.

This is a valid PDF because the area under the curve is just the area of a rectangle with width $b - a$ and height $1/(b - a)$. The CDF is the accumulated area under the PDF:

$$F(x) = \begin{cases} 0 & \text{if } x \leq a \\ \frac{x-a}{b-a} & \text{if } a < x < b \\ 1 & \text{if } x \geq b \end{cases}$$

The Uniform distribution that we will most frequently use is the $\text{Unif}(0, 1)$ distribution, also called the standard Uniform. The $\text{Unif}(0, 1)$ PDF and CDF are particularly simple: $f(x) = 1$ and $F(x) = x$ for $0 < x < 1$. Figure 5.6 shows the $\text{Unif}(0, 1)$ PDF and CDF side by side.

For a general $\text{Unif}(a, b)$ distribution, the PDF is constant on (a, b) , and the CDF is ramp-shaped, increasing linearly from 0 to 1 as x ranges from a to b .

For Uniform distributions, probability is proportional to length.

Proposition 5.2.2. Let $U \sim \text{Unif}(a, b)$, and let (c, d) be a subinterval of (a, b) , of length l (so $l = d - c$). Then the probability of U being in (c, d) is proportional to l . For example, a subinterval that is twice as long has twice the probability of containing U , and a subinterval of the same length has the same probability.




FIGURE 5.6
 $\text{Unif}(0,1)$ PDF and CDF .

Proof. Since the PDF of U is the constant $\frac{1}{b-a}$ on (a, b) , the area under the PDF from c to d is $\frac{l}{b-a}$, which is a constant times l . $\square$

The above proposition is a very special property of the Uniform; for any other distribution, there are intervals of the same length that have different probabilities. Even after conditioning on a Uniform r.v. being in a certain subinterval, we still have a Uniform distribution and thus still have probability proportional to length (within that subinterval); we show this below.

Proposition 5.2.3. Let $U \sim \text{Unif}(a, b)$, and let (c, d) be a subinterval of (a, b) . Then the conditional distribution of U given $U \in (c, d)$ is $\text{Unif}(c, d)$.

Proof. For u in (c, d) , the conditional CDF at u is

$$P(U \leq u \mid U \in (c, d)) = \frac{P(U \leq u, c < U < d)}{P(U \in (c, d))} = \frac{P(U \in (c, u])}{P(U \in (c, d))} = \frac{u - c}{d - c}.$$

The conditional CDF is 0 for $u \leq c$ and 1 for $u \geq d$. So the conditional distribution of U is as claimed. $\square$

Example 5.2.4. Let's illustrate the above propositions for $U \sim \text{Unif}(0, 1)$. In this special case, the support has length 1, so probability is length: the probability of U falling into the interval $(0, 0.3)$ is 0.3, as is the probability of falling into $(0.3, 0.6)$, $(0.4, 0.7)$, or any other interval of length 0.3 within $(0, 1)$.

Now suppose that we learn that $U \in (0.4, 0.7)$. Given this information, the conditional distribution of U is $\text{Unif}(0.4, 0.7)$. Then the conditional probability of $U \in (0.4, 0.6)$ is $2/3$, since $(0.4, 0.6)$ provides $2/3$ of the length of $(0.4, 0.7)$. The conditional probability of $U \in (0.6, 0.7)$ is also $1/3$, since we discard the points to the left of 0.4 when conditioning on $U \in (0.4, 0.7)$. $\square$

Next, let's derive the mean and variance of $U \sim \text{Unif}(a, b)$. The expectation is extremely intuitive: the PDF is constant, so its balancing point should be the midpoint of (a, b) . This is exactly what we find by using the definition of expectation for continuous r.v.s:

$$E(U) = \int_a^b x \cdot \frac{1}{b-a} dx = \frac{1}{b-a} \left(\frac{b^2}{2} - \frac{a^2}{2} \right) = \frac{a+b}{2}.$$

For the variance, we first find $E(U^2)$ using the continuous version of LOTUS:

$$E(U^2) = \int_a^b x^2 \frac{1}{b-a} dx = \frac{1}{3} \cdot \frac{b^3 - a^3}{b-a}$$

Then

$$\text{Var}(U) = E(U^2) - (EU)^2 = \frac{1}{3} \cdot \frac{b^3 - a^3}{b-a} - \left(\frac{a+b}{2} \right)^2,$$

which reduces, after factoring $b^3 - a^3 = (b-a)(a^2 + ab + b^2)$ and simplifying, to

$$\text{Var}(U) = \frac{(b-a)^2}{12}$$

The above derivation isn't terribly painful, but there is an easier path, using a technique that is often useful for continuous distributions. The technique is called location-scale transformation, and it relies on the observation that shifting and scaling a Uniform r.v. produces another Uniform r.v. Shifting is considered a change of location and scaling is a change of scale, hence the term location-scale. For example, if X is Uniform on the interval $(1, 2)$, then $X + 5$ is Uniform on the interval $(6, 7)$, $2X$ is Uniform on the interval $(2, 4)$, and $2X + 5$ is Uniform on $(7, 9)$.

Definition 5.2.5 (Location-scale transformation). Let X be a random variable and $Y = \sigma X + \mu$, where σ and μ are constants with $\sigma > 0$. Then we say that Y has been obtained as a location-scale transformation of X . Here μ controls how the location is changed and σ controls how the scale is changed.

5.2.6. In a location-scale transformation, starting with $X \sim \text{Unif}(a, b)$ and transforming it to $Y = cX + d$ where c and d are constants with $c > 0$, Y is a linear function of X and Uniformity is preserved: $Y \sim \text{Unif}(ca + d, cb + d)$. But if Y is defined as a nonlinear transformation of X , then Y will not be Uniform in general. For example, for $X \sim \text{Unif}(a, b)$ with $0 \leq a < b$, the transformed r.v. $Y = X^2$ has support (a^2, b^2) but is not Uniform on that interval. Chapter 8 explores transformations of r.v.s in detail.

In studying Uniform distributions, a useful strategy is to start with an r.v. that has the simplest Uniform distribution, figure things out in the friendly simple case, and then use a location-scale transformation to handle the general case.

Let's see how this works for finding the expectation and variance of the $\text{Unif}(a, b)$ distribution. The location-scale strategy says to start with $U \sim \text{Unif}(0, 1)$. Since the PDF of U is just 1 on the interval $(0, 1)$, it is easy to see that

$$\begin{aligned} E(U) &= \int_0^1 x dx = \frac{1}{2} \\ E(U^2) &= \int_0^1 x^2 dx = \frac{1}{3} \\ \text{Var}(U) &= \frac{1}{3} - \frac{1}{4} = \frac{1}{12} \end{aligned}$$

Now that we know the answers for U , transforming U into a general $\text{Unif}(a, b)$ r.v. takes just two steps. First we change the support from an interval of length 1 to an interval of length $b - a$, so we multiply U by the scaling factor $b - a$ to obtain a $\text{Unif}(0, b - a)$ r.v. Then we shift

everything until the left endpoint of the support is at a . Thus, if $U \sim \text{Unif}(0, 1)$, the random variable

$$\tilde{U} = a + (b - a)U$$

is distributed $\text{Unif}(a, b)$. Now the mean and variance of $\tilde{U}$ follow directly from properties of expectation and variance. By linearity of expectation,

$$E(\tilde{U}) = E(a + (b - a)U) = a + (b - a)E(U) = a + \frac{b - a}{2} = \frac{a + b}{2}.$$

By the fact that additive constants don't affect the variance while multiplicative constants come out squared,

$$\text{Var}(\tilde{U}) = \text{Var}(a + (b - a)U) = \text{Var}((b - a)U) = (b - a)^2 \text{Var}(U) = \frac{(b - a)^2}{12}.$$

These agree with our previous answers.

The technique of location-scale transformation will work for any family of distributions such that shifting and scaling an r.v. whose distribution in the family produces another r.v. whose distribution is in the family. This technique does not apply to families of discrete distributions (with a fixed support) since, for example, shifting or scaling $X \sim \text{Bin}(n, p)$ changes the support and produces an r.v. that is no longer Binomial. A Binomial r.v. must be able to take on all integer values between 0 and some upper bound, but $X + 4$ can't take on any value in $\{0, 1, 2, 3\}$ and $2X$ can only take even values, so neither of these r.v.s has a Binomial distribution.

5.2.7 (Beware of sympathetic magic). When using location-scale transformations, the shifting and scaling should be applied to the random variables themselves, not to their PDFs. To confuse these two would be an instance of sympathetic magic (see 3.7.7), and would result in invalid PDFs. For example, let $U \sim \text{Unif}(0, 1)$, so the PDF f has $f(x) = 1$ on $(0, 1)$ (and $f(x) = 0$ elsewhere). Then $3U + 1 \sim \text{Unif}(1, 4)$, but $3f + 1$ is the function that equals 4 on $(0, 1)$ and 1 elsewhere, which is not a valid PDF since it does not integrate to 1.

4.5.3 5.3 Universality of the Uniform

In this section, we will discuss a remarkable property of the Uniform distribution: given a $\text{Unif}(0, 1)$ r.v., we can construct an r.v. with any continuous distribution we want. Conversely, given an r.v. with an arbitrary continuous distribution, we can create a $\text{Unif}(0, 1)$ r.v. We call this the universality of the Uniform, because it tells us the Uniform is a universal starting point for building r.v.s with other distributions. Universality of the Uniform also goes by many other names, such as the probability integral transform, inverse transform sampling, the quantile transformation, and even the fundamental theorem of simulation.

To keep the proofs simple, we will state universality of the Uniform for a case where we know the inverse of the desired CDF exists. Similar ideas can be used to simulate a random draw from any desired CDF as a function of a $\text{Unif}(0, 1)$ r.v.

Theorem 5.3.1 (Universality of the Uniform). Let F be a CDF which is a continuous function and strictly increasing on the support of the distribution. This ensures that the inverse function F^{-1} exists, as a function from $(0, 1)$ to $\mathbb{R}$. We then have the following results.

1. Let $U \sim \text{Unif}(0, 1)$ and $X = F^{-1}(U)$. Then X is an r.v. with CDF F .
2. Let X be an r.v. with CDF F . Then $F(X) \sim \text{Unif}(0, 1)$.

Let's make sure we understand what each part of the theorem is saying. The first part says that if we start with $U \sim \text{Unif}(0, 1)$ and a CDF F , then we can create an r.v. whose CDF is

F by plugging U into the inverse CDF F^{-1} . Since F^{-1} is a function (known as the quantile function), U is an r.v., and a function of an r.v. is an r.v., $F^{-1}(U)$ is an r.v.; universality of the Uniform says its CDF is F .

The second part of the theorem goes in the reverse direction, starting from an r.v. X whose CDF is F and then creating a $\text{Unif}(0, 1)$ r.v. Again, F is a function, X is an r.v., and a function of an r.v. is an r.v., so $F(X)$ is an r.v. Since any CDF is between 0 and 1 everywhere, $F(X)$ must take values between 0 and 1. Universality of the Uniform says that the distribution of $F(X)$ is Uniform on $(0, 1)$.

5.3.2. The second part of universality of the Uniform involves plugging a random variable X into its own CDF F . This may seem strangely self-referential, but it makes sense because F is just a function (that satisfies the properties of a valid CDF), and a function of an r.v. is an r.v. There is a potential notational confusion, however: $F(x) = P(X \leq x)$ by definition, but it would be incorrect to say " $F(X) = P(X \leq X) = 1$ ". Rather, we should first find an expression for the CDF as a function of x , then replace x with X to obtain an r.v. For example, if the CDF of X is $F(x) = 1 - e^{-x}$ for $x > 0$, then $F(X) = 1 - e^{-X}$.

Understanding the statement of the theorem is the difficult part; the proof is just a couple of lines for each direction.

Proof.

1. Let $U \sim \text{Unif}(0, 1)$ and $X = F^{-1}(U)$. For all real x ,

$$P(X \leq x) = P(F^{-1}(U) \leq x) = P(U \leq F(x)) = F(x),$$

so the CDF of X is F , as claimed. For the last equality, we used the fact that $P(U \leq u) = u$ for $u \in (0, 1)$.

2. Let X have CDF F , and find the CDF of $Y = F(X)$. Since Y takes values in $(0, 1)$, $P(Y \leq y)$ equals 0 for $y \leq 0$ and equals 1 for $y \geq 1$. For $y \in (0, 1)$,

$$P(Y \leq y) = P(F(X) \leq y) = P(X \leq F^{-1}(y)) = F(F^{-1}(y)) = y.$$

Thus Y has the $\text{Unif}(0, 1)$ CDF.

To gain more insight into what the quantile function F^{-1} and universality of the Uniform mean, let's consider an example that is familiar to millions of students: percentiles on an exam.

Example 5.3.3 (Percentiles). A large number of students take a certain exam, graded on a scale from 0 to 100. Let X be the score of a random student. Continuous distributions are easier to deal with here, so let's approximate the discrete distribution of scores using a continuous distribution. Suppose that X is continuous, with a CDF F that is strictly increasing on $(0, 100)$. In reality, there are only finitely many students and only finitely many possible scores, but a continuous distribution may be a good approximation.

Suppose that the median score on the exam is 60, i.e., half of the students score above 60 and the other half score below 60 (a convenient aspect of assuming a continuous distribution is that we don't need to worry about how many students had scores equal to 60). That is, $F(60) = 1/2$, or, equivalently, $F^{-1}(1/2) = 60$.

If Fred scores a 72 on the exam, then his percentile is the fraction of students who score below a 72. This is $F(72)$, which is some number in $(1/2, 1)$ since 72 is above the median. In general, a student with score x has percentile $F(x)$. Going the other way, if we start with a percentile, say 0.95, then $F^{-1}(0.95)$ is the score that has that percentile. A percentile is also called a quantile, which is why F^{-1} is called the quantile function. The function F converts scores to quantiles, and the function F^{-1} converts quantiles to scores.

The strange operation of plugging X into its own CDF now has a natural interpretation: $F(X)$ is the percentile attained by a random student. It often happens that the distribution of scores on an exam looks very non-Uniform. For example, there is no reason to think that 10% of the scores are between 70 and 80, even though $(70, 80)$ covers 10% of the range of possible scores.

On the other hand, the distribution of percentiles of the students is Uniform: the universality property says that $F(X) \sim \text{Unif}(0, 1)$. For example, 50% of the students have a percentile of at least 0.5. Universality of the Uniform is expressing the fact that 10% of the students have a percentile between 0 and 0.1, 10% have a percentile between 0.1 and 0.2, 10% have a percentile between 0.2 and 0.3, and so on—a fact that is clear from the definition of percentile.

To illustrate universality of the Uniform, we will apply it to the two distributions we encountered in the previous section, the Logistic and Rayleigh.

Example 5.3.4 (Universality with Logistic). The Logistic CDF is

$$F(x) = \frac{e^x}{1 + e^x}, \quad x \in \mathbb{R}$$

Suppose we have $U \sim \text{Unif}(0, 1)$ and wish to generate a Logistic r.v. Part 1 of the universality property says that $F^{-1}(U) \sim \text{Logistic}$, so we first invert the CDF to get F^{-1} :

$$F^{-1}(u) = \log\left(\frac{u}{1-u}\right)$$

Then we plug in U for u :

$$F^{-1}(U) = \log\left(\frac{U}{1-U}\right)$$

Therefore $\log\left(\frac{U}{1-U}\right) \sim \text{Logistic}$.

We can verify directly that $\log\left(\frac{U}{1-U}\right)$ has the required CDF: start from the definition of CDF, do some algebra to isolate U on one side of the inequality, and then use the CDF of the Uniform distribution. Let's work through these calculations once for practice:

$$\begin{aligned} P\left(\log\left(\frac{U}{1-U}\right) \leq x\right) &= P\left(\frac{U}{1-U} \leq e^x\right) \\ &= P(U \leq e^x(1-U)) \\ &= P\left(U \leq \frac{e^x}{1+e^x}\right) \\ &= \frac{e^x}{1+e^x} \end{aligned}$$

which is indeed the Logistic CDF.

We can also use simulation to visualize how universality of the Uniform works. To this end, we generated 1 million $\text{Unif}(0, 1)$ random variables. We then transformed each of these values u into $\log\left(\frac{u}{1-u}\right)$; if the universality of the Uniform is correct, the transformed numbers should follow a Logistic distribution.




FIGURE 5.7

Top: Histogram of 10^6 draws of $U \sim \text{Unif}(0, 1)$, with $\text{Unif}(0, 1)$ PDF for comparison. Bottom: Histogram of 10^6 draws of $\log\left(\frac{U}{1-U}\right)$, with Logistic PDF for comparison.

Figure 5.7 displays a histogram of the realizations of U alongside the $\text{Unif}(0, 1)$ PDF; below that, we have a histogram of the realizations of $\log\left(\frac{U}{1-U}\right)$ next to the Logistic PDF. As we can see, the second histogram looks very much like the Logistic PDF. Thus, by applying F^{-1} , we were able to transform our Uniform draws into Logistic draws, exactly as claimed by the universality of the Uniform.

Conversely, Part 2 of the universality property states that if $X \sim \text{Logistic}$, then

$$F(X) = \frac{e^X}{1 + e^X} \sim \text{Unif}(0, 1).$$

□

Example 5.3.5 (Universality with Rayleigh). The Rayleigh CDF is

$$F(x) = 1 - e^{-x^2/2}, \quad x > 0.$$

The quantile function (the inverse of the CDF) is

$$F^{-1}(u) = \sqrt{-2 \log(1 - u)},$$

so if $U \sim \text{Unif}(0, 1)$, then $F^{-1}(U) = \sqrt{-2 \log(1 - U)} \sim \text{Rayleigh}$.

We again generated 1 million realizations of $U \sim \text{Unif}(0, 1)$ and transformed them




FIGURE 5.8 Top: Histogram of 1 million draws from $U \sim \text{Unif}(0, 1)$, with $\text{Unif}(0, 1)$ PDF for comparison. Bottom: Histogram of 1 million draws from $\sqrt{-2\log(1-U)}$, with Rayleigh PDF for comparison.

to produce 1 million realizations of $\sqrt{-2\log(1-U)}$. As Figure 5.8 shows, the realizations of $\sqrt{-2\log(1-U)}$ look very similar to the Rayleigh PDF, as predicted by the universality of the Uniform.

Conversely, if $X \sim \text{Rayleigh}$, then $F(X) = 1 - e^{-X^2/2} \sim \text{Unif}(0, 1)$. $\square$

Next, let us consider the extent to which universality of the Uniform holds for discrete random variables. The CDF F of a discrete r.v. has jumps and flat regions, so F^{-1} does not exist (in the usual sense). But Part 1 still holds in the sense that given a Uniform random variable, we can construct an r.v. with any discrete distribution we want. The difference is that instead of working with the CDF, which is not invertible, it is more straightforward to work with the PMF.

Suppose we want to use $U \sim \text{Unif}(0, 1)$ to construct a discrete r.v. X with PMF $p_j = P(X = j)$ for $j = 0, 1, 2, \dots, n$. As illustrated in Figure 5.9, we can chop up the interval $(0, 1)$ into pieces of lengths $p_0, p_1, \dots, p_n$. By the properties of a valid PMF, the sum of the p_j 's is 1, so this perfectly divides up the interval, without overshooting or undershooting.

Now define X to be the r.v. which equals 0 if U falls into the p_0 interval, 1 if U falls into the p_1 interval, 2 if U falls into the p_2 interval, and so on. Then X is a discrete




FIGURE 5.9 Given a PMF, chop up the interval $(0, 1)$ into pieces, with lengths given by the PMF values.

r.v. taking on values 0 through n . The probability that $X = j$ is the probability that U falls into the interval of length p_j . But for a $\text{Unif}(0, 1)$ r.v., probability is length, so $P(X = j)$ is precisely p_j , as desired!

The same trick will work for a discrete r.v. that can take on infinitely many values, such as a Poisson; we'll need to chop $(0, 1)$ into infinitely many pieces, but the total length of the pieces is still 1.

We now know how to take an arbitrary PMF and create an r.v. with that PMF. This fulfills our promise from Chapter 3 that any function with the properties given in Theorem 3.2.7 is the PMF of some r.v.

5.3.6. Part 2 of universality of the Uniform, on the other hand, fails for discrete r.v.s. A function of a discrete r.v. is still discrete, so if X is discrete, then $F(X)$ is still discrete. So $F(X)$ doesn't have a Uniform distribution. For example, if $X \sim \text{Bern}(p)$, then $F(X)$ has only two possible values: $F(0) = 1 - p$ and $F(1) = 1$.

The upshot of universality is that we can use a Uniform r.v. U to generate r.v.s from both continuous and discrete distributions: in the continuous case, we can plug U into the inverse CDF, and in the discrete case, we can chop up the unit interval according to the desired PMF. Part 1 of universality of the Uniform is often useful in practice when running simulations (since the software being used may know how to generate Uniform r.v.s but not know how to generate r.v.s with the distribution of interest), though the extent to which it is useful depends on how tractable it is to compute the inverse CDF. Part 2 is important for certain widely used techniques in statistical inference, by providing a transformation that converts an r.v. with an unknown distribution to an r.v. with a known, simple distribution: the Uniform.

Using our analogy of distributions as blueprints and r.v.s as houses, the beauty of the universality property is that the Uniform distribution is a very simple blueprint, and it's easy to create a house from that blueprint; universality of the Uniform then gives us a simple rule for remodeling the Uniform house into a house with any other blueprint, no matter how complicated!

To conclude this section, we give an elegant identity that is often useful for finding the expectation of a nonnegative r.v. The identity also has a neat visual interpretation related to universality of the Uniform, LOTUS, and the relationship between CDFs and quantile functions.

Definition 5.3.7. The survival function of an r.v. X with CDF F is the function G given by $G(x) = 1 - F(x) = P(X > x)$.

Theorem 5.3.8 (Expectation by integrating the survival function). Let X be a nonnegative r.v. Its expectation can be found by integrating its survival function:

$$E(X) = \int_0^\infty P(X > x) dx$$

This result is the continuous analog of Theorem 4.4.8 (note though that it holds for any nonnegative r.v., not just for continuous nonnegative r.v.s). Actuaries sometimes call it the Darth Vader rule, for obscure reasons; statisticians are more likely to refer to it as finding the expectation by integrating the survival function.

Proof. For any number $x \geq 0$, we can write

$$x = \int_0^x dt = \int_0^\infty I(x > t) dt$$

where $I(x > t)$ is 1 if $x \geq t$ and 0 otherwise. So

$$X(s) = \int_0^\infty I(X(s) > t) dt$$

for each s in the sample space. We can write this more compactly as

$$X = \int_0^\infty I(X > t) dt$$

Taking the expectation of both sides and swapping the E with the integral (which can be justified using results from real analysis), we have

$$E(X) = E\left(\int_0^\infty I(X > t) dt\right) = \int_0^\infty E(I(X > t)) dt = \int_0^\infty P(X > t) dt$$

For a visual explanation of this identity, we can graph a CDF and interpret a certain area in two different ways: as the integral of the survival function, and as the integral of the quantile function.

A prototypical CDF of a nonnegative, continuous r.v. with CDF F is shown in Figure 5.10, with the area between the CDF curve and the horizontal line $p = 1$ shaded. This area can be found by integrating $1 - F(x)$, the difference between the line and the curve, from 0 to ∞ .

But another way to find this area is to turn your head sideways and integrate with




FIGURE 5.10

The area above a certain CDF and below the line $p = 1$ is shaded. This area can be interpreted in two ways: as the integral of the survival function, or as the integral of the quantile function.

respect to the vertical axis variable p rather than the horizontal axis variable x . This gives the integral of the quantile function, which, letting $U \sim \text{Unif}(0, 1)$, is

$$\int_0^1 F^{-1}(p) dp = E(F^{-1}(U)) = E(X)$$

by LOTUS and universality of the Uniform. So again we have

$$\int_0^\infty (1 - F(x)) dx = \int_0^1 F^{-1}(p) dp = E(X)$$

4.5.4 5.4 Normal

The Normal distribution is a famous continuous distribution with a bell-shaped PDF. It is extremely widely used in statistics because of a theorem, the central limit theorem, which says that under very weak assumptions, the sum of a large number of i.i.d. random variables has an approximately Normal distribution, regardless of the distribution of the individual r.v.s. This means we can start with independent r.v.s from almost any distribution, discrete or continuous, but once we add up a bunch of them, the distribution of the resulting r.v. looks like a Normal distribution.

The central limit theorem is a topic for Chapter 10, but in the meantime, we'll introduce the properties of the Normal PDF and CDF and derive the expectation and variance of the Normal distribution. To do this, we will again use the strategy of location-scale transformation

by starting with the simplest Normal distribution, the standard Normal, which is centered at 0 and has variance 1 . After deriving the properties of the standard Normal, we'll be able to get to any Normal distribution we want by shifting and scaling.

Definition 5.4.1 (Standard Normal distribution). A continuous r.v. Z is said to have the standard Normal distribution if its PDF φ is given by

$$\varphi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}, \quad -\infty < z < \infty.$$

We write this as $Z \sim \mathcal{N}(0, 1)$ since, as we will show, Z has mean 0 and variance 1. The constant $\frac{1}{\sqrt{2\pi}}$ in front of the PDF may look surprising (why is something with π needed in front of something with e , when there are no circles in sight?), but it's exactly what is needed to make the PDF integrate to 1 . Such constants are called normalizing constants because they normalize the total area under the PDF to 1 . We'll verify soon that this is a valid PDF.

The standard Normal CDF Φ is the accumulated area under the PDF:

$$\Phi(z) = \int_{-\infty}^z \varphi(t) dt = \int_{-\infty}^z \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt$$

Some people, upon seeing the function Φ for the first time, express dismay that it is left in terms of an integral. Unfortunately, we have little choice in the matter: it turns out to be mathematically impossible to find a closed-form expression for the antiderivative of φ , meaning that we cannot express Φ as a finite sum of more familiar functions like polynomials or exponentials. But closed-form or no, it's still a well-defined function: if we give Φ an input z , it returns the accumulated area under the PDF from $-\infty$ up to z .

Notation 5.4.2. We can tell the Normal distribution must be special because the standard Normal PDF and CDF get their own Greek letters. By convention, we use φ for the standard Normal PDF and Φ for the CDF. We will often use Z to denote a standard Normal random variable.

The standard Normal PDF and CDF are plotted in Figure 5.11. The PDF is bellshaped and symmetric about 0 , and the CDF is S -shaped. These have the same general shape as the Logistic PDF and CDF that we saw Example 5.1.6, but the Normal PDF decays to 0 much more quickly.

There are several important symmetry properties that can be deduced from the standard Normal PDF and CDF.

1. Symmetry of PDF: φ satisfies $\varphi(z) = \varphi(-z)$, i.e., φ is an even function.




FIGURE 5.11
Standard Normal PDF φ (left) and CDF Φ (right).

2. Symmetry of tail areas: The area under the PDF curve to the left of -2 , which is $P(Z \leq -2) = \Phi(-2)$ by definition, equals the area to the right of 2 , which is $P(Z \geq 2) = 1 - \Phi(2)$. In general, we have

$$\Phi(z) = 1 - \Phi(-z)$$

for all z . This can be seen visually by looking at the PDF curve, and mathematically by substituting $u = -t$ below and using the fact that PDFs integrate to 1 :

$$\Phi(-z) = \int_{-\infty}^{-z} \varphi(t) dt = \int_z^{\infty} \varphi(u) du = 1 - \int_{-\infty}^z \varphi(u) du = 1 - \Phi(z).$$

3. Symmetry of Z and $-Z$: If $Z \sim \mathcal{N}(0, 1)$, then $-Z \sim \mathcal{N}(0, 1)$ as well. To see this, note that the CDF of $-Z$ is

$$P(-Z \leq z) = P(Z \geq -z) = 1 - \Phi(-z),$$

but that is $\Phi(z)$, according to what we just argued. So $-Z$ has CDF Φ .

We need to prove three key facts about the standard Normal, and then we'll be ready to handle general Normal distributions: we need to show that φ is a valid PDF, that $E(Z) = 0$, and that $\text{Var}(Z) = 1$.

To verify the validity of φ , we'll show that the total area under $e^{-z^2/2}$ is $\sqrt{2\pi}$. However, we can't find the antiderivative of $e^{-z^2/2}$ directly, again because of the annoying fact that the antiderivative isn't expressible in closed form. But this doesn't mean we can't do definite integrals, with some ingenuity.

An amazing trick saves the day here: write down the integral twice. Usually, writing down the same problem repeatedly is more a sign of frustration than a problemsolving strategy. But in this case, it allows a neat conversion to polar coordinates:

$$\begin{aligned} \left(\int_{-\infty}^{\infty} e^{-z^2/2} dz \right) \left(\int_{-\infty}^{\infty} e^{-z^2/2} dz \right) &= \left(\int_{-\infty}^{\infty} e^{-x^2/2} dx \right) \left(\int_{-\infty}^{\infty} e^{-y^2/2} dy \right) \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} e^{-\frac{x^2+y^2}{2}} dx dy \\ &= \int_0^{2\pi} \int_0^{\infty} e^{-r^2/2} r dr d\theta \end{aligned}$$

In the first step, we used the fact that z is just a dummy variable in each integral, so we are allowed to give it a different name (or two different names, one for each integral). The extra r that appears in the final step comes from the Jacobian of the transformation to polar coordinates, as explained in Section A.7.2 of the math appendix. That r is also what saves us from the impossibility of the original integral, since we can now use the substitution $u = r^2/2, du = r dr$. This gives

$$\begin{aligned} \int_0^{2\pi} \int_0^{\infty} e^{-r^2/2} r dr d\theta &= \int_0^{2\pi} \left(\int_0^{\infty} e^{-u} du \right) d\theta \\ &= \int_0^{2\pi} 1 d\theta = 2\pi \end{aligned}$$

Therefore,

$$\int_{-\infty}^{\infty} e^{-z^2/2} dz = \sqrt{2\pi}$$

as we wanted to show.

The expectation of the standard Normal has to be 0, by the symmetry of the PDF; no other balancing point would make sense. We can also see this symmetry by looking at the definition of $E(Z)$:

$$E(Z) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} ze^{-z^2/2} dz$$

and since $g(z) = ze^{-z^2/2}$ is an odd function (see Section A.2.3 of the math appendix for more on even and odd functions), the area under g from $-\infty$ to 0 cancels the area under g from 0 to ∞ . Therefore $E(Z) = 0$. In fact, the same argument shows that $E(Z^n) = 0$ for any odd positive integer n .¹

Getting the mean was easy (one might even say it was EZ), but the variance calculation is a bit more involved. By LOTUS,

$$\begin{aligned} \text{Var}(Z) &= E(Z^2) - (EZ)^2 = E(Z^2) \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} z^2 e^{-z^2/2} dz \\ &= \frac{2}{\sqrt{2\pi}} \int_0^{\infty} z^2 e^{-z^2/2} dz \end{aligned}$$

The last step uses the fact that $z^2 e^{-z^2/2}$ is an even function. Now we use integration by parts with $u = z$ and $dv = ze^{-z^2/2} dz$, so $du = dz$ and $v = -e^{-z^2/2}$:

$$\begin{aligned} \text{Var}(Z) &= \frac{2}{\sqrt{2\pi}} \left(-ze^{-z^2/2} \Big|_0^{\infty} + \int_0^{\infty} e^{-z^2/2} dz \right) \\ &= \frac{2}{\sqrt{2\pi}} \left(0 + \frac{\sqrt{2\pi}}{2} \right) \\ &= 1 \end{aligned}$$

The first term of the integration by parts equals 0 because $e^{-z^2/2}$ decays much faster than z grows, and the second term is $\sqrt{2\pi}/2$ because it's half of the total area under $e^{-z^2/2}$, which we've already proved is $\sqrt{2\pi}$. So indeed, the standard Normal distribution has mean 0 and variance 1.

The general Normal distribution has two parameters, denoted μ and σ^2 , which correspond to the mean and variance (so the standard Normal is the special case where $\mu = 0$ and $\sigma^2 = 1$). Starting with a standard Normal r.v. $Z \sim \mathcal{N}(0, 1)$, we can get a Normal r.v. with any mean and variance by a location-scale transformation (shifting and scaling).

Definition 5.4.3 (Normal distribution). If $Z \sim \mathcal{N}(0, 1)$, then

$$X = \mu + \sigma Z$$

is said to have the Normal distribution with mean μ and variance σ^2 , for any real μ and σ^2 with $\sigma > 0$. We denote this by $X \sim \mathcal{N}(\mu, \sigma^2)$.

It's clear by properties of expectation and variance that X does in fact have mean μ and variance σ^2 :

$$\begin{aligned} E(\mu + \sigma Z) &= E(\mu) + \sigma E(Z) = \mu, \\ \text{Var}(\mu + \sigma Z) &= \text{Var}(\sigma Z) = \sigma^2 \text{Var}(Z) = \sigma^2. \end{aligned}$$

⁰¹ A subtlety is that $\infty - \infty$ is undefined, so we also want to check that the area under the curve $z^n e^{-z^2/2}$ from 0 to ∞ is finite. But this is true since $e^{-z^2/2}$ goes to 0 extremely quickly (faster than exponential decay), more than offsetting the growth of the polynomial z^n .

Note that we multiply Z by the standard deviation σ , not σ^2 ; else the units would be wrong and X would have variance σ^4 .

Of course, if we can get from Z to X , then we can get from X back to Z . The process of getting a standard Normal from a non-standard Normal is called, appropriately enough, standardization. For $X \sim \mathcal{N}(\mu, \sigma^2)$, the standardized version of X is

$$\frac{X - \mu}{\sigma} \sim \mathcal{N}(0, 1)$$

We can use standardization to find the CDF and PDF of X in terms of the standard Normal CDF and PDF.

Theorem 5.4.4 (Normal CDF and PDF). Let $X \sim \mathcal{N}(\mu, \sigma^2)$. Then the CDF of X is

$$F(x) = \Phi\left(\frac{x - \mu}{\sigma}\right)$$

and the PDF of X is

$$f(x) = \varphi\left(\frac{x - \mu}{\sigma}\right) \frac{1}{\sigma}.$$

Proof. For the CDF, we start from the definition $F(x) = P(X \leq x)$, standardize, and use the CDF of the standard Normal:

$$F(x) = P(X \leq x) = P\left(\frac{X - \mu}{\sigma} \leq \frac{x - \mu}{\sigma}\right) = \Phi\left(\frac{x - \mu}{\sigma}\right).$$

Then we differentiate to get the PDF, remembering to apply the chain rule:

$$\begin{aligned} f(x) &= \frac{d}{dx} \Phi\left(\frac{x - \mu}{\sigma}\right) \\ &= \varphi\left(\frac{x - \mu}{\sigma}\right) \frac{1}{\sigma}. \end{aligned}$$

We can also write out the PDF as

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right).$$

Finally, three important benchmarks for the Normal distribution are the probabilities of falling within one, two, and three standard deviations of the mean. The 68-95-99.7% rule tells us that these probabilities are what the name suggests.

Theorem 5.4.5 (68-95-99.7% rule). If $X \sim \mathcal{N}(\mu, \sigma^2)$, then

$$\begin{aligned} P(|X - \mu| < \sigma) &\approx 0.68 \\ P(|X - \mu| < 2\sigma) &\approx 0.95 \\ P(|X - \mu| < 3\sigma) &\approx 0.997. \end{aligned}$$

We can use this rule to get quick approximations of Normal probabilities. ² Often it is easier to apply the rule after standardizing, in which case we have

$$\begin{aligned} P(|Z| < 1) &\approx 0.68, \\ P(|Z| < 2) &\approx 0.95, \\ P(|Z| < 3) &\approx 0.997. \end{aligned}$$

Example 5.4.6 (Practice with the standard Normal CDF). Let $X \sim \mathcal{N}(-1, 4)$. What is $P(|X| < 3)$, exactly (in terms of Φ) and approximately?

Solution:

The event $|X| < 3$ is the same as the event $-3 < X < 3$. We use standardization to express this event in terms of the standard Normal r.v. $Z = (X - (-1))/2$, then apply the 68 – 95 – 99.7% rule to get an approximation. The exact answer is

$$P(-3 < X < 3) = P\left(\frac{-3 - (-1)}{2} < \frac{X - (-1)}{2} < \frac{3 - (-1)}{2}\right) = P(-1 < Z < 2),$$

which is $\Phi(2) - \Phi(-1)$. The 68-95-99.7% rule tells us that $P(-1 < Z < 1) \approx 0.68$ and $P(-2 < Z < 2) \approx 0.95$. In other words, going from ± 1 standard deviation to ± 2 standard deviations adds approximately $0.95 - 0.68 = 0.27$ to the area under the curve. By symmetry, this is evenly divided between the areas $P(-2 < Z < -1)$ and $P(1 < Z < 2)$. Therefore,

$$P(-1 < Z < 2) = P(-1 < Z < 1) + P(1 < Z < 2) \approx 0.68 + \frac{0.27}{2} = 0.815.$$

This is close to the correct value, $\Phi(2) - \Phi(-1) \approx 0.8186$.

As we will see later in the book, several important distributions can be obtained through transforming Normal r.v.s in natural ways, e.g., squaring or exponentiating. Chapter 8 delves into transformations in depth, but meanwhile there is a lot that we can do just using LOTUS and properties of CDFs.

Example 5.4.7 (Folded Normal). Let $Y = |Z|$ with $Z \sim \mathcal{N}(0, 1)$. The distribution of Y is called a Folded Normal with parameters $\mu = 0$ and $\sigma^2 = 1$. In this example, we will derive the mean, variance, and distribution of Y . At first sight, Y may seem tricky to deal with since the absolute value function is not differentiable at 0 (due to its sharp corner), but Y has a perfectly valid continuous distribution.

- (a) Find $E(Y)$.
- (b) Find $\text{Var}(Y)$.
- (c) Find the CDF and PDF of Y .

Solution:

(a) We will derive the PDF of Y later in this example, but to find $E(Y)$, LOTUS says we can work directly with the PDF of Z :

$$E(Y) = E|Z| = \int_{-\infty}^{\infty} |z| \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz = 2 \int_0^{\infty} z \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz = \sqrt{\frac{2}{\pi}}$$

(b) Note that $Y^2 = Z^2$, so we do not need to do another integral! We have

$$E(Y^2) = E(Z^2) = \text{Var}(Z) = 1$$

so

⁰² The 68 – 95 – 99.7% rule says that 95% of the time, a Normal random variable will fall within ± 2 standard deviations of its mean. An even more accurate approximation says that 95% of the time, a Normal r.v. is within ± 1.96 SDs of its mean. This explains why the number 1.96 comes up very often in statistics in the context of 95% confidence intervals, which are often created by taking an estimate and putting a buffer zone of 1.96 SDs on either side.

$$\text{Var}(Y) = E(Y^2) - (E(Y))^2 = 1 - \frac{2}{\pi}.$$

(c) For $y \leq 0$, the CDF of Y is $F_Y(y) = P(Y \leq y) = 0$. For $y > 0$, the CDF is $F_Y(y) = P(Y \leq y) = P(|Z| \leq y) = P(-y \leq Z \leq y) = \Phi(y) - \Phi(-y) = 2\Phi(y) - 1$. So the PDF of Y is $2\varphi(y)$ for $y \geq 0$, and 0 otherwise, where φ is the $\mathcal{N}(0, 1)$ PDF. Sanity check: Note that $2\Phi(y) - 1 \rightarrow 2 - 1 = 1$ as $y \rightarrow \infty$, as it must, and that the CDF of Y is a continuous function since Φ is continuous and at 0 there is not a jump: $2\Phi(0) - 1 = 0$. Also, the PDF of Y makes sense since taking the absolute value of Z "folds" the probability mass of a negative range of values of Z over to the positive side, e.g., the probability for Z values between -2 and -1 contributes to the probability for Y values between 1 and 2. This results in zero density for negative values and double the density for positive values.

4.5.5 5.5 Exponential

The Exponential distribution is the continuous counterpart to the Geometric distribution. Recall that a Geometric random variable counts the number of failures before the first success in a sequence of Bernoulli trials. The story of the Exponential distribution is analogous, but we are now waiting for a success in continuous time, where successes arrive at a rate of λ successes per unit of time. The average number of successes in a time interval of length t is λt , though the actual number of successes varies randomly. An Exponential random variable represents the waiting time until the first arrival of a success.

Definition 5.5.1 (Exponential distribution). A continuous r.v. X is said to have the Exponential distribution with parameter λ , where $\lambda > 0$, if its PDF is

$$f(x) = \lambda e^{-\lambda x}, \quad x > 0.$$

We denote this by $X \sim \text{Expo}(\lambda)$.
The corresponding CDF is

$$F(x) = 1 - e^{-\lambda x}, \quad x > 0.$$

The $\text{Expo}(1)$ PDF and CDF are plotted in Figure 5.12. Note the resemblance to the Geometric PMF and CDF pictured in Chapter 4. Exercise 43 explores the sense in which the Geometric converges to the Exponential, in the limit where the Bernoulli trials are performed faster and faster but with smaller and smaller success probabilities.




FIGURE 5.12
Expo(1) PDF and CDF.

We've seen how all Uniform and Normal distributions are related to one another via location-scale transformations, and we might wonder whether the Exponential distribution allows this

too. Exponential r.v.s are defined to have support $(0, \infty)$, and shifting would change the left endpoint. But scale transformations work nicely, and we can use scaling to get from the simple $\text{Expo}(1)$ to the general $\text{Expo}(\lambda)$: if $X \sim \text{Expo}(1)$, then

$$Y = \frac{X}{\lambda} \sim \text{Expo}(\lambda),$$

since

$$P(Y \leq y) = P\left(\frac{X}{\lambda} \leq y\right) = P(X \leq \lambda y) = 1 - e^{-\lambda y}, \quad y > 0.$$

Conversely, if $Y \sim \text{Expo}(\lambda)$, then $\lambda Y \sim \text{Expo}(1)$.

This means that just as we did for the Uniform and the Normal, we can get the mean and variance of the Exponential distribution by starting with $X \sim \text{Expo}(1)$. Both $E(X)$ and $\text{Var}(X)$ are obtained using standard integration by parts calculations. This gives

$$\begin{aligned} E(X) &= \int_0^\infty x e^{-x} dx = 1 \\ E(X^2) &= \int_0^\infty x^2 e^{-x} dx = 2 \\ \text{Var}(X) &= E(X^2) - (E(X))^2 = 1 \end{aligned}$$

In the next chapter we'll introduce a new tool called the moment generating function, which will let us get these results without integration.

For $Y = X/\lambda \sim \text{Expo}(\lambda)$ we then have

$$\begin{aligned} E(Y) &= \frac{1}{\lambda} E(X) = \frac{1}{\lambda} \\ \text{Var}(Y) &= \frac{1}{\lambda^2} \text{Var}(X) = \frac{1}{\lambda^2}, \end{aligned}$$

so the mean and variance of the $\text{Expo}(\lambda)$ distribution are $1/\lambda$ and $1/\lambda^2$, respectively. As we'd expect intuitively, the faster the rate of arrivals λ , the shorter the average waiting time.

The Exponential distribution has a very special property called the memoryless property, which says that even if you've waited for hours or days without success, the success isn't any more likely to arrive soon. In fact, you might as well have just started waiting 10 seconds ago. The definition formalizes this idea.

Definition 5.5.2 (Memoryless property). A continuous distribution is said to have the memoryless property if a random variable X from that distribution satisfies

$$P(X \geq s + t \mid X \geq s) = P(X \geq t)$$

for all $s, t \geq 0$.

Here s represents the time you've already spent waiting; the definition says that after you've waited s minutes, the probability you'll have to wait another t minutes is exactly the same as the probability of having to wait t minutes with no previous waiting time under your belt. Another way to state the memoryless property is that conditional on $X \geq s$, the additional waiting time $X - s$ is still distributed $\text{Expo}(\lambda)$. In particular, this implies

$$E(X \mid X \geq s) = s + E(X) = s + \frac{1}{\lambda}$$

(Conditional expectation is explained in detail in Chapter 9, but the meaning should already be clear: for any r.v. X and event A , $E(X \mid A)$ is the expected value of X given A ; this can

be defined by replacing the unconditional PMF or PDF of X in the definition of $E(X)$ by the conditional PMF or PDF of X given A .)

Using the definition of conditional probability, we can directly verify that the Exponential distribution has the memoryless property. Let $X \sim \text{Expo}(\lambda)$. Then

$$P(X \geq s + t \mid X \geq s) = \frac{P(X \geq s + t)}{P(X \geq s)} = \frac{e^{-\lambda(s+t)}}{e^{-\lambda s}} = e^{-\lambda t} = P(X \geq t).$$

What are the implications of the memoryless property? If you're waiting at a bus stop and the time until the bus arrives has an Exponential distribution, then conditional on your having waited 30 minutes, the bus isn't due to arrive soon. The distribution simply forgets that you've been waiting for half an hour, and your remaining wait time is the same as if you had just shown up to the bus stop. If the

lifetime of a machine has an Exponential distribution, then no matter how long the machine has been functional, conditional on having lived that long, the machine is as good as new: there is no wear-and-tear effect that makes the machine more likely to break down soon. If human lifetimes were Exponential, then conditional on having survived to the age of 80, your remaining lifetime would have the same distribution as that of a newborn baby!

Clearly, the memoryless property is not an appropriate description for human or machine lifetimes. Why then do we care about the Exponential distribution?

1. Some physical phenomena, such as radioactive decay, truly do exhibit the memoryless property, so the Exponential is an important model in its own right.
2. The Exponential distribution is well-connected to other named distributions. In the next section, we'll see how the Exponential and Poisson distributions can be united by a shared story, and we'll discover many more connections in later chapters.
3. The Exponential serves as a building block for more flexible distributions, such as the Weibull distribution (introduced in Chapter 6), that allow for a wear-and-tear effect (where older units are due to break down) or a survival-of-the-fittest effect (where the longer you've lived, the stronger you get). To understand these distributions, we first have to understand the Exponential.

The memoryless property is a very special property of the Exponential distribution: no other continuous distribution on $(0, \infty)$ is memoryless! Let's prove this.

Theorem 5.5.3. If X is a positive continuous random variable with the memoryless property, then X has an Exponential distribution.

Proof. Suppose X is a positive continuous r.v. with the memoryless property. Let F be the CDF of X and G be the survival function of X , given by $G(x) = 1 - F(x)$. We will show that $G(x) = e^{-\lambda x}$ for some λ , by first showing that $G(xt) = G(t)^x$ for all real $x > 0$. The memoryless property says that

$$G(s + t) = G(s)G(t)$$

for all $s, t \geq 0$. Putting $s = t$, we have

$$G(2t) = G(t)^2$$

so

$$G(3t) = G(2t + t) = G(2t)G(t) = G(t)^3, G(4t) = G(t)^4, \dots$$

We now have

$$G(mt) = G(t)^m$$

for m a positive integer. Let's extend this in stages. Replacing t by $t/2$ in $G(2t) = G(t)^2$, we have $G(t/2) = G(t)^{1/2}$. Similarly,

$$G\left(\frac{t}{n}\right) = G(t)^{1/n}$$

for any positive integer n . It follows that

$$G\left(\frac{m}{n}t\right) = \left(G\left(\frac{t}{n}\right)\right)^m = G(t)^{m/n}$$

for any positive integers m, n , so

$$G(xt) = G(t)^x$$

for all positive rational numbers x . Any positive real number can be written as a limit of positive rational numbers so, using the fact that G is a continuous function, the above equation holds for all positive real numbers x . Taking $t = 1$, we have

$$G(x) = G(1)^x = e^{-\lambda x}$$

where $\lambda = -\log(G(1)) > 0$. This is exactly the form we wanted for G , so X has an Exponential distribution.

The memoryless property is defined analogously for discrete distributions: a discrete distribution is memoryless if for X an r.v. with that distribution,

$$P(X \geq j + k \mid X \geq j) = P(X \geq k)$$

for all nonnegative integers j, k . In view of the analogy between the Geometric and Exponential stories (or if you have solved Exercise 32 from Chapter 4), you might guess that the Geometric distribution is memoryless. If so, you would be correct! If we're waiting for the first Heads in a sequence of fair coin tosses, and in a streak of bad luck we happen to get ten Tails in a row, this has no impact on how many additional tosses we'll need: the coin isn't due for a Heads, nor conspiring against us to perpetually land Tails. The coin is memoryless. The Geometric is the only memoryless discrete distribution taking values in $\{0, 1, 2, \dots\}$, and the Exponential is the only memoryless continuous distribution taking values in $(0, \infty)$.

As practice with the memoryless property, the following example chronicles the adventures of Fred, who experiences firsthand the frustrations of the memoryless property after moving to a town with a memoryless public transportation system.

Example 5.5.4 (Blissville and Blotchville). Fred lives in Blissville, where buses always arrive exactly on time, with the time between successive buses fixed at 10 minutes. Having lost his watch, he arrives at the bus stop at a uniformly random time on a certain day (assume that buses run 24 hours a day, every day, and that the time that Fred arrives is independent of the bus arrival process).

(a) What is the distribution of how long Fred has to wait for the next bus? What is the average time that Fred has to wait?

(b) Given that the bus has not yet arrived after 6 minutes, what is the probability that Fred will have to wait at least 3 more minutes?

(c) Fred moves to Blotchville, a city with inferior urban planning and where buses are much more erratic. Now, when any bus arrives, the time until the next bus arrives is an Exponential random variable with mean 10 minutes. Fred arrives at the bus stop at a random time (assume

that Blotchville has followed and will follow this system for all of eternity, and that the time that Fred arrives is independent of the bus arrival process). What is the distribution of Fred's waiting time for the next bus? What is the average time that Fred has to wait?

(d) When Fred complains to a friend how much worse transportation is in Blotchville, the friend says: "Stop whining so much! You arrive at a uniform instant between the previous bus arrival and the next bus arrival. The average length of that interval between buses is 10 minutes, but since you are equally likely to arrive at any time in that interval, your average waiting time is only 5 minutes."

Fred disagrees, both from experience and from solving Part (c) while waiting for the bus. Explain what is wrong with the friend's reasoning.

Solution:

(a) The distribution is Uniform on $(0, 10)$, so the mean is 5 minutes.

(b) Let T be the waiting time. Then

$$P(T \geq 6 + 3 \mid T > 6) = \frac{P(T \geq 9, T > 6)}{P(T > 6)} = \frac{P(T \geq 9)}{P(T > 6)} = \frac{1/10}{4/10} = \frac{1}{4}.$$

In particular, Fred's waiting time in Blissville is not memoryless; conditional on having waited 6 minutes already, there's only a $1/4$ chance that he'll have to wait at least another 3 minutes, whereas if he had just showed up, there would be a $P(T \geq 3) = 7/10$ chance of having to wait at least 3 minutes.

(c) By the memoryless property, the distribution is Exponential with parameter $1/10$ (and mean 10 minutes) regardless of when Fred arrives; how much longer the next bus will take to arrive is independent of how long ago the previous bus arrived. The average time that Fred has to wait is 10 minutes.

(d) Fred's friend is making the mistake, explained in 4.1.3, of replacing a random variable (the time between buses) by its expectation (10 minutes), thereby ignoring the variability in interarrival times. The average length of a time interval between two buses is 10 minutes, but Fred is not equally likely to arrive at any of these intervals: Fred is more likely to arrive during a long interval between buses than to arrive during a short interval between buses. For example, if one interval between buses is 50 minutes and another interval is 5 minutes, then Fred is 10 times more likely to arrive during the 50-minute interval.

This phenomenon is known as length-biased sampling, and it comes up in many real-life situations. For example, asking randomly chosen mothers how many children they have yields a different distribution from asking randomly chosen people how many siblings they have, including themselves. Asking students the sizes of their classes and averaging those results may give a much higher value than taking a list of classes and averaging the sizes of each; this is called the class size paradox. See exercises 16 and 17 from Chapter 4 for more about the class size paradox and length-biased sampling.

Fred's adventures in Blissville and Blotchville continue in the exercises (see also MacKay [17] for more of Fred's adventures). The bus arrivals in Blotchville follow a Poisson process, which is the topic of the next section. $\square$

4.5.6 5.6 Poisson Processes

The Exponential distribution is closely connected to the Poisson distribution, as suggested by our use of λ for the parameters of both distributions. In this section we will see that the Exponential and Poisson are linked by a common story, which is the story of the Poisson process. A Poisson process is a sequence of arrivals occurring at different points on a timeline,

such that the number of arrivals in a particular interval of time has a Poisson distribution. Poisson processes are discussed in much greater detail in Chapter 13, but we already have the tools to understand the definition and basic properties.

Definition 5.6.1 (Poisson process). A process of arrivals in continuous time is called a Poisson process with rate λ if the following two conditions hold:

1. The number of arrivals that occur in an interval of length t is a $\text{Pois}(\lambda t)$ random variable.
2. The numbers of arrivals that occur in disjoint intervals are independent of each other. For example, the numbers of arrivals in the intervals $(0, 10)$, $[10, 12)$, and $[15, \infty)$ are independent.

In this section, we will focus on Poisson processes on $(0, \infty)$, but we can also define Poisson processes on $(-\infty, \infty)$ or other intervals, and in Chapter 13 we will introduce Poisson processes in more than one dimension. A sketch of a Poisson process on $(0, \infty)$ is pictured in Figure 5.13. Each X marks the spot of an arrival.

For concreteness, suppose that the arrivals are emails landing in an inbox according to a Poisson process with rate λ . There are several things we might want to know about this process. One question we could ask is: in one hour, how many emails will arrive? The answer comes directly from the definition, which tells us that the




FIGURE 5.13 A Poisson process on $(0, \infty)$. Each X corresponds to an arrival.

number of emails in an hour follows a $\text{Pois}(\lambda)$ distribution. Notice that the number of emails is a nonnegative integer, so a discrete distribution is appropriate.

But we could also flip the question around and ask: how long does it take until the first email arrives (measured relative to some fixed starting point)? The waiting time for the first email is a positive real number, so a continuous distribution on $(0, \infty)$ is appropriate. Let T_1 be the time until the first email arrives. To find the distribution of T_1 , we just need to understand one crucial fact: saying that the waiting time for the first email is greater than t is the same as saying that no emails have arrived between 0 and t . In other words, if N_t is the number of emails that arrive at or before time t , then

$$T_1 > t \text{ is the same event as } N_t = 0.$$

We call this the count-time duality because it connects a discrete r.v., N_t , which counts the number of arrivals, with a continuous r.v., T_1 , which marks the time of the first arrival. More generally, the count-time duality says that

$$T_n > t \text{ is the same event as } N_t < n.$$

Saying that the n th arrival has not happened yet as of time t is equivalent to saying that, up until time t , there have been fewer than n arrivals.

If two events are the same, they have the same probability. Since $N_t \sim \text{Pois}(\lambda t)$ by the definition of Poisson process,

$$P(T_1 > t) = P(N_t = 0) = \frac{e^{-\lambda t}(\lambda t)^0}{0!} = e^{-\lambda t}.$$

Therefore $P(T_1 \leq t) = 1 - e^{-\lambda t}$, so $T_1 \sim \text{Expo}(\lambda)$! The time until the first arrival in a Poisson process of rate λ has an Exponential distribution with parameter λ .

What about $T_2 - T_1$, the time between the first and second arrivals? Since disjoint intervals in a Poisson process are independent by definition, the past is irrelevant once the first arrival

occurs. Thus $T_2 - T_1$ is independent of the time until the first arrival, and by the same argument as before, $T_2 - T_1$ also has an Exponential distribution with rate λ .

Similarly, $T_3 - T_2 \sim \text{Expo}(\lambda)$ independently of T_1 and $T_2 - T_1$. Continuing in this way, we deduce that all the interarrival times are i.i.d. $\text{Expo}(\lambda)$ random variables. Thus, Poisson processes tie together two important distributions, one discrete and one continuous, and the use of a common symbol λ for both the Poisson and Exponential parameters is felicitous notation, for λ is the arrival rate in the process that unites the two distributions.

5.6.2. The total time until the second arrival, T_2 , is the sum of two independent $\text{Expo}(\lambda)$ r.v.s, T_1 and $T_2 - T_1$. This does not have an Exponential distribution, but rather a Gamma distribution, which is introduced in Chapter 8.

The story of the Poisson process provides intuition for the fact, shown below, that the minimum of independent Exponential r.v.s is another Exponential r.v.

Example 5.6.3 (Minimum of independent Expos). Let $X_1, \dots, X_n$ be independent, with $X_j \sim \text{Expo}(\lambda_j)$. Let $L = \min(X_1, \dots, X_n)$. Show that $L \sim \text{Expo}(\lambda_1 + \dots + \lambda_n)$, and interpret this intuitively.

Solution:

We can find the distribution of L by considering its survival function $P(L > t)$, since the survival function is 1 minus the CDF.

$$\begin{aligned} P(L > t) &= P(\min(X_1, \dots, X_n) > t) = P(X_1 > t, \dots, X_n > t) \\ &= P(X_1 > t) \cdots P(X_n > t) = e^{-\lambda_1 t} \cdots e^{-\lambda_n t} = e^{-(\lambda_1 + \dots + \lambda_n)t} \end{aligned}$$

The second equality holds since saying that the minimum of the X_j is greater than t is the same as saying that all of the X_j are greater than t . The third equality holds by independence of the X_j . Thus, L has the survival function (and the CDF) of an Exponential distribution with parameter $\lambda_1 + \dots + \lambda_n$.

Intuitively, we can interpret the λ_j as the rates of n independent Poisson processes. We can imagine, for example, X_1 as the waiting time for a green car to pass by, X_2 as the waiting time for a blue car to pass by, and so on, assigning a color to each X_j . Then L is the waiting time for a car with any of these colors to pass by, so it makes sense that L has a combined rate of $\lambda_1 + \dots + \lambda_n$.

5.6.4. The minimum of independent Exponentials is Exponential, but the maximum of independent Exponentials is not Exponential. However, the result about such a minimum turns out to be useful in studying such a maximum, as illustrated in the next two examples.

Example 5.6.5 (Maximum of 3 independent Exponentials). Three students are working independently on their probability homework. All 3 start at 1 pm on a certain day, and each takes an Exponential time with mean 6 hours to complete the homework. What is the earliest time at which all 3 students will have completed the homework, on average?

Solution: Label the students as 1, 2, 3, and let X_j be how long it takes student j to finish the homework. Let $\lambda = 1/6$, and let T be the time when all 3 students will have completed the homework, so $T = \max(X_1, X_2, X_3)$ with $X_i \sim \text{Expo}(\lambda)$. The CDF of T is

$$P(T \leq t) = P(X_1 \leq t, X_2 \leq t, X_3 \leq t) = (1 - e^{-\lambda t})^3.$$

So the PDF of T is

$$f_T(t) = 3\lambda e^{-\lambda t} (1 - e^{-\lambda t})^2$$

In particular, T is not Exponential.

Finding $E(T)$ by integrating $tf_T(t)$ is possible but not especially pleasant. A neater approach

is to use the memoryless property and the fact that the minimum of independent Exponentials is Exponential. We can decompose

$$T = T_1 + T_2 + T_3,$$

where $T_1 = \min(X_1, X_2, X_3)$ is how long it takes for one student to complete the homework, T_2 is the additional time it takes for a second student to complete the homework, and T_3 is the additional time until all 3 have completed the homework. Then $T_1 \sim \text{Expo}(3\lambda)$, by the result of Example 5.6.3.

By the memoryless property, at the first time when a student completes the homework the other two students are starting from fresh, so $T_2 \sim \text{Expo}(2\lambda)$. Again by the memoryless property, $T_3 \sim \text{Expo}(\lambda)$. The memoryless property also implies that T_1, T_2, T_3 are independent (which would be very useful if we were finding $\text{Var}(T)$). By linearity,

$$E(T) = \frac{1}{3\lambda} + \frac{1}{2\lambda} + \frac{1}{\lambda} = 2 + 3 + 6 = 11,$$

which shows that on average, the 3 students will have all completed the homework at midnight, 11 hours after they started.

Example 5.6.6 (Machine repair). A certain machine often breaks down and needs to be fixed. At time 0, the machine is working. It works for an $\text{Expo}(\lambda)$ period of time (measured in days), and then breaks down. It then takes an $\text{Expo}(\lambda)$ amount of time to get it fixed, after which it will work for an $\text{Expo}(\lambda)$ time until it breaks down again, after which it will take an $\text{Expo}(\lambda)$ time to get it fixed, etc. Assume that these $\text{Expo}(\lambda)$ r.v.s are i.i.d.

(a) A transition occurs when the machine switches from working to being broken, or switches from being broken to working. Find the distribution of the number of transitions that occur in the time interval $(0, t)$.

(b) Hoping to reduce the frequency of breakdowns, the machine is redesigned so that it can continue to function even if one component has failed. The redesigned machine has 5 components, each of which works for an $\text{Expo}(\lambda)$ amount of time and then fails, independently. The machine works properly if and only if at most one component has failed. Currently, all 5 components are working (none have failed). Find the expected time until the machine breaks down.

Solution:

(a) The times between transitions are i.i.d. $\text{Expo}(\lambda)$, so the times at which transitions occur follow a Poisson process of rate λ . So the desired distribution is $\text{Pois}(\lambda t)$.

(b) The time until a component fails is $\text{Expo}(5\lambda)$. Then by the memoryless property, the additional time until another component fails is $\text{Expo}(4\lambda)$. So the expected time until the machine breaks down is

$$\frac{1}{5\lambda} + \frac{1}{4\lambda} = \frac{9}{20\lambda}$$

4.5.7 5.7 Symmetry of i.i.d. Continuous r.v.s

Continuous r.v.s that are independent and identically distributed have an important symmetry property: all possible orderings are equally likely. Intuitively, this is because if all we are told is that $X_1, \dots, X_n$ are i.i.d., then they are interchangeable, in the sense that we have been given no information that distinguishes one X_i from another X_j . This is reminiscent of the fact that it is common for someone to say "Do you want the good news first or the bad news first?" but rare for someone to say "I have two pieces of news. Which do you want to hear first?",

since in the latter case no distinguishing information has been provided for the two pieces of news.

Proposition 5.7.1. Let $X_1, \dots, X_n$ be i.i.d. from a continuous distribution. Then

$$P(X_{a_1} < X_{a_2} < \dots < X_{a_n}) = \frac{1}{n!}$$

for any permutation $a_1, a_2, \dots, a_n$ of $1, 2, \dots, n$.

Proof. Let F be the CDF of X_j . By symmetry, all orderings of $X_1, \dots, X_n$ are equally likely. For example, $P(X_3 < X_2 < X_1) = P(X_1 < X_2 < X_3)$ since both sides have exactly the same structure: they are both of the form $P(A < B < C)$ where A, B, C are i.i.d. draws from F . For any i and j with $i \neq j$, the probability of the tie $X_i = X_j$ is 0 since X_i and X_j are independent continuous r.v.s. So the probability of there being at least one tie among $X_1, \dots, X_n$ is also 0, since

$$P\left(\bigcup_{i \neq j} \{X_i = X_j\}\right) \leq \sum_{i \neq j} P(X_i = X_j) = 0.$$

Thus, $X_1, \dots, X_n$ are distinct with probability 1, and the probability of any particular ordering is $1/n!$.

5.7.2. This proposition may fail if the r.v.s are dependent. Let $n = 2$, and consider the extreme case where X_1 and X_2 are so dependent that they are always equal, i.e., $X_1 = X_2$ with probability 1. Then $P(X_1 < X_2) = P(X_2 < X_1) = 0$. For dependent X_1, X_2 we can also make $P(X_1 < X_2) \neq P(X_2 < X_1)$. For an example, see Exercise 42 from Chapter 3.

5.7.3. If X and Y are i.i.d. continuous r.v.s, then

$$P(X < Y) = P(Y < X) = \frac{1}{2}$$

by symmetry and since the probability of a tie is 0. In contrast, if X and Y are i.i.d. discrete r.v.s, it is still true that $P(X < Y) = P(Y < X)$ by symmetry, but this number is less than $1/2$ because of the possibility of a tie. For example, if X and Y are i.i.d. nonnegative integer-valued r.v.s with $P(X = j) = c_j$, then

$$1 = P(X < Y) + P(X = Y) + P(Y < X) = 2P(X < Y) + P(X = Y)$$

so

$$P(X < Y) = \frac{1}{2} \cdot (1 - P(X = Y)) = \frac{1}{2} \cdot \left(1 - \sum_{j=0}^{\infty} c_j^2\right) < \frac{1}{2}.$$

The ranks of a list of distinct numbers are defined by giving the smallest number a rank of 1, the second smallest a rank of 2, and so on. For example, the ranks for 3.14, 2.72, 1.41, 1.62 are 4, 3, 1, 2. Proposition 5.7.1 says that the ranks of i.i.d. continuous $X_1, \dots, X_n$ are a uniformly random permutation of the numbers $1, \dots, n$. The next example shows how we can use this symmetry property in conjunction with indicator r.v.s in problems involving records, such as the record level of rainfall or the record performance on a high jump.

Example 5.7.4 (Records). Athletes compete one at a time at the high jump. Let X_j be how high the j th jumper jumped, with $X_1, X_2, \dots$ i.i.d. with a continuous distribution. We say that the j th jumper sets a record if X_j is greater than all of $X_{j-1}, \dots, X_1$.

(a) Is the event "the 110th jumper sets a record" independent of the event "the 111th jumper sets a record"?

(b) Find the mean number of records among the first n jumpers. What happens to the mean as $n \rightarrow \infty$?

(c) A double record occurs at time j if both the j th and $(j-1)$ st jumpers set records. Find the mean number of double records among the first n jumpers. What happens to the mean as $n \rightarrow \infty$?

Solution:

(a) Let I_j be the indicator r.v. for the j th jumper setting a record. By symmetry, $P(I_j = 1) = 1/j$ (as any of the first j jumps is equally likely to be the highest of those jumps). Also,

$$P(I_{110} = 1, I_{111} = 1) = \frac{109!}{111!} = \frac{1}{110 \cdot 111},$$

since in order for both the 110th and 111th jumps to be records, we need the highest of the first 111 jumps to be in position 111 and the second highest to be in position 110, and the remaining 109 can be in any order. So

$$P(I_{110} = 1, I_{111} = 1) = P(I_{110} = 1)P(I_{111} = 1)$$

which shows that the 110th jumper setting a record is independent of the 111th jumper setting a record. Intuitively, this makes sense since learning that the 111th jumper sets a record gives us no information about the "internal" matter of how the first 110 jumps are arranged amongst themselves.

(b) By linearity, the expected number of records among the first n jumpers is $\sum_{j=1}^n \frac{1}{j}$, which goes to ∞ as $n \rightarrow \infty$ since the harmonic series diverges.

(c) Let J_j be the indicator r.v. for a double record occurring at time j , for $2 \leq j \leq n$. Then $P(J_j = 1) = \frac{1}{j(j-1)}$, following the logic of Part (a). So the expected number of double records is

$$\sum_{j=2}^n \frac{1}{j(j-1)} = \sum_{j=2}^n \left(\frac{1}{j-1} - \frac{1}{j} \right) = 1 - \frac{1}{n},$$

since all the other terms cancel out. Thus, the expected number of records goes to ∞ as $n \rightarrow \infty$, but the expected number of double records goes to 1. $\square$

4.5.8 5.8 Recap

A continuous r.v. can take on any value in an interval, although the probability that it equals any particular value is 0. The CDF of a continuous r.v. is differentiable, and the derivative is called the probability density function (PDF). Probability is given by area under the PDF curve, not by the value of the PDF at a point. We must integrate the PDF to get a probability. The table below summarizes and compares some important concepts in the discrete case and the continuous case.

	Discrete r.v.	Continuous r.v.
CDF	$F(x) = P(X \leq x)$	$F(x) = P(X \leq x)$
PMF/PDF	- PMF is height of jump of F at x .	- $P(X \in A) = \int_A f(x)dx$.
Expectation	$E(X) = \sum_x xP(X = x)$	$E(X) = \int_{-\infty}^{\infty} xf(x)dx$
LOTUS	$E(g(X)) = \sum_x g(x)P(X = x)$	$E(g(X)) = \int_{-\infty}^{\infty} g(x)f(x)dx$

Three important continuous distributions are the Uniform, Normal, and Exponential. A $\text{Unif}(a, b)$ r.v. is a "completely random" number in the interval (a, b) , and it has the property that probability is proportional to length. The universality of the Uniform tells us how we can

use a $\text{Unif}(0, 1)$ r.v. to construct r.v.s from other distributions we may be interested in; it also says that if we plug a continuous r.v. into its own CDF, the resulting r.v. has a $\text{Unif}(0, 1)$ distribution.

A $\mathcal{N}(\mu, \sigma^2)$ r.v. has a symmetric bell-shaped PDF centered at μ , with σ controlling how spread out the curve is. The mean is μ and standard deviation is σ . The 68-95-99.7% rule gives important benchmarks for the probability of a Normal r.v. falling within 1, 2, and 3 standard deviations of its mean.

An $\text{Expo}(\lambda)$ r.v. represents the waiting time for the first success in continuous time, analogous to how a Geometric r.v. represents the number of failures before the first success in discrete time; the parameter λ can be interpreted as the rate at which successes arrive. The Exponential distribution has the memoryless property, which says that conditional on our having waited a certain amount of time without success, the distribution of the remaining wait time is exactly the same as if we hadn't waited at all. In fact, the Exponential is the only positive continuous distribution with the memoryless property.


A Poisson process is a sequence of arrivals in continuous time such that the number of arrivals in an interval is Poisson (with mean proportional to the length of the interval) and disjoint intervals have independent numbers of arrivals. The interarrival times in a Poisson process of rate λ are i.i.d. $\text{Expo}(\lambda)$ r.v.s.

A useful symmetry property of i.i.d. r.v.s $X_1, X_2, \dots, X_n$ is that all orderings are equally likely. For example, $P(X_1 < X_2 < X_3) = P(X_3 < X_2 < X_1)$. If the X_j are continuous in addition to being i.i.d., then we can also conclude, e.g., that $P(X_1 < X_2 < X_3) = 1/6$, whereas in the discrete case we also have to account for the possibility of ties.

A new strategy that we learned for continuous distributions is location-scale transformation, which says that if shifting and scaling will not take us outside the family of distributions we're studying, then we can start with the simplest member of the family, find the answer for the simple case, then use shifting and scaling to arrive at the general case. For the three main distributions of this chapter, this approach works as follows.

- Uniform: If $U \sim \text{Unif}(0, 1)$, then $\tilde{U} = a + (b - a)U \sim \text{Unif}(a, b)$.
- Normal: If $Z \sim \mathcal{N}(0, 1)$, then $X = \mu + \sigma Z \sim \mathcal{N}(\mu, \sigma^2)$.
- Exponential: If $X \sim \text{Expo}(1)$, then $Y = X/\lambda \sim \text{Expo}(\lambda)$. We do not consider shifts here since a nonzero shift would prevent the support from being $(0, \infty)$.

We can now add the Exponential and Geometric distributions to our diagram of connections between distributions: the Exponential is a continuous limit of the Geometric, and the Poisson and Exponential are connected by the Poisson process.



And in our map of the four fundamental objects in probability, we add the PDF as another blueprint for continuous random variables.




FIGURE 5.14

Four fundamental objects in probability: distributions, random variables, events, and numbers. For a continuous r.v. X , we have $P(X = x) = 0$, so we use the PDF as a blueprint in place of the PMF.

4.5.9 5.9r

In this section we will introduce continuous distributions in R, learn how to make basic plots, demonstrate the universality of the Uniform by simulation, and simulate arrival times in a Poisson process.

Uniform, Normal, and Exponential Distributions

For continuous distributions, the function that starts with d is the PDF instead of the PMF. Thus, we have the following functions:

- `dunif`, `punif`, `runif`. To evaluate the $\text{Unif}(a, b)$ PDF at x , we use `dunif(x, a, b)`. For the CDF, we use `punif(x, a, b)`. To generate n realizations from the $\text{Unif}(a, b)$ distribution, we use `runif(n, a, b)`.
 - `dnorm`, `pnorm`, `rnorm`. To evaluate the $\mathcal{N}(\mu, \sigma^2)$ PDF at x , we use `dnorm(x, mu, sigma)`, where `mu` is the mean μ and `sigma` is the standard deviation (not variance) σ . For the CDF, we use `pnorm(x, mu, sigma)`. To generate n realizations from the $\mathcal{N}(\mu, \sigma^2)$ distribution, we use `rnorm(n, mu, sigma)`.
- 5.9.1 (Normal parameters in R). Note that we have to input the standard deviation, not the variance! For example, to get the $\mathcal{N}(10, 3)$ CDF at 12, we use `pnorm(12, 10, sqrt(3))`. Ignoring this is a common, disastrous coding error.
- `dexp`, `pexp`, `rexp`. To evaluate the $\text{Expo}(\lambda)$ PDF at x , we use `dexp(x, lambda)`. For the CDF, we use `pexp(x, lambda)`. To generate n realizations from the $\text{Expo}(\lambda)$ distribution, we use `rexp(n, lambda)`.

Due to the importance of location-scale transformations for continuous distributions, R has default parameter settings for each of these three families. The default for the Uniform is $\text{Unif}(0, 1)$, the default for the Normal is $\mathcal{N}(0, 1)$, and the default for the Exponential is $\text{Expo}(1)$. For example, `dunif(0.5)`, with no additional inputs, evaluates the $\text{Unif}(0, 1)$ PDF at 0.5, and `rnorm(10)`, with no additional inputs, generates 10 realizations from the $\mathcal{N}(0, 1)$ distribution.

This means there are two ways to generate a $\mathcal{N}(\mu, \sigma^2)$ random variable in R . After choosing our values of μ and σ , such as

```
mu <- 1
sigma <- 2
```

we can do either of the following:

```
rnorm(1,mu,sigma)
mu + sigma*rnorm(1)
```

Either way, we end up generating a draw from the $\mathcal{N}(\mu, \sigma^2)$ distribution.

Plots in R

A simple way to plot a function in R is with the `curve` command. For example,

```
curve(dnorm, from=-3, to=3, n=1000)
```

creates a plot of the standard Normal PDF from -3 to 3 . What is actually happening is that R evaluates the function at a finite number of closely spaced points and connects the points with very short lines to create the illusion of a curve. The input `n = 1000` tells R to evaluate at 1000 points so that the curve looks very smooth; if we were to choose $n = 20$, the piecewise linearity would become very apparent.

Another command that creates plots is called, fittingly, `plot`. This command has many, many possible inputs to customize what the plot looks like; for the sake of demonstration, we'll plot the standard Normal PDF once again, using `plot` instead of `curve`.

The most important inputs to `plot` are a vector of x values and a vector of y values to plot. A useful command for this purpose is `seq`. As introduced in Chapter 1, `seq(a, b, d)` creates the vector of values ranging from a to b , with successive entries spaced apart by d .

```
x <- seq(-3,3,0.01)
y <- dnorm(x)
```

So `x` consists of all numbers from -3 to 3 , spaced 0.01 apart, and `y` contains the values of the Normal PDF at each of the points in `x` . Now we simply plot the two with `plot(x,y)`. The default is a scatterplot. For a line plot, we use `plot(x,y, type = "l")`. We can also set the axis labels and plot title with `xlab`, `ylab`, and `main`.

```
plot(x,y,type="l",xlab="x",ylab="dnorm(x)",main="N(0,1) PDF")
```

The axis limits can be set manually with `xlim` and `ylim`. If, for example, you wanted the vertical axis to range from 0 to 1 , you would add `ylim = c(0,1)` inside the plot command.

Finally, to change the color of the plot, add `col="orange"` or `col="green"`, or whatever your favorite color is!

Universality With Logistic

We proved in Example 5.3.4 that for $U \sim \text{Unif}(0, 1)$, the r.v. $\log(U/(1 - U))$ follows a Logistic distribution. In R, we can simply generate a large number of $\text{Unif}(0, 1)$ realizations and transform them.

```
u <- runif(10^4)
x <- log(u/(1-u))
```

Now x contains 10^4 realizations from the distribution of $\log(U/(1 - U))$. We can visualize them with a histogram, using the command `hist(x)`. The histogram resembles a Logistic PDF, which is reassuring. To control how fine-grained the histogram is, we can set the number of breaks in the histogram: `hist(x, breaks = 100)` produces a finer histogram, while `hist(x, breaks = 10)` produces a coarser histogram.

Poisson Process Simulation

To simulate n arrivals in a Poisson process with rate λ , we first generate the interarrival times as i.i.d. Exponentials and store them in a vector:

```
n <- 50
lambda <- 10
x <- rexp(n, lambda)
```

Then we convert the interarrival times into arrival times using the `cumsum` function, which stands for "cumulative sum".

```
t <- cumsum(x)
```

The vector t now contains all the simulated arrival times.

4.5.10 5.10 Exercises

Exercises marked with (S) have detailed solutions at <http://stat110.net>.

PDFs and Cdfs

1. The Rayleigh distribution from Example 5.1.7 has PDF

$$f(x) = xe^{-x^2/2}, \quad x > 0$$

Let X have the Rayleigh distribution.

- (a) Find $P(1 < X < 3)$.
 - (b) Find the first quartile, median, and third quartile of X ; these are defined to be the values q_1, q_2, q_3 (respectively) such that $P(X \leq q_j) = j/4$ for $j = 1, 2, 3$.
2. (a) Make up a PDF f , with an application for which that PDF would be plausible, where $f(x) > 1$ for all x in a certain interval.
(b) Show that if a PDF f has $f(x) > 1$ for all x in a certain interval, then that interval must have length less than 1.
 3. Let F be the CDF of a continuous r.v., and $f = F'$ be the PDF.
(a) Show that g defined by $g(x) = 2F(x)f(x)$ is also a valid PDF.
(b) Show that h defined by $h(x) = \frac{1}{2}f(-x) + \frac{1}{2}f(x)$ is also a valid PDF.
 4. Let X be a continuous r.v. with CDF F and PDF f .

- (a) Find the conditional CDF of X given $X > c$, for c a constant with $P(X > c) \neq 0$. That is, find $P(X \leq x \mid X > c)$ for all c , in terms of F .
- (b) Find the conditional PDF of X given $X > c$ (this is the derivative of the conditional CDF).
- (c) Check that the conditional PDF from (b) is a valid PDF, by showing directly that it is nonnegative and integrates to 1.
5. A circle with a random radius $R \sim \text{Unif}(0, 1)$ is generated. Let A be its area.
- (a) Find the mean and variance of A , without first finding the CDF or PDF of A .
- (b) Find the CDF and PDF of A .
6. The 68 – 95 – 99.7% rule gives approximate probabilities of a Normal r.v. being within 1, 2, and 3 standard deviations of its mean. Derive analogous rules for the following distributions.
- (a) $\text{Unif}(0, 1)$.
- (b) $\text{Expo}(1)$.
- (c) $\text{Expo}(1/2)$. Discuss whether there is one such rule that applies to all Exponential distributions, just as the 68-95-99.7% rule applies to all Normal distributions, not just to the standard Normal.
7. Let

$$F(x) = \frac{2}{\pi} \sin^{-1}(\sqrt{x}), \text{ for } 0 < x < 1$$

$F(x) = 0$ for $x \leq 0$, and $F(x) = 1$ for $x \geq 1$.

- (a) Check that F is a valid CDF, and find the corresponding PDF f . This distribution is called the Arcsine distribution, though it also goes by the name $\text{Beta}(1/2, 1/2)$ (we will explore the Beta in depth in Chapter 8).
- (b) Explain how it is possible for f to be a valid PDF even though $f(x)$ goes to ∞ as x approaches 0 from the right and as x approaches 1 from the left.
8. The Beta distribution with parameters $a = 3, b = 2$ has PDF

$$f(x) = 12x^2(1 - x), \text{ for } 0 < x < 1.$$

(We will discuss the Beta in detail in Chapter 8.) Let X have this distribution.

- (a) Find the CDF of X .
- (b) Find $P(0 < X < 1/2)$.
- (c) Find the mean and variance of X (without quoting results about the Beta distribution).
9. The Cauchy distribution has PDF

$$f(x) = \frac{1}{\pi(1 + x^2)}$$

for all real x . (We will introduce the Cauchy from another point of view in Chapter 7.)

Find the CDF of a random variable with the Cauchy PDF.

Hint: Recall that the derivative of the inverse tangent function $\arctan(x)$ is $\frac{1}{1+x^2}$.

Uniform and Universality of the Uniform

10. Let $U \sim \text{Unif}(0, 8)$.
- (a) Find $P(U \in (0, 2) \cup (3, 7))$ without using calculus.
- (b) Find the conditional distribution of U given $U \in (3, 7)$.
11. (S) Let U be a Uniform r.v. on the interval $(-1, 1)$ (be careful about minus signs).
- (a) Compute $E(U)$, $\text{Var}(U)$, and $E(U^4)$.
- (b) Find the CDF and PDF of U^2 . Is the distribution of U^2 Uniform on $(0, 1)$?

12. (S) A stick is broken into two pieces, at a uniformly random breakpoint. Find the CDF and average of the length of the longer piece.
13. A stick of length 1 is broken at a uniformly random point, yielding two pieces. Let X and Y be the lengths of the shorter and longer pieces, respectively, and let $R = X/Y$ be the ratio of the lengths X and Y .
- Find the CDF and PDF of R .
 - Find the expected value of R (if it exists).
 - Find the expected value of $1/R$ (if it exists).
14. Let $U_1, \dots, U_n$ be i.i.d. $\text{Unif}(0, 1)$, and $X = \max(U_1, \dots, U_n)$. What is the PDF of X ? What is EX ?

Hint: Find the CDF of X first, by translating the event $X \leq x$ into an event involving $U_1, \dots, U_n$.

15. Let $U \sim \text{Unif}(0, 1)$. Using U , construct $X \sim \text{Expo}(\lambda)$.
16. (S) Let $U \sim \text{Unif}(0, 1)$, and

$$X = \log\left(\frac{U}{1-U}\right)$$

Then X has the Logistic distribution, as defined in Example 5.1.6.

- Write down (but do not compute) an integral giving $E(X^2)$.
- Find $E(X)$ without using calculus.

Hint: A useful symmetry property here is that $1 - U$ has the same distribution as U .

17. Let $U \sim \text{Unif}(0, 1)$. As a function of U , create an r.v. X with CDF $F(x) = 1 - e^{-x^3}$ for $x > 0$.
18. The Pareto distribution with parameter $a > 0$ has PDF $f(x) = a/x^{a+1}$ for $x \geq 1$ (and 0 otherwise). This distribution is often used in statistical modeling.
- Find the CDF of a Pareto r.v. with parameter a ; check that it is a valid CDF.
 - Suppose that for a simulation you want to run, you need to generate i.i.d. $\text{Pareto}(a)$ r.v.s. You have a computer that knows how to generate i.i.d. $\text{Unif}(0, 1)$ r.v.s but does not know how to generate Pareto r.v.s. Show how to do this.

Normal

19. Let $Z \sim \mathcal{N}(0, 1)$. Create an r.v. $Y \sim \mathcal{N}(1, 4)$, as a simple-looking function of Z . Make sure to check that your Y has the correct mean and variance.
20. Engineers sometimes work with the "error function"

$$\text{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-x^2} dx$$

instead of the standard Normal CDF Φ .

- Show that the following conversion between Φ and erf holds for all z :

$$\Phi(z) = \frac{1}{2} + \frac{1}{2} \text{erf}\left(\frac{z}{\sqrt{2}}\right)$$

- Show that erf is an odd function, i.e., $\text{erf}(-z) = -\text{erf}(z)$.

21. (a) Find the points of inflection of the $\mathcal{N}(0, 1)$ PDF φ , i.e., the points where the curve switches from convex (second derivative positive) to concave (second derivative negative) or vice versa.

(b) Use the result of (a) and a location-scale transformation to find the points of inflection of the $\mathcal{N}(\mu, \sigma^2)$ PDF.

22. The distance between two points needs to be measured, in meters. The true distance between the points is 10 meters, but due to measurement error we can't measure the distance exactly. Instead, we will observe a value of $10 + \epsilon$, where the error ϵ is distributed $\mathcal{N}(0, 0.04)$. Find the probability that the observed distance is within 0.4 meters of the true distance (10 meters). Give both an exact answer in terms of Φ and an approximate numerical answer.

23. Alice is trying to transmit to Bob the answer to a yes-no question, using a noisy channel. She encodes "yes" as 1 and "no" as 0, and sends the appropriate value. However, the channel adds noise; specifically, Bob receives what Alice sends plus a $\mathcal{N}(0, \sigma^2)$ noise term (the noise is independent of what Alice sends). If Bob receives a value greater than $1/2$ he interprets it as "yes"; otherwise, he interprets it as "no".

(a) Find the probability that Bob understands Alice correctly.

(b) What happens to the result from (a) if σ is very small? What about if σ is very large? Explain intuitively why the results in these extreme cases make sense.

24. A woman is pregnant, with a due date of January 10, 2020. Of course, the actual date on which she will give birth is not necessarily the due date. On a timeline, define time 0 to be the instant when January 10, 2020 begins. Suppose that the time T when the woman gives birth has a Normal distribution, centered at 0 and with standard deviation 8 days. What is the probability that she gives birth on her due date? (Your answer should be in terms of Φ , and simplified.)

25. We will show in the next chapter that if X_1 and X_2 are independent with $X_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$, then $X_1 + X_2 \sim \mathcal{N}(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$. Use this result to find $P(X < Y)$ for $X \sim \mathcal{N}(a, b)$, $Y \sim \mathcal{N}(c, d)$ with X and Y independent.

Hint: Write $P(X < Y) = P(X - Y < 0)$ and then standardize $X - Y$. Check that your answer makes sense in the special case where X and Y are i.i.d.

26. Walter and Carl both often need to travel from Location A to Location B. Walter walks, and his travel time is Normal with mean w minutes and standard deviation σ minutes (travel time can't be negative without using a tachyon beam, but assume that w is so much larger than σ that the chance of a negative travel time is negligible).

Carl drives his car, and his travel time is Normal with mean c minutes and standard deviation 2σ minutes (the standard deviation is larger for Carl due to variability in traffic conditions). Walter's travel time is independent of Carl's. On a certain day, Walter and Carl leave from Location A to Location B at the same time.

(a) Find the probability that Carl arrives first (in terms of Φ and the parameters). For this you can use the important fact, proven in the next chapter, that if X_1 and X_2 are independent with $X_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$, then $X_1 + X_2 \sim \mathcal{N}(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$.

(b) Give a fully simplified criterion (not in terms of Φ), such that Carl has more than a 50% chance of arriving first if and only if the criterion is satisfied.

(c) Walter and Carl want to make it to a meeting at Location B that is scheduled to begin $w + 10$ minutes after they depart from Location A. Give a fully simplified criterion (not in terms of Φ) such that Carl is more likely than Walter to make it on time for the meeting if and only if the criterion is satisfied.

27. Let $Z \sim \mathcal{N}(0, 1)$. We know from the 68-95-99.7% rule that there is a 68% chance of Z being in the interval $(-1, 1)$. Give a visual explanation of whether or not there is an interval (a, b) that is shorter than the interval $(-1, 1)$, yet which has at least as large a chance as $(-1, 1)$ of containing Z .

28. Let $Y \sim \mathcal{N}(\mu, \sigma^2)$. Use the fact that $P(|Y - \mu| < 1.96\sigma) \approx 0.95$ to construct a random interval $(a(Y), b(Y))$ (that is, an interval whose endpoints are r.v.s), such that the probability that μ is in the interval is approximately 0.95. This interval is called a confidence interval for

μ ; such intervals are often desired in statistics when estimating unknown parameters based on data.

29. Let $Y = |X|$, with $X \sim \mathcal{N}(\mu, \sigma^2)$. This is a well-defined continuous r.v., even though the absolute value function is not differentiable at 0 (due to the sharp corner).

(a) Find the CDF of Y in terms of Φ . Be sure to specify the CDF everywhere.

(b) Find the PDF of Y .

(c) Is the PDF of Y continuous at 0? If not, is this a problem as far as using the PDF to find probabilities?

30. (S) Let $Z \sim \mathcal{N}(0, 1)$ and let S be a random sign independent of Z , i.e., S is 1 with probability 1/2 and -1 with probability 1/2. Show that $SZ \sim \mathcal{N}(0, 1)$.

31. (S) Let $Z \sim \mathcal{N}(0, 1)$. Find $E(\Phi(Z))$ without using LOTUS, where Φ is the CDF of Z .

32. (S) Let $Z \sim \mathcal{N}(0, 1)$ and $X = Z^2$. Then the distribution of X is called Chi-Square with 1 degree of freedom. This distribution appears in many statistical methods.

(a) Find a good numerical approximation to $P(1 \leq X \leq 4)$ using facts about the Normal distribution, without querying a calculator/computer/table about values of the Normal CDF.

(b) Let Φ and φ be the CDF and PDF of Z , respectively. Show that for any $t > 0$, $I(Z > t) \leq (Z/t)I(Z > t)$. Using this and LOTUS, derive Mills' inequality, which is the following lower bound on Φ :

$$\Phi(t) \geq 1 - \varphi(t)/t$$

33. Let $Z \sim \mathcal{N}(0, 1)$, with CDF Φ . We will show in Chapter 8 that the PDF of Z^2 is the function g given by

$$g(w) = \frac{1}{\sqrt{2\pi w}} e^{-w/2}$$

for $w > 0$, and $g(w) = 0$ for $w \leq 0$.

(a) Find expressions for $E(Z^4)$ as integrals in two different ways, one based on the PDF of Z and the other based on the PDF of Z^2 .

(b) Find $E(Z^2 + Z + \Phi(Z))$.

34. (S) Let $Z \sim \mathcal{N}(0, 1)$. A measuring device is used to observe Z , but the device can only handle positive values, and gives a reading of 0 if $Z \leq 0$; this is an example of censored data. So assume that $X = ZI_{Z>0}$ is observed rather than Z , where $I_{Z>0}$ is the indicator of $Z > 0$. Find $E(X)$ and $\text{Var}(X)$.

35. Let $Z \sim \mathcal{N}(0, 1)$, and c be a nonnegative constant. Find $E(\max(Z - c, 0))$, in terms of the standard Normal CDF Φ and PDF φ . (This kind of calculation often comes up in quantitative finance.)

Hint: Use LOTUS, and handle the max symbol by adjusting the limits of integration appropriately. As a check, make sure that your answer reduces to $1/\sqrt{2\pi}$ when $c = 0$; this must be the case since we showed in Example 5.4.7 that $E|Z| = \sqrt{2/\pi}$, and we have $|Z| = \max(Z, 0) + \max(-Z, 0)$ so by symmetry

$$E|Z| = E(\max(Z, 0)) + E(\max(-Z, 0)) = 2E(\max(Z, 0))$$

Exponential

36. (S) A post office has 2 clerks. Alice enters the post office while 2 other customers, Bob and Claire, are being served by the 2 clerks. She is next in line. Assume that the time a clerk spends serving a customer has an $\text{Expo}(\lambda)$ distribution.

(a) What is the probability that Alice is the last of the 3 customers to be done being served?

Hint: No integrals are needed.

(b) What is the expected total time that Alice needs to spend at the post office?

37. Let T be the time until a radioactive particle decays, and suppose (as is often done in physics and chemistry) that $T \sim \text{Expo}(\lambda)$.

(a) The half-life of the particle is the time at which there is a 50% chance that the particle has decayed (in statistical terminology, this is the median of the distribution of T). Find the half-life of the particle.

(b) Show that for ϵ a small, positive constant, the probability that the particle decays in the time interval $[t, t + \epsilon]$, given that it has survived until time t , does not depend on t and is approximately proportional to ϵ .

Hint: $e^x \approx 1 + x$ if $x \approx 0$.

(c) Now consider n radioactive particles, with i.i.d. times until decay $T_1, \dots, T_n \sim \text{Expo}(\lambda)$. Let L be the first time at which one of the particles decays. Find the CDF of L . Also, find $E(L)$ and $\text{Var}(L)$.

(d) Continuing (c), find the mean and variance of $M = \max(T_1, \dots, T_n)$, the last time at which one of the particles decays, without using calculus.

Hint: Draw a timeline, apply (c), and remember the memoryless property.

38. (S) Fred wants to sell his car, after moving back to Blissville (where he is happy with the bus system). He decides to sell it to the first person to offer at least \$18,000 for it. Assume that the offers are independent Exponential random variables with mean \$12,000, and that Fred is able to keep getting offers until he obtains one that meets his criterion.

(a) Find the expected number of offers Fred will have.

(b) Find the expected amount of money that Fred will get for the car.

39. As in the previous problem, Fred wants to sell his car, and the offers for his car are i.i.d. Exponential r.v.s with mean \$12,000. Assume now though that he will wait until he has 3 offers (no matter how large or small they are), and then accept the largest of the 3 offers. Find the expected amount of money that Fred will get for his car.

40. (a) Fred visits Blotchville again. He finds that the city has installed an electronic display at the bus stop, showing the time when the previous bus arrived. The times between arrivals of buses are still independent Exponentials with mean 10 minutes. Fred waits for the next bus, and then records the time between that bus and the previous bus. On average, what length of time between buses does he see?

(b) Fred then visits Blunderville, where the times between buses are also 10 minutes on average, and independent. Yet to his dismay, he finds that on average he has to wait more than 1 hour for the next bus when he arrives at the bus stop! How is it possible that the average Fred-to-bus time is greater than the average bus-to-bus time even though Fred arrives at some time between two bus arrivals? Explain this intuitively, and construct a specific discrete distribution for the times between buses showing that this is possible.

41. Fred and Gretchen are waiting at a bus stop in Blotchville. Two bus routes, Route 1 and Route 2, have buses that stop at this bus stop. For Route i , buses arrive according to a Poisson process with rate λ_i buses / minute. The Route 1 process is independent of the Route 2 process. Fred is waiting for a Route 1 bus, and Gretchen is waiting for a Route 2 bus.

(a) Given that Fred has already waited for 20 minutes, on average how much longer will he have to wait for his bus?

(b) Find the probability that at least n Route 1 buses will pass by before the first Route 2 bus arrives. The following result from Chapter 7 may be useful here: for independent random variables $X_1 \sim \text{Expo}(\lambda_1)$, $X_2 \sim \text{Expo}(\lambda_2)$, we have $P(X_1 < X_2) = \lambda_1 / (\lambda_1 + \lambda_2)$.

(c) For this part only, assume that $\lambda_1 = \lambda_2 = \lambda$. Find the expected time it will take until both Fred and Gretchen have caught their buses.

42. (S) Joe is waiting in continuous time for a book called *The Winds of Winter* to be released.

Suppose that the waiting time T until news of the book's release is posted, measured in years relative to some starting point, has an Exponential distribution with $\lambda = 1/5$.

Joe is not so obsessive as to check multiple times a day; instead, he checks the website once at the end of each day. Therefore, he observes the day on which the news was posted, rather than the exact time T . Let X be this measurement, where $X = 0$ means that the news was posted within the first day (after the starting point), $X = 1$ means it was posted on the second day, etc. (assume that there are 365 days in a year). Find the PMF of X . Is this a named distribution that we have studied?

43. The Exponential is the analog of the Geometric in continuous time. This problem explores the connection between Exponential and Geometric in more detail, asking what happens to a Geometric in a limit where the Bernoulli trials are performed faster and faster but with smaller and smaller success probabilities.

Suppose that Bernoulli trials are being performed in continuous time; rather than only thinking about first trial, second trial, etc., imagine that the trials take place at points on a timeline. Assume that the trials are at regularly spaced times $0, \Delta t, 2\Delta t, \dots$, where Δt is a small positive number. Let the probability of success of each trial be $\lambda\Delta t$, where λ is a positive constant. Let G be the number of failures before the first success (in discrete time), and T be the time of the first success (in continuous time).

(a) Find a simple equation relating G to T .

Hint: Draw a timeline and try out a simple example.

(b) Find the CDF of T .

Hint: First find $P(T > t)$.

(c) Show that as $\Delta t \rightarrow 0$, the CDF of T converges to the $\text{Expo}(\lambda)$ CDF, evaluating all the CDFs at a fixed $t \geq 0$.

Hint: Use the compound interest limit (see Section A.2.5 of the math appendix).

44. The Laplace distribution has PDF

$$f(x) = \frac{1}{2}e^{-|x|}$$

for all real x . The Laplace distribution is also called a symmetrized Exponential distribution. Explain this in the following two ways.

(a) Plot the PDFs and explain how they relate.

(b) Let $X \sim \text{Expo}(1)$ and S be a random sign (1 or -1, with equal probabilities), with S and X independent. Find the PDF of SX (by first finding the CDF), and compare the PDF of SX and the Laplace PDF.

45. Emails arrive in an inbox according to a Poisson process with rate 20 emails per hour. Let T be the time at which the 3rd email arrives, measured in hours after a certain fixed starting time. Find $P(T > 0.1)$ without using calculus.

Hint: Apply the count-time duality.

46. Let T be the lifetime of a certain person (how long that person lives), and let T have CDF F and PDF f . The hazard function of T is defined by

$$h(t) = \frac{f(t)}{1 - F(t)}$$

(a) Explain why h is called the hazard function and in particular, why $h(t)$ is the probability density for death at time t , given that the person survived up until then.

(b) Show that an Exponential r.v. has constant hazard function and conversely, if the hazard function of T is a constant then T must be $\text{Expo}(\lambda)$ for some λ .

47. Let T be the lifetime of a person (or animal or gadget), with CDF F and PDF f . Let h be the hazard function, defined as in the previous problem. If we know F then we can calculate

f , and then in turn we can calculate h . In this problem, we consider the reverse problem: how to recover F and f from knowing h .

(a) Show that the CDF and hazard function are related by

$$F(t) = 1 - \exp\left(-\int_0^t h(s)ds\right)$$

for all $t > 0$.

Hint: Let $G(t) = 1 - F(t)$ be the survival function, and consider the derivative of $\log G(t)$.

(b) Show that the PDF and hazard function are related by

$$f(t) = h(t) \exp\left(-\int_0^t h(s)ds\right)$$

for all $t > 0$.

Hint: Apply the result of (a).

48. (S) Find $E(X^3)$ for $X \sim \text{Expo}(\lambda)$, using LOTUS and the fact that $E(X) = 1/\lambda$ and $\text{Var}(X) = 1/\lambda^2$, and integration by parts at most once. In the next chapter, we'll learn how to find $E(X^n)$ for all n .

49. (S) The Gumbel distribution is the distribution of $-\log X$ with $X \sim \text{Expo}(1)$.

(a) Find the CDF of the Gumbel distribution.

(b) Let $X_1, X_2, \dots$ be i.i.d. $\text{Expo}(1)$ and let $M_n = \max(X_1, \dots, X_n)$. Show that $M_n \log n$ converges in distribution to the Gumbel distribution, i.e., as $n \rightarrow \infty$ the CDF of $M_n - \log n$ converges to the Gumbel CDF.

Mixed Practice

50. Explain intuitively why $P(X < Y) = P(Y < X)$ if X and Y are i.i.d., but equality may not hold if X and Y are not independent or not identically distributed.

51. Let X be an r.v. (discrete or continuous) such that $0 \leq X \leq 1$ always holds. Let $\mu = E(X)$.

(a) Show that

$$\text{Var}(X) \leq \mu - \mu^2 \leq \frac{1}{4}$$

Hint: With probability 1, we have $X^2 \leq X$.

(b) Show that there is only one possible distribution for X for which $\text{Var}(X) = 1/4$. What is the name of this distribution?

52. The Rayleigh distribution from Example 5.1.7 has PDF

$$f(x) = xe^{-x^2/2}, \quad x > 0.$$

Let X have the Rayleigh distribution.

(a) Find $E(X)$ without using much calculus, by interpreting the integral in terms of known results about the Normal distribution.

(b) Find $E(X^2)$.

Hint: A nice approach is to use LOTUS and the substitution $u = x^2/2$, and then interpret the resulting integral in terms of known results about the Exponential distribution.

53. (S) Consider an experiment where we observe the value of a random variable X , and estimate the value of an unknown constant θ using some random variable $T = g(X)$ that is a function of X . The r.v. T is called an estimator. Think of X as the data observed in the experiment, and θ as an unknown parameter related to the distribution of X .

For example, consider the experiment of flipping a coin n times, where the coin has an unknown probability θ of Heads. After the experiment is performed, we have observed the value of $X \sim \text{Bin}(n, \theta)$. The most natural estimator for θ is then X/n .

The bias of an estimator T for θ is defined as $b(T) = E(T) - \theta$. The mean squared error is the average squared error when using $T(X)$ to estimate θ :

$$\text{MSE}(T) = E(T - \theta)^2$$

Show that

$$\text{MSE}(T) = \text{Var}(T) + (b(T))^2.$$

This implies that for fixed MSE, lower bias can only be attained at the cost of higher variance and vice versa; this is a form of the bias-variance tradeoff, a phenomenon which arises throughout statistics.

54. (S) (a) Suppose that we have a list of the populations of every country in the world.

Guess, without looking at data yet, what percentage of the populations have the digit 1 as their first digit (e.g., a country with a population of 1,234,567 has first digit 1 and a country with population 89,012,345 does not).

(b) After having done (a), look through a list of populations and count how many start with a 1. What percentage of countries is this? Benford's law states that in a very large variety of real-life data sets, the first digit approximately follows a particular distribution with about a 30% chance of a 1, an 18% chance of a 2, and in general

$$P(D = j) = \log_{10} \left(\frac{j+1}{j} \right), \text{ for } j \in \{1, 2, 3, \dots, 9\}$$

where D is the first digit of a randomly chosen element. (Exercise 6 from Chapter 3 asks for a proof that this is a valid PMF.) How closely does the percentage found in the data agree with that predicted by Benford's law?

(c) Suppose that we write the random value in some problem (e.g., the population of a random country) in scientific notation as $X \times 10^N$, where N is a nonnegative integer and $1 \leq X < 10$. Assume that X is a continuous r.v. with PDF

$$f(x) = c/x, \text{ for } 1 \leq x \leq 10$$

and 0 otherwise, with c a constant. What is the value of c (be careful with the bases of logs)? Intuitively, we might hope that the distribution of X does not depend on the choice of units in which X is measured. To see whether this holds, let $Y = aX$ with $a > 0$. What is the PDF of Y (specifying where it is nonzero)?

(d) Show that if we have a random number $X \times 10^N$ (written in scientific notation) and X has the PDF f from (c), then the first digit (which is also the first digit of X) has Benford's law as its PMF.

Hint: What does $D = j$ correspond to in terms of the values of X ?

55. (S) (a) Let $X_1, X_2, \dots$ be independent $\mathcal{N}(0, 4)$ r.v.s., and let J be the smallest value of j such that $X_j > 4$ (i.e., the index of the first X_j exceeding 4). In terms of Φ , find $E(J)$.

(b) Let f and g be PDFs with $f(x) > 0$ and $g(x) > 0$ for all x . Let X be a random variable with PDF f . Find the expected value of the ratio

$$R = \frac{g(X)}{f(X)}$$

Such ratios come up very often in statistics, when working with a quantity known as a likelihood ratio and when using a computational technique known as importance sampling.

(c) Define

$$F(x) = e^{-e^{-x}}$$

This is a CDF and is a continuous, strictly increasing function. Let X have CDF F , and define $W = F(X)$. What are the mean and variance of W ?

56. Let $X, Y, Z \sim \mathcal{N}(0, 1)$ be i.i.d.

(a) Find an expression for $E(Z^2\Phi(Z))$ as an integral.

(b) Find $P(\Phi(Z) < 2/3)$.

(c) Find $P(X < Y < Z)$.

57. Let $X, Y, Z \sim \mathcal{N}(0, 1)$ be i.i.d., and $W = (\Phi(Z))^2$.

(a) Find the CDF and PDF of W .

(b) Let f_W be the PDF of W and φ be the PDF of Z . Find unsimplified expressions for $E(W^3)$ as integrals in two different ways, one based on f_W and one based on φ .

(c) Find $P(X + 2Y < 2Z + 3)$, in terms of Φ .

Hint: Move all of the r.v.s to one side of the inequality.

58. Let $Z \sim \mathcal{N}(0, 1)$ and $Y = \max(Z, 0)$. So Y is Z if $Z > 0$, and Y is 0 if $Z \leq 0$.

(a) Find an expression for $E(Y)$ as an integral (which can be unsimplified).

(b) Let $Y_1, Y_2, \dots$ be independent r.v.s, each with the same distribution as Y . Let $N = \min\{n : Y_n = 0\}$, i.e., N is the smallest value such that $Y_N = 0$. Find $E(N)$.

(c) Find the CDF of Y in terms of Φ . (Be sure to specify it for all real numbers.)

59. The unit circle $\{(x, y) : x^2 + y^2 = 1\}$ is divided into three arcs by choosing three random points A, B, C on the circle (independently and uniformly), forming arcs between A and B , between A and C , and between B and C . Let L be the length of the arc containing the point $(1, 0)$. What is $E(L)$? Study this by working through the following steps.

(a) Explain what is wrong with the following argument: "The total length of the arcs is 2π , the circumference of the circle. So by symmetry and linearity, each arc has length $2\pi/3$ on average. Referring to the arc containing $(1, 0)$ is just a way to specify one of the arcs (it wouldn't matter if $(1, 0)$ were replaced by $(0, -1)$ or any other specific point on the circle in the statement of the problem). So the expected value of L is $2\pi/3$."

(b) Let the arc containing $(1, 0)$ be divided into two pieces: the piece extending counterclockwise from $(1, 0)$ and the piece extending clockwise from $(1, 0)$. Write $L = L_1 + L_2$, where L_1 and L_2 are the lengths of the counterclockwise and clockwise pieces, respectively. Find the CDF, PDF, and expected value of L_1 .

(c) Use (b) to find $E(L)$.

60. (S) As in Example 5.7.4, athletes compete one at a time at the high jump. Let X_j be how high the j th jumper jumped, with $X_1, X_2, \dots$ i.i.d. with a continuous distribution. We say that the j th jumper is "best in recent memory" if they jump higher than the previous 2 jumpers (for $j \geq 3$; the first 2 jumpers don't qualify).

(a) Find the expected number of best in recent memory jumpers among the 3rd through n th jumpers.

(b) Let A_j be the event that the j th jumper is the best in recent memory. Find $P(A_3)$, $P(A_4)$, and $P(A_3 \cap A_4)$. Are A_3 and A_4 independent?

61. Tyrion, Cersei, and n other guests arrive at a party at i.i.d. times drawn from a continuous distribution with support $[0, 1]$, and stay until the end (time 0 is the party's start time and time 1 is the end time). The party will be boring at times when neither Tyrion nor Cersei is there, fun when exactly one of them is there, and awkward when both Tyrion and Cersei are there.

(a) On average, how many of the n other guests will arrive at times when the party is fun?

(b) Jaime and Robert are two of the other guests. By computing both sides in the definition of independence, determine whether the event "Jaime arrives at a fun time" is independent of the event "Robert arrives at a fun time".

(c) Give a clear intuitive explanation of whether the two events from (b) are independent, and whether they are conditionally independent given the arrival times of everyone else, i.e., everyone except Jaime and Robert.

62. Let $X_1, X_2, \dots$ be the annual rainfalls in Boston (measured in inches) in the years 2101, 2102, ..., respectively. Assume that annual rainfalls are i.i.d. draws from a continuous distribution. A rainfall value is a record high if it is greater than those in all previous years (starting with 2101), and a record low if it is lower than those in all previous years.

(a) In the 22nd century (the years 2101 through 2200, inclusive), find the expected number of years that have either a record low or a record high rainfall.

(b) On average, in how many years in the 22nd century is there a record low followed in the next year by a record high? (Only the record low is required to be in the 22nd century, not the record high.)

(c) By definition, the year 2101 is a record high (and record low). Let N be the number of years required to get a new record high. Find $P(N > n)$ for all positive integers n , and use this to find the PMF of N .

Hint: Note that $P(N = n) + P(N > n) = P(N > n - 1)$.

(d) With notation as above, show that $E(N)$ is infinite.

4.6 6 Moments

The n th moment of an r.v. X is $E(X^n)$. In this chapter, we explore how the moments of an r.v. shed light on its distribution. We have already seen that the first two moments are useful since they provide the mean $E(X)$ and variance $E(X^2) - (EX)^2$, which are important summaries of the average value of X and how spread out its distribution is. But there is much more to a distribution than its mean and variance. We'll see that the third and fourth moments tell us about the asymmetry of a distribution and the behavior of the tails or extreme values, two properties that are not captured by the mean and variance. After introducing moments, we'll discuss the moment generating function (MGF), which not only helps us compute moments but also provides a useful alternative way to specify a distribution.

4.6.1 6.1 Summaries of a Distribution

The mean is called a measure of central tendency because it tells us something about the center of a distribution, specifically its center of mass. Other measures of central tendency that are commonly used in statistics are the median and the mode, which we now define.

Definition 6.1.1 (Median). We say that c is a median of a random variable X if $P(X \leq c) \geq 1/2$ and $P(X \geq c) \geq 1/2$. (The simplest way this can happen is if the CDF of X hits $1/2$ exactly at c , but we know that some CDFs have jumps.)

Definition 6.1.2 (Mode). For a discrete r.v. X , we say that c is a mode of X if it maximizes the PMF: $P(X = c) \geq P(X = x)$ for all x . For a continuous r.v. X with PDF f , we say that c is a mode if it maximizes the PDF: $f(c) \geq f(x)$ for all x .

As with the mean, the median and mode of an r.v. depend only on its distribution, so we can talk about the mean, median, or mode of a distribution without referring to any particular r.v. that has that distribution. For example, if $Z \sim \mathcal{N}(0, 1)$ then the median of Z is 0 (since $\Phi(0) = 1/2$ by symmetry), and we also say that the standard Normal distribution has median 0.

Intuitively, the median is a value c such that half the mass of the distribution falls on either

side of c (or as close to half as possible, for discrete r.v.s), and the mode is a value that has the greatest mass or density out of all values in the support of

X . If the CDF F is a continuous, strictly increasing function, then $F^{-1}(1/2)$ is the median (and is unique).

Note that a distribution can have multiple medians and multiple modes. Medians have to occur side by side; modes can occur all over the distribution. In Figure 6.1, we show a distribution supported on $[-5, -1] \cup [1, 5]$ that has two modes and infinitely many medians. The PDF is 0 between -1 and 1, so all values between -1 and 1 are medians of the distribution because half of the mass falls on either side. The two modes are at -3 and 3.




FIGURE 6.1 A distribution with two modes (-3 and 3) and infinitely many medians (all x in the interval $[-1, 1]$).

Example 6.1.3 (Mean, median, and mode for salaries). A certain company has 100 employees. Let $s_1, s_2, \dots, s_{100}$ be their salaries, sorted in increasing order (we can still do this even if some salaries appear more than once). Let X be the salary of a randomly selected employee (chosen uniformly). The mean, median, and mode for the data set $s_1, s_2, \dots, s_{100}$ are defined to be the corresponding quantities for X .

What is a typical salary? What is the most useful one-number summary of the salary data? The answer, as is often the case, is it depends on the goal. Different summaries reveal different characteristics of the data, so it may be hard to choose just one number-and it is often unnecessary to do so, since usually we can provide several summaries (and plot the data too). Here we briefly compare the mean, median, and mode, though often it makes sense to report all three (and other summaries too).

If the salaries are all different, the mode doesn't give a useful one-number summary since there are 100 modes. If there are only a few possible salaries in the company, the mode becomes more useful. But even then it could happen that, for example, 34 people receive salary a , 33 receive salary b , and 33 receive salary c . Then a is the unique mode, but if we only report a we are ignoring b and c , which just barely missed being modes and which together account for almost 2/3 of the data.

Next, let's consider the median. There are two numbers "in the middle", s_{50} and s_{51} . In fact, any number m with $s_{50} \leq m \leq s_{51}$ is a median, since there is at least a 50% chance that the random employee's salary is in $\{s_1, \dots, s_{50}\}$ (in which case it is at most m) and at least a 50% chance that it is in $\{s_{51}, \dots, s_{100}\}$ (in which case it is at least m). The usual convention is to choose $m = (s_{50} + s_{51})/2$, the mean of the two numbers in the middle. If the number of employees had been odd, this issue would not have come up; in that case, there is a unique median, the number in the middle when all the salaries are listed in increasing order.

Compared with the mean, the median is much less sensitive to extreme values. For example, if the CEO's salary is changed from being slightly more than anyone else's to vastly more than anyone else's, that could have a large impact on the mean but it has no impact on the median. This robustness is a reason that the median could be a more sensible summary than the mean of what the typical salary is. On the other hand, suppose that we want to know the total cost

the company is paying for its employees' salaries. If we only know a mode or a median, we can't extract this information, but if we know the mean we can just multiply it by 100 .

Suppose that we are trying to guess what a not-yet-observed r.v. X will be, by making a prediction c . The mean and the median both seem like natural guesses for c , but which is better? That depends on how "better" is defined. Two natural ways to judge how good c is are the mean squared error $E(X - c)^2$ and the mean absolute error $E|X - c|$. The following result says what the best guesses are in both cases.

Theorem 6.1.4. Let X be an r.v. with mean μ , and let m be a median of X .

- The value of c that minimizes the mean squared error $E(X - c)^2$ is $c = \mu$.
- A value of c that minimizes the mean absolute error $E|X - c|$ is $c = m$.

Proof. We will first prove a useful identity:

$$E(X - c)^2 = \text{Var}(X) + (\mu - c)^2$$

This can be checked by expanding everything out, but a faster way is to use the fact that adding a constant doesn't affect the variance:

$$\text{Var}(X) = \text{Var}(X - c) = E(X - c)^2 - (E(X - c))^2 = E(X - c)^2 - (\mu - c)^2.$$

It follows that $c = \mu$ is the unique choice that minimizes $E(X - c)^2$, since that choice makes $(\mu - c)^2 = 0$ and any other choice makes $(\mu - c)^2 > 0$.

Next, let's consider the mean absolute error. Let $a \neq m$. We need to show that $E|X - m| \leq E|X - a|$, which is equivalent to $E(|X - a| - |X - m|) \geq 0$. Assume that $m < a$ (the case $m > a$ can be handled similarly). If $X \leq m$ then

$$|X - a| - |X - m| = a - X - (m - X) = a - m$$

and if $X > m$ then

$$|X - a| - |X - m| \geq X - a - (X - m) = m - a.$$

Let

$$Y = |X - a| - |X - m|.$$

We can split the definition of $E(Y)$ into 2 parts based on whether $X \leq m$ occurs, using indicator r.v.s. Let I be the indicator r.v. for $X \leq m$, so $1 - I$ is the indicator r.v. for $X > m$. Then

$$\begin{aligned} E(Y) &= E(YI) + E(Y(1 - I)) \\ &\geq (a - m)E(I) + (m - a)E(1 - I) \\ &= (a - m)P(X \leq m) + (m - a)P(X > m) \\ &= (a - m)P(X \leq m) - (a - m)(1 - P(X \leq m)) \\ &= (a - m)(2P(X \leq m) - 1) \end{aligned}$$

By definition of median, we have $2P(X \leq m) - 1 \geq 0$. Thus, $E(Y) \geq 0$, which implies $E(|X - m|) \leq E(|X - a|)$. $\square$




FIGURE 6.2

Two PDFs with mean 2 and variance 12. The light curve is Normal and symmetric; the dark curve is Log-Normal and right-skewed.

Regardless of which measure of central tendency we use in a particular application, it is usually important also to know about the spread of the distribution as measured by, for example, the variance. However, there are also major features of a distribution that are not captured by the mean and variance. For example, the two PDFs in Figure 6.2 both have mean 2 and variance 12. The light curve is the $\mathcal{N}(2, 12)$ PDF and the dark curve belongs to the Log-Normal family of distributions (the LogNormal is defined later in this chapter). The Normal curve is symmetric about 2, so

its mean, median, and mode are all 2. In contrast, the Log-Normal is heavily skewed to the right; this means its right tail is very long compared to its left tail. It has mean 2, but median 1 and mode 0.25. From the mean and variance alone, we would not be able to capture the difference between the asymmetry of the Log-Normal and the symmetry of the Normal.

Now consider Figure 6.3, which shows the PMF of a $\text{Bin}(10, 0.9)$ r.v. on the left and the PMF of 8 plus a $\text{Bin}(10, 0.1)$ r.v. on the right. Both of these distributions have mean, median, and mode equal to 9 and variance equal to 0.9, but they look drastically different. We say that the PMF on the left is left-skewed and the PMF on the right is right-skewed. In this chapter we'll learn that a standard measure of the asymmetry of a distribution is based on the third moment.




FIGURE 6.3

Left: $\text{Bin}(10, 0.9)$ is left-skewed. Right: $\text{Bin}(10, 0.1)$, shifted to the right by 8, is right-skewed but has the same mean, median, mode, and variance as $\text{Bin}(10, 0.9)$.

The previous two examples considered asymmetric distributions, but symmetric distributions with the same mean and variance can also look very different. The left plot in Figure 6.4 shows two symmetric PDFs with mean 0 and variance 1. The light curve is the $\mathcal{N}(0, 1)$ PDF

and the dark curve is the t_3 PDF, scaling it to have variance 1 (we'll define the t_3 distribution in Chapter 10). The dark curve has a sharper peak and heavier tails than the light curve. The tail behavior is magnified in the right plot, where it is easy to see that the dark curve decays to 0 much more gradually, making outcomes far out in the tail much more likely than for the light curve. As we'll learn in this chapter, a standard measure of the heaviness of the tails of a distribution is based on the fourth moment.

In the next section, we will go into further detail of how to interpret moments, especially the first four moments.




FIGURE 6.4

Left: $\mathcal{N}(0,1)$ PDF (light) and a scaled t_3 PDF (dark). Both have mean 0 and variance 1, but the latter has a sharper peak and heavier tails. Right: magnified version of right tail behavior.

4.6.2 6.2 Interpreting Moments

Definition 6.2.1 (Kinds of moments). Let X be an r.v. with mean μ and variance σ^2 . For any positive integer n , the n th moment of X is $E(X^n)$, the n th central moment is $E((X - \mu)^n)$, and the n th standardized moment is $E\left(\left(\frac{X - \mu}{\sigma}\right)^n\right)$. Throughout the previous sentence, "if it exists" is left implicit.

In particular, the mean is the first moment and the variance is the second central moment. The term moment is borrowed from physics. Let X be a discrete r.v. with distinct possible values $x_1, \dots, x_n$, and imagine a pebble with mass $m_j = P(X = x_j)$ positioned at x_j on a number line, for each j (as illustrated in Figure 6.5). In physics,

$$E(X) = \sum_{j=1}^n m_j x_j$$

is called the center of mass of the system, and

$$\text{Var}(X) = \sum_{j=1}^n m_j (x_j - E(X))^2$$

is called the moment of inertia about the center of mass.

We'll now define skewness, a single-number summary of asymmetry which, as alluded to earlier, is based on the third moment. In fact, skewness is defined to be the third standardized moment.




FIGURE 6.5 Physics interpretation of moments. The mean (first moment) of an r.v. corresponds to the center of mass of a collection of pebbles, and the variance (second central moment) corresponds to the moment of inertia about the center of mass.

Definition 6.2.2 (Skewness). The skewness of an r.v. X with mean μ and variance σ^2 is the third standardized moment of X :

$$\text{Skew}(X) = E\left(\frac{X - \mu}{\sigma}\right)^3.$$

By standardizing first, we make the definition of $\text{Skew}(X)$ not depend on the location or the scale of X , which is reasonable since we already have μ and σ to provide information about location and scale. Also, standardizing first gives the nice property that the units in which X is measured (e.g., inches vs. meters) will not affect the skewness.

To understand how skewness measures asymmetry, we first need to discuss what it means for an r.v. to be symmetric.

Definition 6.2.3 (Symmetry of an r.v.). We say that an r.v. X has a symmetric distribution about μ if $X - \mu$ has the same distribution as $\mu - X$. We also say that X is symmetric or that the distribution of X is symmetric; these all have the same meaning.

The number μ in the above definition must be $E(X)$ if the mean exists, since

$$E(X) - \mu = E(X - \mu) = E(\mu - X) = \mu - E(X)$$

simplifies to $E(X) = \mu$. Because of this, it is common to say " X is symmetric" as shorthand for " X is symmetric about its mean" (if the mean exists). The number μ is also a median of the distribution, since if $X - \mu$ has the same distribution as $\mu - X$, then

$$P(X - \mu \leq 0) = P(\mu - X \leq 0),$$

so

$$P(X \leq \mu) = P(X \geq \mu),$$

which implies that

$$P(X \leq \mu) = 1 - P(X > \mu) \geq 1 - P(X \geq \mu) = 1 - P(X \leq \mu),$$

showing that $P(X \leq \mu) \geq 1/2$ and $P(X \geq \mu) \geq 1/2$.

6.2.4. Sometimes people say " X is symmetric" to mean " X is symmetric about 0 ". Note that if X is symmetric about μ , then $X - \mu$ is symmetric about 0 . Symmetry about 0 is especially convenient since then $-X$ and X have the same distribution, and the PDF of X (if X is continuous) is an even function, as shown below.

Intuitively, symmetry means that the PDF of X to the left of μ is the mirror image of the PDF of X to the right of μ (for X continuous, and the same holds for the PMF if X is discrete). For example, we have seen before that $X \sim \mathcal{N}(\mu, \sigma^2)$ is symmetric; in terms of the definition, this is because $X - \mu$ and $\mu - X$ are both $\mathcal{N}(0, \sigma^2)$. We have also seen from Corollary 3.3.8 that $X \sim \text{Bin}(n, p)$ is symmetric when $p = 1/2$.

We can also give an algebraic description of what the PDF of a symmetric continuous r.v. looks like.

Proposition 6.2.5 (Symmetry in terms of the PDF). Let X be a continuous r.v. with PDF f . Then X is symmetric about μ if and only if $f(x) = f(2\mu - x)$ for all x .

Proof. Let F be the CDF of X . If symmetry holds, we have $F(x) = P(X - \mu \leq x - \mu) = P(\mu - X \leq x - \mu) = P(X \geq 2\mu - x) = 1 - F(2\mu - x)$. Taking the derivative of both sides yields $f(x) = f(2\mu - x)$. Conversely, suppose that $f(x) = f(2\mu - x)$ holds for all x . Integrating both sides, we have

$$P(X - \mu \leq t) = P(X \leq \mu + t) = \int_{-\infty}^{\mu+t} f(x)dx = \int_{-\infty}^{\mu+t} f(2\mu - x)dx$$

which, after the substitution $w = 2\mu - x$, becomes

$$\int_{\mu-t}^{\infty} f(w)dw = P(X \geq \mu - t) = P(\mu - X \leq t)$$

Odd central moments give some information about symmetry.

Proposition 6.2.6 (Odd central moments of a symmetric distribution). Let X be symmetric about its mean μ . Then for any odd number m , the m th central moment $E(X - \mu)^m$ is 0 if it exists.

Proof. Since $X - \mu$ has the same distribution as $\mu - X$, they have the same m th moment (if it exists):

$$E(X - \mu)^m = E(\mu - X)^m$$

Let $Y = (X - \mu)^m$. Then $(\mu - X)^m = -(X - \mu)^m = (-1)^m Y = -Y$, so the above equation just says $E(Y) = -E(Y)$. So $E(Y) = 0$.

This leads us to consider using an odd standardized moment as a measure of the skew of a distribution. The first standardized moment is always 0, so the third standardized moment is taken as the definition of skewness. Positive skewness is indicative of having a long right tail relative to the left tail, and negative skewness is indicative of the reverse. (The converse of the above proposition is false, though: there exist asymmetric distributions whose odd central moments are all 0.)

Why not use, say, the fifth standardized moment instead of the third? One reason is that the third standardized moment is usually easier to calculate. Another reason is that we may want to estimate skewness from a data set. It is usually easier to estimate lower moments than higher moments in a stable way since, for example, a large, noisy observation will have a very large, very noisy fifth power. Nevertheless, just as the mean isn't the only useful notion of average and the variance isn't the only useful notion of spread, the third standardized moment isn't the only useful notion of skew.

Another important descriptive feature of a distribution is how heavy its tails are. For a given variance, is the variability explained more by a few rare (extreme) events or by a moderate number of moderate deviations from the mean? This is an important consideration for risk management in finance: for many financial assets, the distribution of returns has a heavy left tail caused by rare but severe crisis events, and failure to account for these rare events can have disastrous consequences, as demonstrated by the 2008 financial crisis.

As with measuring skew, no single measure can perfectly capture the tail behavior, but there is a widely used summary based on the fourth standardized moment.

Definition 6.2.7 (Kurtosis). The kurtosis of an r.v. X with mean μ and variance σ^2 is a shifted version of the fourth standardized moment of X :

$$\text{Kurt}(X) = E\left(\frac{X - \mu}{\sigma}\right)^4 - 3$$

6.2.8. The reason for subtracting 3 is that this makes any Normal distribution have kurtosis 0 (as shown in Section 6.5). This provides a convenient basis for comparison. However, some sources define the kurtosis without the 3, in which case they call our version "excess kurtosis".

Figure 6.6 shows three named distributions and lists the skewness and kurtosis of each. The Expo(1) and Pois(4) distributions (left and middle) both have positive skewness and positive kurtosis, indicating that they are right-skewed and their tails are heavier than those of a Normal distribution. The Unif(0,1) distribution (right) has zero skewness and negative kurtosis: zero skewness because the distribution is symmetric about its mean, and negative kurtosis because it has no tails!




FIGURE 6.6

Skewness and kurtosis of some named distributions. Left: $\text{Expo}(1)$ PDF, skewness = 2, kurtosis = 6. Middle: $\text{Pois}(4)$ PMF, skewness = 0.5, kurtosis = 0.25. Right: $\text{Unif}(0, 1)$ PDF, skewness = 0, kurtosis = -1.2 .

4.6.3 6.3 Sample Moments

In statistical inference, a central problem is how to use data to estimate unknown parameters of a distribution, or functions of unknown parameters. It is especially common to want to estimate the mean and variance of a distribution. If the data are i.i.d. random variables $X_1, \dots, X_n$ where the mean $E(X_j)$ is unknown, then the most obvious way to estimate the mean is simply to average the X_j , taking the arithmetic mean.

For example, if the observed data are 3, 1, 1, 5, then a simple, natural way to estimate the mean of the distribution that generated the data is to use $(3 + 1 + 1 + 5)/4 = 2.5$. This is called the sample mean. Similarly, if we want to estimate the second moment of the distribution that generated the data 3, 1, 1, 5, then a simple, natural way is to use $(3^2 + 1^2 + 1^2 + 5^2)/4 = 9$. This is called the sample second moment. In general, sample moments are defined as follows.

Definition 6.3.1 (Sample moments). Let $X_1, \dots, X_n$ be i.i.d. random variables. The k th sample moment is the r.v.

$$M_k = \frac{1}{n} \sum_{j=1}^n X_j^k.$$

The sample mean $\bar{X}_n$ is the first sample moment:

$$\bar{X}_n = \frac{1}{n} \sum_{j=1}^n X_j.$$

In contrast, the population mean or true mean is $E(X_j)$, the mean of the distribution from which the X_j were drawn.

The law of large numbers, which we prove in Chapter 10, shows that the k th sample moment of i.i.d. random variables $X_1, \dots, X_n$ converges to the k th moment $E(X_1^k)$ as $n \rightarrow \infty$. Also, the expected value of the k th sample moment is the k th moment. In statistical terms, we say that the k th sample moment is unbiased for estimating the k th moment. It is easy to check this by linearity:

$$E\left(\frac{1}{n} \sum_{j=1}^n X_j^k\right) = \frac{1}{n} (E(X_1^k) + \dots + E(X_n^k)) = E(X_1^k).$$

The mean and variance of the sample mean are often needed, and have nice expressions that are often needed in statistics.

Theorem 6.3.2 (Mean and variance of sample mean). Let $X_1, \dots, X_n$ be i.i.d. r.v.s with mean μ and variance σ^2 . Then the sample mean $\bar{X}_n$ is unbiased for estimating μ . That is,

$$E(\bar{X}_n) = \mu$$

The variance of $\bar{X}_n$ is given by

$$\text{Var}(\bar{X}_n) = \frac{\sigma^2}{n}$$

Proof. We have $E(\bar{X}_n) = \mu$ since we showed above that the k th sample moment is unbiased for estimating the k th moment. For the variance, we will use the fact (shown in the next chapter) that the variance of the sum of independent r.v.s is the sum of the variances:

$$\text{Var}(\bar{X}_n) = \frac{1}{n^2} \text{Var}(X_1 + \dots + X_n) = \frac{n}{n^2} \text{Var}(X_1) = \frac{\sigma^2}{n}$$

For estimating the variance of the distribution of i.i.d. r.v.s $X_1, \dots, X_n$, a natural approach building on the above concepts is to mimic the formula $\text{Var}(X) = E(X^2) - (EX)^2$ by taking the second sample moment and subtracting the square of the sample mean. There are advantages to this method, but a more common method is as follows.

Definition 6.3.3 (Sample variance and sample standard deviation). Let $X_1, \dots, X_n$ be i.i.d. random variables. The sample variance is the r.v.

$$S_n^2 = \frac{1}{n-1} \sum_{j=1}^n (X_j - \bar{X}_n)^2$$

The sample standard deviation is the square root of the sample variance.

The idea of the above definition is to mimic the formula $\text{Var}(X) = E(X - E(X))^2$ by averaging the squared distances of the X_j from the sample mean, except with $n-1$ rather than n in the denominator. The motivation for the $n-1$ is that this makes

the sample variance S_n^2 unbiased for estimating σ^2 , i.e., it is correct on average. However, the sample standard deviation S_n is not unbiased for estimating σ ; we will see in Chapter 10 which way the inequality goes. In any case, unbiasedness is only one of several criteria by which to judge an estimation procedure. For example, in some problems we can get a lower mean squared error in return for allowing a little bit of bias, and this tradeoff may be worthwhile.

Theorem 6.3.4 (Unbiasedness of sample variance). Let $X_1, \dots, X_n$ be i.i.d. r.v.s with mean μ and variance σ^2 . Then the sample variance S_n^2 is unbiased for estimating σ^2 , i.e.,

$$E(S_n^2) = \sigma^2$$

Proof. The key to the proof is the handy identity

$$\sum_{j=1}^n (X_j - c)^2 = \sum_{j=1}^n (X_j - \bar{X}_n)^2 + n (\bar{X}_n - c)^2$$

which holds for all c . To verify the identity, add and subtract $\bar{X}_n$ in the left-hand sum:

$$\begin{aligned} \sum_{j=1}^n (X_j - c)^2 &= \sum_{j=1}^n ((X_j - \bar{X}_n) + (\bar{X}_n - c))^2 \\ &= \sum_{j=1}^n (X_j - \bar{X}_n)^2 + 2 \sum_{j=1}^n (X_j - \bar{X}_n) (\bar{X}_n - c) + \sum_{j=1}^n (\bar{X}_n - c)^2 \\ &= \sum_{j=1}^n (X_j - \bar{X}_n)^2 + n (\bar{X}_n - c)^2 \end{aligned}$$

For the last line, we used the fact that $\bar{X}_n - c$ does not depend on j and the fact that

$$\sum_{j=1}^n (X_j - \bar{X}_n) = \sum_{j=1}^n X_j - \sum_{j=1}^n \bar{X}_n = n\bar{X}_n - n\bar{X}_n = 0.$$

Now let us apply the identity, choosing $c = \mu$. Taking the expectation of both sides,

$$nE(X_1 - \mu)^2 = E\left(\sum_{j=1}^n (X_j - \bar{X}_n)^2\right) + nE(\bar{X}_n - \mu)^2.$$

By definition of variance, $E(X_1 - \mu)^2 = \text{Var}(X_1) = \sigma^2$, and $E(\bar{X}_n - \mu)^2 = \text{Var}(\bar{X}_n) = \sigma^2/n$. Plugging these results in above and simplifying, we have $E(S_n^2) = \sigma^2$.

Similarly, we can define the sample skewness to be

$$\frac{\frac{1}{n} \sum_{j=1}^n (X_j - \bar{X}_n)^3}{S_n^3},$$

and the sample kurtosis to be

$$\frac{\frac{1}{n} \sum_{j=1}^n (X_j - \bar{X}_n)^4}{S_n^4} - 3.$$

Beyond the fourth moment, it rapidly gets harder to interpret moments graphically and harder to estimate them well from data if they are unknown. However, in the rest of this chapter we will see that it can still be useful to know all the moments of a distribution. We will also study a way of computing moments that is often easier than LOTUS. Both the usefulness and the computation of moments are closely connected to a blueprint called the moment generating function, to which we devote most of the rest of this chapter.

4.6.4 6.4 Moment Generating Functions

A generating function is a clothesline on which we hang up a sequence of numbers for display.
- Herbert Wilf [28]

Generating functions are a powerful tool in combinatorics and probability, bridging between sequences of numbers and the world of calculus. In probability, they are useful for studying both discrete and continuous distributions. The general idea behind a generating function is as follows: starting with a sequence of numbers, create a continuous function - the generating

function - that encodes the sequence. We then have all the tools of calculus at our disposal for manipulating the generating function.

A moment generating function, as its name suggests, is a generating function that encodes the moments of a distribution. Here is the definition, followed by a few examples.

Definition 6.4.1 (Moment generating function). The moment generating function (MGF) of an r.v. X is $M(t) = E(e^{tX})$, as a function of t , if this is finite on some open interval $(-a, a)$ containing 0. Otherwise we say the MGF of X does not exist.

A natural question at this point is "What is the interpretation of t ?" The answer is that t has no interpretation in particular; it's just a bookkeeping device that we introduce in order to be able to use calculus instead of working with a discrete sequence of moments.

Note that $M(0) = 1$ for any valid MGF M ; whenever you compute an MGF, plug in 0 and see if you get 1, as a quick check!

Example 6.4.2 (Bernoulli MGF). For $X \sim \text{Bern}(p)$, e^{tX} takes on the value e^t with probability p and the value 1 with probability q , so $M(t) = E(e^{tX}) = pe^t + q$. Since this is finite for all values of t , the MGF is defined on the entire real line.

Example 6.4.3 (Geometric MGF). For $X \sim \text{Geom}(p)$,

$$M(t) = E(e^{tX}) = \sum_{k=0}^{\infty} e^{tk} q^k p = p \sum_{k=0}^{\infty} (qe^t)^k = \frac{p}{1 - qe^t}$$

for $qe^t < 1$, i.e., for t in $(-\infty, \log(1/q))$, which is an open interval containing 0.

Example 6.4.4 (Uniform MGF). Let $U \sim \text{Unif}(a, b)$. Then the MGF of U is

$$M(t) = E(e^{tU}) = \frac{1}{b-a} \int_a^b e^{tu} du = \frac{e^{tb} - e^{ta}}{t(b-a)}$$

for $t \neq 0$, and $M(0) = 1$.

The next three theorems give three reasons why the MGF is important. First, the MGF encodes the moments of an r.v. Second, the MGF of an r.v. determines its distribution, like the CDF and PMF/PDF. Third, MGFs make it easy to find the distribution of a sum of independent r.v.s. Let's take these one by one.

Theorem 6.4.5 (Moments via derivatives of the MGF). Given the MGF of X , we can get the n th moment of X by evaluating the n th derivative of the MGF at 0: $E(X^n) = M^{(n)}(0)$.

Proof. This can be seen by noting that the Taylor expansion of $M(t)$ about 0 is

$$M(t) = \sum_{n=0}^{\infty} M^{(n)}(0) \frac{t^n}{n!},$$

while on the other hand, we also have

$$M(t) = E(e^{tX}) = E\left(\sum_{n=0}^{\infty} X^n \frac{t^n}{n!}\right).$$

We are allowed to interchange the expectation and the infinite sum because certain technical conditions are satisfied (this is where we invoke the condition that $E(e^{tX})$ is finite in an interval around 0), so

$$M(t) = \sum_{n=0}^{\infty} E(X^n) \frac{t^n}{n!}.$$

Matching the coefficients of the two expansions, we get $E(X^n) = M^{(n)}(0)$.

The above theorem is surprising in that for a continuous r.v. X , to compute moments would

seemingly require doing integrals with LOTUS, but with the MGF it is possible to find moments by taking derivatives rather than doing integrals!

Theorem 6.4.6 (MGF determines the distribution). The MGF of a random variable determines its distribution: if two r.v.s have the same MGF, they must have the same distribution. In fact, if there is even a tiny interval $(-a, a)$ containing 0 on which the MGFs are equal, then the r.v.s must have the same distribution.

The above theorem is a difficult result in analysis, so we will not prove it here.

Theorem 6.4.7 (MGF of a sum of independent r.v.s). If X and Y are independent, then the MGF of $X + Y$ is the product of the individual MGFs:

$$M_{X+Y}(t) = M_X(t)M_Y(t)$$

This is true because if X and Y are independent, then $E(e^{t(X+Y)}) = E(e^{tX})E(e^{tY})$ (this follows from results discussed in Chapter 7). Using this fact, we can get the MGFs of the Binomial and Negative Binomial, which are sums of independent Bernoullis and Geometrics, respectively.

Example 6.4.8 (Binomial MGF). The MGF of a $\text{Bern}(p)$ r.v. is $pe^t + q$, so the MGF of a $\text{Bin}(n, p)$ r.v. is

$$M(t) = (pe^t + q)^n.$$




FIGURE 6.7

Bin

(2,1 / 2) MGF, $M(t) = (\frac{1}{2}e^t + \frac{1}{2})^2$. The slope of the MGF at $t = 0$ is 1, so the mean of the distribution is 1. The concavity or second derivative of the MGF at $t = 0$ is $3/2$, so the second moment of the distribution is $3/2$.

Figure 6.7 plots the MGF of the $\text{Bin}(2, 1/2)$ distribution, $M(t) = (\frac{1}{2}e^t + \frac{1}{2})^2$, between $t = -1$ and $t = 1$. As with all MGFs, the value of the MGF is 1 at $t = 0$. The first and second moments of the distribution are the first and second derivatives of the MGF, evaluated at $t = 0$; these correspond to the slope and concavity of the plotted curve at $t = 0$. These two derivatives are 1 and $3/2$, corresponding to the fact that the $\text{Bin}(2, 1/2)$ distribution has mean 1 and variance $3/2 - 1^2 = 1/2$. $\square$

Example 6.4.9 (Negative Binomial MGF). We know the MGF of a $\text{Geom}(p)$ r.v. is $\frac{p}{1-qe^t}$ for $qe^t < 1$, so the MGF of $X \sim \text{NBin}(r, p)$ is

$$M(t) = \left(\frac{p}{1-qe^t} \right)^r, \text{ for } qe^t < 1.$$

$\square$

6.4.10. Not all r.v.s have an MGF. Some r.v.s X don't even have $E(X)$ exist, or don't have $E(X^n)$ exist for some $n > 1$, in which case the MGF clearly will not exist. But even if all the moments of X exist, the MGF may not exist if the moments grow too quickly. Luckily, there is a way to fix this: inserting an imaginary number! The function $\psi(t) = E(e^{itX})$ with $i = \sqrt{-1}$ is called the characteristic function by statisticians and the Fourier transform by everyone else. It turns out that the characteristic function always exists. In this book we will focus on the MGF rather than the characteristic function, to avoid having to handle imaginary numbers.

As we saw in the previous chapter, location and scale transformations are a fundamental way to build a family of distributions from an initial distribution. For example, starting with $Z \sim \mathcal{N}(0, 1)$, we can scale by σ and shift by μ to obtain $X = \mu + \sigma Z \sim \mathcal{N}(\mu, \sigma^2)$. In general, if we have an r.v. X with mean μ and standard deviation $\sigma > 0$, we can create the standardized version $(X - \mu)/\sigma$, and vice versa. Conveniently, it is easy to relate the MGFs of two r.v.s connected by such a transformation.

Proposition 6.4.11 (MGF of location-scale transformation). If X has MGF $M(t)$, then the MGF of $a + bX$ is

$$E(e^{t(a+bX)}) = e^{at} E(e^{btX}) = e^{at} M(bt).$$

For example, let's use this proposition to help obtain the MGFs of the Normal and Exponential distributions.

Example 6.4.12 (Normal MGF). The MGF of a standard Normal r.v. Z is

$$M_Z(t) = E(e^{tZ}) = \int_{-\infty}^{\infty} e^{tz} \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz$$

After completing the square, we have

$$M_Z(t) = e^{t^2/2} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-(z-t)^2/2} dz = e^{t^2/2}$$

since the $\mathcal{N}(t, 1)$ PDF integrates to 1. Thus, the MGF of $X = \mu + \sigma Z \sim \mathcal{N}(\mu, \sigma^2)$ is

$$M_X(t) = e^{\mu t} M_Z(\sigma t) = e^{\mu t} e^{(\sigma t)^2/2} = e^{\mu t + \frac{1}{2}\sigma^2 t^2}$$

Example 6.4.13 (Exponential MGF). The MGF of $X \sim \text{Expo}(1)$ is

$$M(t) = E(e^{tX}) = \int_0^{\infty} e^{tx} e^{-x} dx = \int_0^{\infty} e^{-x(1-t)} dx = \frac{1}{1-t} \quad \text{for } t < 1$$

So the MGF of $Y = X/\lambda \sim \text{Expo}(\lambda)$ is

$$M_Y(t) = M_X\left(\frac{t}{\lambda}\right) = \frac{\lambda}{\lambda - t} \quad \text{for } t < \lambda$$

4.6.5 6.5 Generating Moments with Mgf's

We now give some examples of using an MGF to find moments. Theorem 6.4.5 shows that we can get moments by differentiating the MGF and evaluating at 0, rather than doing a complicated sum or integral by LOTUS. Better yet, in some cases we can simultaneously find all the moments of a distribution via a Taylor expansion, rather than differentiating over and over again.

Example 6.5.1 (Exponential moments). In this example we will show how to use the Exponential MGF to get all the moments of the Exponential distribution simultaneously! Let $X \sim \text{Expo}(1)$. The MGF of X is $M(t) = 1/(1-t)$ for $t < 1$.

As shown in Theorem 6.4.5, we could obtain the moments by taking derivatives of the MGF and evaluating at 0. In this case, though, we recognize $1/(1-t)$ as a geometric series, valid in an interval around 0. For $|t| < 1$,

$$M(t) = \frac{1}{1-t} = \sum_{n=0}^{\infty} t^n = \sum_{n=0}^{\infty} n! \frac{t^n}{n!}$$

On the other hand, we know that $E(X^n)$ is the coefficient of the term involving t^n in the Taylor expansion of $M(t)$:

$$M(t) = \sum_{n=0}^{\infty} E(X^n) \frac{t^n}{n!}$$

Thus we can match coefficients to conclude that $E(X^n) = n!$ for all n . We not only did not have to do a LOTUS integral, but also we did not, for example, have to take 10 derivatives to get the 10th moment—we got the moments all at once.

To find the moments of $Y \sim \text{Expo}(\lambda)$, use a scale transformation: we can express $Y = X/\lambda$ where $X \sim \text{Expo}(1)$. Therefore $Y^n = X^n/\lambda^n$ and

$$E(Y^n) = \frac{n!}{\lambda^n}$$

In particular, we have found the mean and variance of Y , making good on our promise from Chapter 5:

$$E(Y) = \frac{1}{\lambda}$$

$$\text{Var}(Y) = E(Y^2) - (EY)^2 = \frac{2}{\lambda^2} - \frac{1}{\lambda^2} = \frac{1}{\lambda^2}.$$

Example 6.5.2 (Standard Normal moments). In this example we will find all the moments of the standard Normal distribution. Let $Z \sim \mathcal{N}(0, 1)$. We can use the same trick of matching the coefficients of the Taylor expansion.

$$M(t) = e^{t^2/2} = \sum_{n=0}^{\infty} \frac{(t^2/2)^n}{n!} = \sum_{n=0}^{\infty} \frac{t^{2n}}{2^n \cdot n!} = \sum_{n=0}^{\infty} \frac{(2n)!}{2^n \cdot n!} \frac{t^{2n}}{(2n)!}$$

Therefore

$$E(Z^{2n}) = \frac{(2n)!}{2^n \cdot n!}$$

and the odd moments of Z are equal to 0, which must be true due to the symmetry of the standard Normal. From the story proof about partnerships in Example 1.5.4, we know that $\frac{(2n)!}{2^n \cdot n!}$ is the product of the odd numbers from 1 through $2n-1$, so

$$E(Z^2) = 1, E(Z^4) = 1 \cdot 3, E(Z^6) = 1 \cdot 3 \cdot 5, \dots$$

This result also shows that the kurtosis of a Normal r.v. is 0. For $X \sim \mathcal{N}(\mu, \sigma^2)$,

$$\text{Kurt}(X) = E\left(\frac{X-\mu}{\sigma}\right)^4 - 3 = E(Z^4) - 3 = 3 - 3 = 0.$$

Example 6.5.3 (Log-Normal moments). Now let's consider the Log-Normal distribution. We say that Y is Log-Normal with parameters μ and σ^2 , denoted by $Y \sim \mathcal{LN}(\mu, \sigma^2)$, if $Y = e^X$ where $X \sim \mathcal{N}(\mu, \sigma^2)$.

6.5.4. Log-Normal does not mean "log of a Normal", since a Normal can be negative. Rather, Log-Normal means "log is Normal". It is important to distinguish between the mean and variance of the Log-Normal and the mean and variance of the underlying Normal. Here we are defining μ and σ^2 to be the mean and variance of the underlying Normal, which is the most common convention.

Interestingly, the Log-Normal MGF does not exist, since $E(e^{tY})$ is infinite for all $t > 0$. Consider the case where $Y = e^Z$ for $Z \sim \mathcal{N}(0, 1)$; by LOTUS,

$$E(e^{tY}) = E(e^{te^z}) = \int_{-\infty}^{\infty} e^{te^z} \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{te^z - z^2/2} dz$$

For any $t > 0$, $te^z - z^2/2$ goes to infinity as z grows, so the above integral diverges. Since $E(e^{tY})$ is not finite on an open interval around 0, the MGF of Y does not exist. The same reasoning holds for a general Log-Normal distribution.

However, even though the Log-Normal MGF does not exist, we can still obtain all the moments of the Log-Normal, using the MGF of the Normal. For $Y = e^X$ with $X \sim \mathcal{N}(\mu, \sigma^2)$,

$$E(Y^n) = E(e^{nX}) = M_X(n) = e^{n\mu + \frac{1}{2}n^2\sigma^2}.$$

In other words, the n th moment of the Log-Normal is the MGF of the Normal evaluated at $t = n$. Letting

$$m = E(Y) = e^{\mu + \frac{1}{2}\sigma^2}$$

we have, after some algebra,

$$\text{Var}(Y) = E(Y^2) - m^2 = m^2(e^{\sigma^2} - 1)$$

All Log-Normal distributions are right-skewed. For example, Figure 6.2 shows a

Log-Normal PDF in dark, with mean 2 and variance 12. This is the distribution of e^X for $X \sim \mathcal{N}(0, 2 \log 2)$, and it is clearly right-skewed. To quantify this, let us compute the skewness of the Log-Normal r.v. $Y = e^X$ for $X \sim \mathcal{N}(0, \sigma^2)$. Letting $m = E(Y) = e^{\frac{1}{2}\sigma^2}$, we have $E(Y^n) = m^{n^2}$ and $\text{Var}(Y) = m^2(m^2 - 1)$, and the third central moment is

$$\begin{aligned} E(Y - m)^3 &= E(Y^3 - 3mY^2 + 3m^2Y - m^3) \\ &= E(Y^3) - 3mE(Y^2) + 2m^3 \\ &= m^9 - 3m^5 + 2m^3 \end{aligned}$$

Thus, the skewness is

$$\text{Skew}(Y) = \frac{E(Y - m)^3}{\text{SD}^3(Y)} = \frac{m^9 - 3m^5 + 2m^3}{m^3(m^2 - 1)^{3/2}} = (m^2 + 2)\sqrt{m^2 - 1}$$

where in the last step we factored $m^6 - 3m^2 + 2 = (m^2 + 2)(m - 1)^2(m + 1)^2$. The skewness is positive since $m > 1$, and it increases very quickly as σ grows.

In the next example, we introduce the Weibull distribution, which is one of the most widely used distributions in survival analysis (the study of the duration of time until an event occurs, e.g., modeling how long someone with a particular disease will live).

Example 6.5.5 (Weibull distribution). As you may remember from the previous chapter, the Exponential distribution is memoryless, which makes it unrealistic for, e.g., modeling a human lifetime. Remarkably, simply raising an Exponential r.v. to a power dramatically improves the flexibility and applicability of the distribution.

Let $T = X^{1/\gamma}$, with $X \sim \text{Expo}(\lambda)$ and $\lambda, \gamma > 0$. The distribution of T is called the Weibull

distribution, and we denote this by $T \sim \text{Wei}(\lambda, \gamma)$. This generalizes the Exponential, with the case $\gamma = 1$ reducing back to the Exponential.

Weibull distributions are widely used in biostatistics, epidemiology, and engineering; there is even an 800 -page book devoted to this distribution: *The Weibull Distribution: A Handbook* by Horst Rinne [21].

The PDF of T is

$$f(t) = \gamma \lambda e^{-\lambda t^\gamma} t^{\gamma-1}$$

for $t > 0$, as can be shown by relating the CDF of T to the CDF of X , or by using transformation results from Chapter 8. The PDF looks somewhat complicated, but often when working with a Weibull random variable we can use the strategy of first transforming it back to an Exponential.

For simplicity and concreteness, let's look at a specific Weibull distribution, letting $\lambda = 1$ and $\gamma = 1/3$.

- (a) Find $P(T > s + t \mid T > s)$ for $s, t > 0$. Does T have the memoryless property?
- (b) Find the mean and variance of T , and the n th moment $E(T^n)$ for $n = 1, 2, \dots$
- (c) Determine whether or not the MGF of T exists.

Solution:

- (a) The CDF of T is

$$P(T \leq t) = P(X^3 \leq t) = P(X \leq t^{1/3}) = 1 - e^{-t^{1/3}},$$

for $t > 0$. So

$$P(T > s + t \mid T > s) = \frac{P(T > s + t)}{P(T > s)} = \frac{e^{-(s+t)^{1/3}}}{e^{-s^{1/3}}},$$

which is not the same as $P(T > t) = e^{-t^{1/3}}$. Thus, T does not have the memoryless property. Nor could it, since it is not Exponential.

(b) Example 6.5.1 shows that the moments of X are given by $E(X^n) = n!$. This allows us to find the moments of T without doing any additional work! Specifically,

$$E(T^n) = E(X^{3n}) = (3n)!.$$

The mean and variance of T are

$$E(T) = 3! = 6, \text{Var}(T) = 6! - 6^2 = 684.$$

- (c) By LOTUS,

$$E(e^{tT}) = E(e^{tX^3}) = \int_0^\infty e^{tx^3 - x} dx$$

This integral diverges for $t > 0$ since the tx^3 term dominates over the x ; more precisely, we have $tx^3 - x > x$ for all x sufficiently large (specifically, for $x > \sqrt[3]{2/t}$), so this integral diverges by comparison with the divergent integral $\int_0^\infty e^x dx$. So the MGF of T does not exist, even though all the moments of T do exist.

6.5.6. Several different parameterizations of the Weibull are commonly used, e.g., we are including a scale in the Exponential and then raising to a power, but it is also common to take an Expo(1) to a power and then rescale. So care is needed when reading various references: always check which convention is being used. Our parameterization here is widely used in

medical statistics, and is convenient to work with since we are just raising an $\text{Expo}(\lambda)$ r.v. to a power.

4.6.6 6.6 Sums of Independent r.v.s via Mgf's

Since the MGF of a sum of independent r.v.s is just the product of the individual MGFs, we now have a new strategy for finding the distribution of a sum of independent r.v.s: multiply the individual MGFs together and see if the product is recognizable as the MGF of a named distribution. The next two examples illustrate this strategy.

Example 6.6.1 (Sum of independent Poissons). Using MGFs, we can easily show that the sum of independent Poissons is Poisson. First let's find the MGF of $X \sim \text{Pois}(\lambda)$:

$$E(e^{tX}) = \sum_{k=0}^{\infty} e^{tk} \frac{e^{-\lambda} \lambda^k}{k!} = e^{-\lambda} \sum_{k=0}^{\infty} \frac{(\lambda e^t)^k}{k!} = e^{-\lambda} e^{\lambda e^t} = e^{\lambda(e^t - 1)}.$$

Now let $Y \sim \text{Pois}(\mu)$ be independent of X . The MGF of $X + Y$ is

$$E(e^{tX}) E(e^{tY}) = e^{\lambda(e^t - 1)} e^{\mu(e^t - 1)} = e^{(\lambda + \mu)(e^t - 1)}$$

which is the $\text{Pois}(\lambda + \mu)$ MGF. Since the MGF determines the distribution, we have proven that $X + Y \sim \text{Pois}(\lambda + \mu)$. Contrast this with the proof from Chapter 4 (Theorem 4.8.1), which required using the law of total probability and summing over all possible values of X . The proof using MGFs is far less tedious.

6.6.2. It is important that X and Y be independent in the above example. To see why, consider an extreme form of dependence: $X = Y$. In that case, $X + Y = 2X$, which can't possibly be Poisson since its value is always an even number!

Example 6.6.3 (Sum of independent Normals). If we have $X_1 \sim \mathcal{N}(\mu_1, \sigma_1^2)$ and $X_2 \sim \mathcal{N}(\mu_2, \sigma_2^2)$ independently, then the MGF of $X_1 + X_2$ is

$$M_{X_1+X_2}(t) = M_{X_1}(t) M_{X_2}(t) = e^{\mu_1 t + \frac{1}{2} \sigma_1^2 t^2} \cdot e^{\mu_2 t + \frac{1}{2} \sigma_2^2 t^2} = e^{(\mu_1 + \mu_2)t + \frac{1}{2}(\sigma_1^2 + \sigma_2^2)t^2}$$

which is the $\mathcal{N}(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$ MGF. Again, because the MGF determines the distribution, it must be the case that $X_1 + X_2 \sim \mathcal{N}(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$. Thus the sum of independent Normals is Normal, and the means and variances simply add.

Example 6.6.4 (Sum is Normal). A converse to the previous example also holds: if X_1 and X_2 are independent and $X_1 + X_2$ is Normal, then X_1 and X_2 must be Normal! This is known as Cramér's theorem. Proving this in full generality is difficult, but it becomes much easier if X_1 and X_2 are i.i.d. with MGF $M(t)$. Without loss of generality, we can assume $X_1 + X_2 \sim \mathcal{N}(0, 1)$, and then its MGF is

$$e^{t^2/2} = E(e^{t(X_1+X_2)}) = E(e^{tX_1}) E(e^{tX_2}) = (M(t))^2$$

so $M(t) = e^{t^2/4}$, which is the $\mathcal{N}(0, 1/2)$ MGF. Thus, $X_1, X_2 \sim \mathcal{N}(0, 1/2)$.

In Chapter 8 we'll discuss a more general technique for finding the distribution of a sum of r.v.s, which applies when the individual MGFs don't exist, or when the product of the individual MGFs is not recognizable and we would like to get the PMF/PDF instead.

4.6.7 6.7 *probability Generating Functions

In this section we discuss probability generating functions, which are similar to MGFs but are guaranteed to exist for nonnegative integer-valued r.v.s. First we'll use PGFs to conquer a seemingly intractable counting problem. Then we'll prove that the PGF of a

nonnegative integer-valued r.v. determines its distribution, which we omitted in the more general MGF setting.

Definition 6.7.1 (Probability generating function). The probability generating function (PGF) of a nonnegative integer-valued r.v. X with PMF $p_k = P(X = k)$ is the generating function of the PMF. By LOTUS, this is

$$E(t^X) = \sum_{k=0}^{\infty} p_k t^k$$

The PGF converges to a value in $[-1, 1]$ for all t in $[-1, 1]$ since $\sum_{k=0}^{\infty} p_k = 1$ and $|p_k t^k| \leq p_k$ for $|t| \leq 1$.

The MGF is closely related to the PGF, when both exist: for $t > 0$,

$$E(t^X) = E(e^{X \log t})$$

is the MGF evaluated at $\log t$.

Example 6.7.2 (Generating dice probabilities). Frederick Mosteller, the founder of the Harvard Statistics Department, once recounted the following life-changing moment:

A key moment in my life occurred in one of those classes during my sophomore year. We had the question: When three dice are rolled what is the chance that the sum of the faces will be 10? The students in this course were very good, but we all got the answer largely by counting on our fingers. When we came to class, I said to the teacher, "That's all very well—we got the answer—but if we had been asked about six dice and the probability of getting 18, we would still be home counting. How do you do problems like that?" He said, "I don't know, but I know a man who probably does and I'll ask him."

One day I was in the library and Professor Edwin G. Olds of the Mathematics Department came in. He shouted at me, "I hear you're interested in the three dice problem." He had a huge voice, and you know how libraries are. I was embarrassed. "Well, come and see me," he said, "and I'll show you about it." "Sure," I said. But I was saying to myself, "I'll never go." Then he said, "What are you doing?" I showed him. "That's nothing important," he said. "Let's go now."

So we went to his office, and he showed me a generating function. It was the most marvelous thing I had ever seen in mathematics. It used mathematics that, up to that time, in my heart of hearts, I had thought was something that mathematicians just did to create homework problems for innocent students in high school and college. I don't know where I had got ideas like that about various parts of mathematics. Anyway, I was stunned when I saw how Olds used this mathematics that I hadn't believed in. He used it in such an unusually outrageous way. It was a total retranslation of the meaning of the numbers. [1]

Let X be the total from rolling 6 fair dice, and let $X_1, \dots, X_6$ be the individual rolls. What is $P(X = 18)$? It turns out that there are 3431 ways to obtain a sum of 18, so the probability is $3431/6^6 \approx 0.0735$. Listing out all possibilities is extremely tedious, and the tedium would be compounded with the worry of having somehow missed a case. And what if we laboriously listed out all 3431 cases, and then were asked to find $P(X = 19)$?

The PGF of X lets us count the cases in a systematic way. The PGF of X_1 is

$$E(t^{X_1}) = \frac{1}{6} (t + t^2 + \dots + t^6)$$

Since the X_j are i.i.d., the PGF of X is

$$E(t^X) = E(t^{X_1} \dots t^{X_6}) = E(t^{X_1}) \dots E(t^{X_6}) = \frac{t^6}{6^6} (1 + t + \dots + t^5)^6$$

By definition, the coefficient of t^{18} in the PGF is $P(X = 18)$. So the number of ways to get a sum of 18 is the coefficient of t^{18} in $t^6(1 + t + \cdots + t^5)^6$, which is the coefficient of t^{12} in $(1 + t + \cdots + t^5)^6$. Multiplying this out by hand is tedious, but it is vastly easier than listing out 3431 cases, and it can also be done easily on a computer without having to write a special program.

Better yet, we can use the fact that $1 + t + \cdots + t^5$ is a geometric series to write

$$(1 + t + \cdots + t^5)^6 = \frac{(1 - t^6)^6}{(1 - t)^6}$$

(Assume that $|t| < 1$, which we can do since, as with the MGF, we just need to know how the PGF behaves in an open interval containing 0.) The above equation is just algebra since we have the bookkeeping device t , but would have been hard to fathom if everything were still in sequence notation. By the binomial theorem, the numerator is

$$(1 - t^6)^6 = \sum_{j=0}^6 \binom{6}{j} (-1)^j t^{6j}$$

For the denominator, write

$$\frac{1}{(1 - t)^6} = (1 + t + t^2 + \cdots)^6 = \sum_{k=0}^{\infty} a_k t^k$$

Here a_k is the number of ways to choose one term from each of the six $(1 + t + t^2 + \cdots)$ factors, such that the degrees add up to k . For example, for $k = 20$ one possibility is to choose the $t^3, 1, t^2, t^{10}, 1, t^5$ terms, respectively, since these choices contribute one t^{20} term when the product is expanded out. So a_k is the number of solutions to $y_1 + y_2 + \cdots + y_6 = k$, with the y_j nonnegative integers. We saw how to count this number of solutions in Chapter 1: a_k is the Bose-Einstein value $\binom{6+k-1}{k} = \binom{k+5}{5}$. So

$$\frac{1}{(1 - t)^6} = \sum_{k=0}^{\infty} \binom{k+5}{5} t^k$$

For $0 < t < 1$, another way to see why this equation holds is to write it as

$$\sum_{k=0}^{\infty} \binom{k+5}{5} (1 - t)^6 t^k = 1,$$

which we already knew to be true since the $\text{NBin}(6, 1 - t)$ PMF must sum to 1. (The identity for $(1 - t)^{-6}$ is an example of a generalization of the binomial theorem to allow for negative integer powers; this helps explain why the Negative Binomial is called that despite being neither negative nor Binomial!)

Putting together the above results, we just need the coefficient of t^{12} in

$$\left(\sum_{j=0}^2 \binom{6}{j} (-1)^j t^{6j} \right) \left(\sum_{k=0}^{12} \binom{k+5}{5} t^k \right),$$

where we summed only up to $j = 2$ and $k = 12$ in the two factors since any further terms will not contribute to the coefficient of t^{12} . This lets us reduce the 3431 cases down to just three cases: (j, k) is $(0, 12)$, $(1, 6)$, or $(2, 0)$. The coefficient of t^{12} is

$$\binom{17}{5} - 6 \binom{11}{5} + \binom{6}{2} = 3431$$

since, for example, when $j = 1$ and $k = 6$ we get the term

$$-\binom{6}{1}t^6 \cdot \binom{6+5}{5}t^6 = -6\binom{11}{5}t^{12}.$$

Thus,

$$P(X = 18) = \frac{3431}{6^6}.$$

Since the PGF is just a handy bookkeeping device for the PMF, it fully determines the distribution (for any nonnegative integer-valued r.v.). The theorem below shows how to pluck the PMF values down from the "clothesline" of the PGF.

Theorem 6.7.3. Let X and Y be nonnegative integer-valued r.v.s, with PGFs g_X and g_Y respectively. Suppose that $g_X(t) = g_Y(t)$ for all t in $(-a, a)$, where $0 < a < 1$. Then X and Y have the same distribution, and their PMF can be recovered by taking derivatives of g_X :

$$P(X = k) = P(Y = k) = \frac{g_X^{(k)}(0)}{k!}.$$

Proof. Write

$$g_X(t) = \sum_{k=0}^{\infty} p_k t^k$$

Then $g_X(0) = p_0$, so $P(X = 0)$ has been recovered-from knowing the function g_X , we can extract the value of $P(X = 0)$. The derivative is

$$g'_X(t) = \sum_{k=1}^{\infty} k p_k t^{k-1}$$

so $g'_X(0) = p_1$ (swapping the derivative and the infinite sum is justified by results in real analysis). Then $P(X = 1)$ has been recovered. Continuing in this way, we can recover the entire PMF by taking derivatives.

Example 6.7.4. Let $X \sim \text{Bin}(n, p)$. The PGF of a $\text{Bern}(p)$ r.v. is $pt + q$ (with $q = 1 - p$), so the PGF of X is $g(t) = (pt + q)^n$. The above theorem says that any r.v. with this PGF must in fact be Binomial. Furthermore, we can recover the PMF by computing

$$g(0) = q^n, g'(0) = npq^{n-1}, g''(0)/2! = \binom{n}{2}p^2q^{n-2}, \dots$$

We can avoid having to take derivatives by using the binomial theorem to write

$$g(t) = (pt + q)^n = \sum_{k=0}^n \binom{n}{k} p^k q^{n-k} t^k$$

from which we can directly read off the Binomial PMF.

While we're working with the Binomial PGF, let's see how it can be used to get the moments of a Binomial. Letting $p_k = P(X = k)$, we have

$$g'(t) = np(pt + q)^{n-1} = \sum_{k=1}^n k p_k t^{k-1},$$

so

$$g'(1) = np = \sum_{k=1}^n kp_k = E(X)$$

Taking the derivative again,

$$g''(t) = n(n-1)p^2(pt+q)^{n-2} = \sum_{k=2}^n k(k-1)p_k t^{k-2},$$

so

$$E(X(X-1)) = g''(1) = n(n-1)p^2.$$

Rearranging these results gives another proof that $\text{Var}(X) = npq$. Continuing in this way, we have computed what are called the factorial moments of the Binomial:

$$E(X(X-1)\dots(X-k+1)) = k! \binom{n}{k} p^k$$

Dividing by $k!$ on both sides, this implies $E\binom{X}{k} = \binom{n}{k} p^k$, which can also be seen with a story proof: $\binom{X}{k}$ is the number of ways to choose k out of the X successful Bernoulli trials, which is the number of ways to choose k out of the n original trials such that all k are successes. Creating an indicator r.v. for each of the $\binom{n}{k}$ subsets of size k and using linearity, the result follows.

4.6.8 6.8 Recap

A useful way to study a distribution is via its moments. The first 4 moments are widely used as a basis for quantitatively describing what the distribution looks like, though many other descriptions are also possible. In particular, the first moment is the mean, the second central moment is the variance, the third standardized moment measures skew (asymmetry), and the fourth standardized moment minus 3 is a measure of how heavy the tails are.

Moments are useful for far more than studying the location and shape of a distribution, especially when the moment generating function (MGF) exists (which is stronger than saying that all the moments exist). The MGF of an r.v. X is the function M defined by

$$M(t) = E(e^{tX})$$

if this is finite for all t in some open interval containing 0. If the MGF exists, then

$$M(0) = 1, M'(0) = E(X), M''(0) = E(X^2), M'''(0) = E(X^3), \dots$$

MGFs are useful for three main reasons: for computing moments (as an alternative to LOTUS), for studying sums of independent r.v.s, and since they fully determine the distribution and thus serve as an additional blueprint for a distribution.

We also introduced the Log-Normal and Weibull distributions in this chapter, both of which are widely used in practice. The Log-Normal and Weibull are connected to the Normal and Exponential, respectively, via simple transformations. The log of a log-Normal r.v. is Normal, and raising a $\text{Wei}(\lambda, \gamma)$ r.v. to the power γ yields an $\text{Expo}(\lambda)$ r.v. Often the best way to study a Log-Normal is to transform it back to Normal, and likewise for Weibull and Exponential. Chapter 8 goes into much more detail about how to work with transformations.

Figure 6.8 augments our map of the connections between fundamental objects in probability. If the MGF of X exists, then the sequence $E(X), E(X^2), E(X^3), \dots$ of moments provides enough information (at least in principle) to determine the distribution of X .




FIGURE 6.8

For an r.v. X , we can study its moments $E(X)$, $E(X^2)$, $\dots$. These can be computed using LOTUS or by the MGF (if it exists). If the MGF exists, it determines the distribution, taking us full circle and adding to our list of blueprints.

4.6.9 6.9r

Functions

The MGF of an r.v. is a function. As an example of defining and working with functions in R, let's use the $\mathcal{N}(0, 1)$ MGF, which is given by $M(t) = e^{t^2/2}$. The code

```
M <- function(t) {exp(t^2/2)}
```

defines M to be this function. The function (t) says that we're defining a function of one variable t (called the argument of the function). Then, for example, $M(0)$ evaluates the function at 0, $M(1 : 10)$ evaluates the function at 1, 2, ..., 10, and `curve(M, from = -3, to = 3)` plots the graph of M from -3 to 3. Writing

```
M <- function(x) {exp(x^2/2)}
```

would define the same function M , except that now the argument is named x . Giving the arguments names is helpful for functions of more than one variable, since R then saves us from having to remember the order in which to write the arguments, and allows us to assign default values. For example, the $\mathcal{N}(\mu, \sigma^2)$ MGF is given by $g(t) = \exp(\mu t + \sigma^2 t^2/2)$, which depends on t, μ , and σ . We can define this in R by `g <- function(t, mean = 0, sd = 1) {exp(mean * t + sd^2 * t^2/2)}`. What is `g(1, 2, 3)`? It's the $\mathcal{N}(2, 3^2)$ MGF evaluated at 1, but it may be hard to remember which argument is which, especially when working with many functions with many arguments over the course of many months. So we can also write `g(t = 1, mean = 2, sd = 3)` or `g(mean = 2, sd = 3, t = 1)` or any of the other 4 permutations to mean the same thing.

Also, when defining g we specified default values of 0 for the mean and 1 for the standard deviation, so if we want the $\mathcal{N}(0, 5^2)$ MGF evaluated at 3, we can use `g(t = 3, sd = 5)` as shorthand. It would be bad here to write `g(3, 5)`, since that is ambiguous about which argument is omitted; in fact, R interprets this as `g(t = 3, mean = 5)`.

Moments

LOTUS makes it easy to write down any moment of a continuous r.v. as an integral, and then R can help us do the integral numerically, using the `integrate` command. For example,

let's approximate the 6th moment of a $\mathcal{N}(0, 1)$ r.v. The code

```
g <- function(x) x^6*dnorm(x)
```

```
integrate(g, lower = -Inf, upper = Inf)
```

asks R to compute $\int_{-\infty}^{\infty} g(x)dx$, where $g(x) = x^6\varphi(x)$ with φ the $\mathcal{N}(0, 1)$ PDF. When we ran this, R reported 15 (the correct answer, as we know from this chapter!) and that the absolute error was less than 7.9×10^{-5} . Similarly, to check that the 2nd moment (and variance) of a $\text{Unif}(-1, 1)$ r.v. is $1/3$, we can use

```
h <- function(x) x^2*dunif(x, -1, 1)
```

```
integrate(h, lower = -1, upper = 1)
```

6.9.1. Numerical integration runs into difficulties for some functions; as usual, checking answers in multiple ways is a good idea. Using `upper = Inf` is preferred to using a large number as the upper limit when integrating up to ∞ (and likewise for a lower limit of $-\infty$). For example, on many systems `integrate(dnorm, 0, 10^6)` reports 0 while `integrate(dnorm, 0, Inf)` reports the correct answer, 0.5.

For moments of a discrete r.v., we can use LOTUS and the `sum` command. For example, to find the 2nd moment of $X \sim \text{Pois}(7)$, we can use

```
g <- function(k) k^2*dpois(k,7)
```

```
sum(g(0 : 100))
```

Here we summed up to 100 since it's clear after getting a sense of the terms that the total contribution of all the terms after $k = 100$ is negligible (choosing an upper limit in this way is in contrast to how to use the `integrate` command in the continuous case). The result is extremely close to 56, which is comforting since $E(X^2) = \text{Var}(X) + (EX)^2 = 7 + 49 = 56$.

A sample moment can be found in one line in R. If $\mathbf{x}$ is a vector of data, then `mean(x)` gives its sample mean and, more generally, `mean(x^n)` gives the n th sample mean for any positive integer n . For example,

```
x <- rnorm(100)
```

```
mean(x^6)
```

gives the 6th sample moment of 100 i.i.d. $\mathcal{N}(0, 1)$ r.v.s. How close is it to the true 6th moment? How close are other sample moments to the corresponding true moments?

The sample variance can also be found in one line in R. If $\mathbf{x}$ is a vector of data, then `var(x)` gives its sample variance. This returns NA (not available) if $\mathbf{x}$ has length 1, since the $n - 1$ in the denominator is 0 in this case. It makes sense not to return a numerical value in this case, not only because of the definition but also because it would be insane to try to estimate the variability of a population if we only have one observation!

For a simple demonstration of using the sample mean and sample variance to estimate the true mean and true variance of a distribution, we generate 1000 times from a $\mathcal{N}(0, 1)$ distribution and store the values in $\mathbf{z}$. We then compute the sample mean and sample variance with `mean` and `var`.

```
z <- rnorm(1000)
```

```
mean(z)
```

```
var(z)
```

We find that `mean(z)` is close to 0 and `var(z)` is close to 1. You can try this out for a $\mathcal{N}(\mu, \sigma^2)$ distribution (or other distribution) of your choosing; just remember that `rnorm` takes σ and not σ^2 as an input!

The sample standard deviation of $\mathbf{x}$ can be found using `sd(x)`. This gives the same result as `sqrt(var(x))`.

R does not come with built-in functions for sample skewness or sample kurtosis, but we can define our own functions as follows.

```
skew <- function(x) {
  centralmoment <- mean((x-mean(x))^3)
  centralmoment/(sd(x)^3)
}
kurt <- function(x) {
  centralmoment <- mean((x-mean(x))^4)
  centralmoment/(sd(x)^4) - 3
}
```

Medians and Modes

To find the median of a continuous r.v. with CDF F , we need to solve the equation $F(x) = 1/2$ for x , which is equivalent to finding the root (zero) of the function g given by $g(x) = F(x) - 1/2$. This can be done using `uniroot` in R. For example, let's find the median of the `Expo(1)` distribution. The code

```
g <- function(x) pexp(x) - 1/2
uniroot(g, lower=0, upper=1)
```

asks R to find a root of the desired function between 0 and 1. This returns an answer very close to the true answer of $\log(2) \approx 0.693$. Of course, in this case we can solve $1 - e^{-x} = 1/2$ directly without having to use numerical methods.

6.9.2. The `uniroot` command is useful but it only attempts to find one root (as the name suggests), and there is no guarantee that it will find a root.

An easier way to find the median of the `Expo(1)` in R is to use `qexp(1/2)`. The function `qexp` is the quantile function of the `Expo(1)` distribution, which means that `qexp(p)` is the value of x such that $P(X \leq x) = p$ for $X \sim \text{Expo}(1)$.

For finding the mode of a continuous distribution, we can use the `optimize` function in R. For example, let's find the mode of the `Gamma(6, 1)` distribution, which is an important distribution that we will introduce in the next chapter. Its PDF is proportional to $x^5 e^{-x}$. Using calculus, we can find that the mode is at $x = 5$. Using R, we can find that the mode is very close to $x = 5$ as follows.

```
h <- function(x) x^5 * exp(-x)
optimize(h, lower = 0, upper = 20, maximum=TRUE)
```

If we had wanted to minimize instead of maximize, we could have put `maximum=FALSE`.

Next, let's do a discrete example of median and mode. An interesting fact about the `Bin( $n, p$ )` distribution is that if the mean np is an integer, then the median and mode are also np (even if the distribution is very skewed). To check this fact about the median for the `Bin(50, 0.2)` distribution, we can use the following code.

```
n <- 50; p <- 0.2
which.max(pbinom(0:n, n, p) >= 0.5)
```

The `which.max` function finds the location of the maximum of a vector, giving the index of the first occurrence of a maximum. Since `TRUE` is encoded as 1 and `FALSE` is encoded as 0, the first maximum in `pbinom(0 : n, n, p) >= 0.5` is at the first value for which the CDF is at least 0.5. The output of the above code is 11, but we must be careful to avoid an off-by-one error: the index 11 corresponds to the median being 10, since we started evaluating the CDF at 0. Similarly, `which.max(dbinom(0 : n, n, p))` returns 11, showing that the mode is at 10.

The sample median of a vector $\mathbf{x}$ of data can be found using `median(x)`. But `mode(x)` does not give the sample mode of $\mathbf{x}$ (rather, it gives information about what type of

object $\mathbf{x}$ is). To find the sample mode (or sample modes, in case there are ties), we can use the following function.

```
datamode <- function(x) {  
  t <- table(x)  
  m <- max(t)  
  as.numeric(names(t[t==m]))  
}
```

Log-normal and Weibull Distributions

Analogous to the functions `dnorm`, `pnorm`, and `rnorm` for the Normal distribution, the functions `dlnorm`, `plnorm`, and `rlnorm` give the Log-Normal PDF, the Log-Normal CDF, and random generation of Log-Normal r.v.s, respectively. The parameters used for these functions for the Log-Normal are the mean and standard deviation of the underlying Normal.

For example, `dl norm(x, 1, 2)` gives the PDF of the $\mathcal{LN}(1, 4)$ distribution (not $\mathcal{LN}(1, 2)$, nor a Log-Normal whose mean is 1). Because of the relationship between Normal and Log-Normal, `rlnorm(n, mu, sigma)` is equivalent to `exp(rnorm(n, mu, sigma))`.

For the Weibull, we can obtain the PDF, the CDF, and random generation with the functions `dweibull`, `pweibull`, and `rweibull`, respectively. The parametrization in R is different from the one we are using, but it is easy to convert between them: for the $\text{Wei}(\lambda, \gamma)$ distribution, let $a = \gamma$ and $b = \lambda^{-1/\gamma}$.

Then `dweibull(x, a, b)` gives the $\text{Wei}(\lambda, \gamma)$ PDF, `pweibull(x, a, b)` gives the CDF, and `rweibull(n, a, b)` generates n i.i.d. draws from the distribution. Because of the relationship between Exponential and Weibull, another way to generate $\text{Wei}(\lambda, \gamma)$ r.v.s is to generate $\text{Expo}(\lambda)$ r.v.s and then raise each of them to the $1/\gamma$ power.

Dice Simulation

In the starred Section 6.7, we showed that in rolling 6 fair dice, the probability of a total of 18 is $3431/6^6 \approx 0.07354$. But the proof was complicated. If we only need an approximate answer, simulation is a much easier approach. And we already know how to do it! Here is the code for a million repetitions:

```
r <- replicate(10^6, sum(sample(6, 6, replace=TRUE)))  
sum(r==18)/10^6
```

In our simulation this yielded 0.07346, which is very close to 0.07354.

4.6.10 6.10 Exercises

Exercises marked with (S) have detailed solutions at <http://stat110.net>.

Means, Medians, Modes, and Moments

1. Let $U \sim \text{Unif}(a, b)$. Find the median and mode of U .
2. Let $X \sim \text{Expo}(\lambda)$. Find the median and mode of X .
3. Let X have the Pareto distribution with parameter $a > 0$; this means that X has PDF $f(x) = a/x^{a+1}$ for $x \geq 1$ (and 0 otherwise). Find the median and mode of X .

4. Let $X \sim \text{Bin}(n, p)$.
 - (a) For $n = 5, p = 1/3$, find all medians and all modes of X . How do they compare to the mean?
 - (b) For $n = 6, p = 1/3$, find all medians and all modes of X . How do they compare to the mean?
5. Let X be Discrete Uniform on $1, 2, \dots, n$. Find all medians and all modes of X (your answer can depend on whether n is even or odd).
6. Suppose that we have data giving the amount of rainfall in a city each day in a certain year. We want useful, informative summaries of how rainy the city was that year. On the majority of days in that year, it did not rain at all in the city. Discuss and compare the following six summaries: the mean, median, and mode of the rainfall on a randomly chosen day from that year, and the mean, median, and mode of the rainfall on a randomly chosen rainy day from that year (where by "rainy day" we mean that it did rain that day in the city).
7. Let a and b be positive constants. The Beta distribution with parameters a and b , which we introduce in detail in Chapter 8, has PDF proportional to $x^{a-1}(1-x)^{b-1}$ for $0 < x < 1$ (and the PDF is 0 outside of this range). Show that for $a > 1, b > 1$, the mode of the distribution is $(a-1)/(a+b-2)$.
Hint: Take the log of the PDF first (note that this does not affect where the maximum is achieved).
8. Find the median of the Beta distribution with parameters $a = 3$ and $b = 1$ (see the previous problem for information about the Beta distribution).
9. Let Y be Log-Normal with parameters μ and σ^2 . So $Y = e^X$ with $X \sim \mathcal{N}(\mu, \sigma^2)$. Three students are discussing the median and the mode of Y . Evaluate and explain whether or not each of the following arguments is correct.
 - (a) Student A: The median of Y is e^μ because the median of X is μ and the exponential function is continuous and strictly increasing, so the event $Y \leq e^\mu$ is the same as the event $X \leq \mu$.
 - (b) Student B: The mode of Y is e^μ because the mode of X is μ , which corresponds to e^μ for Y since $Y = e^X$.
 - (c) Student C: The mode of Y is μ because the mode of X is μ and the exponential function is continuous and strictly increasing, so maximizing the PDF of X is equivalent to maximizing the PDF of $Y = e^X$.
10. A distribution is called symmetric unimodal if it is symmetric (about some point) and has a unique mode. For example, any Normal distribution is symmetric unimodal. Let X have a continuous symmetric unimodal distribution for which the mean exists. Show that the mean, median, and mode of X are all equal.
11. Let $X_1, \dots, X_n$ be i.i.d. r.v.s with mean μ , variance σ^2 , and skewness γ .
 - (a) Standardize the X_j by letting

$$Z_j = \frac{X_j - \mu}{\sigma}$$

Let $\bar{X}_n$ and $\bar{Z}_n$ be the sample means of the X_j and Z_j , respectively. Show that Z_j has the same skewness as X_j , and $\bar{Z}_n$ has the same skewness as $\bar{X}_n$.

- (b) Show that the skewness of the sample mean $\bar{X}_n$ is $\gamma/\sqrt{n}$. You can use the fact, shown in

Chapter 7, that if X and Y are independent then $E(XY) = E(X)E(Y)$.

Hint: By (a), we can assume $\mu = 0$ and $\sigma^2 = 1$ without loss of generality; if the X_j are not standardized initially, then we can standardize them. If $(X_1 + X_2 + \cdots + X_n)^3$ is expanded out, there are 3 types of terms: terms such as X_1^3 , terms such as $3X_1^2X_2$, and terms such as $6X_1X_2X_3$.

(c) What does the result of (b) say about the distribution of $\bar{X}_n$ when n is large?

12. Let c be the speed of light in a vacuum. Suppose that c is unknown, and scientists wish to estimate it. But even more so than that, they wish to estimate c^2 , for use in the famous equation $E = mc^2$.

Through careful experiments, they obtain i.i.d. measurements $X_1, \dots, X_n \sim \mathcal{N}(c, \sigma^2)$. Using these data, there are various possible ways to estimate c^2 . Two natural ways are: (1) estimate c using the average of the X_j 's and then square the estimated c , and (2) average the X_j^2 's. So let

$$\bar{X}_n = \frac{1}{n} \sum_{j=1}^n X_j$$

and consider the two estimators

$$T_1 = \bar{X}_n^2 \text{ and } T_2 = \frac{1}{n} \sum_{j=1}^n X_j^2$$

Note that T_1 is the square of the first sample moment and T_2 is the second sample moment.

(a) Find $P(T_1 < T_2)$.

Hint: Start by comparing $\left(\frac{1}{n} \sum_{j=1}^n x_j\right)^2$ and $\frac{1}{n} \sum_{j=1}^n x_j^2$ when $x_1, \dots, x_n$ are numbers, by considering a discrete r.v. whose possible values are $x_1, \dots, x_n$.

(b) When an r.v. T is used to estimate an unknown parameter θ , the bias of the estimator T is defined to be $E(T) - \theta$. Find the bias of T_1 and the bias of T_2 .

Hint: First find the distribution of $\bar{X}_n$. In general, for finding $E(Y^2)$ for an r.v. Y , it is often useful to write it as $E(Y^2) = \text{Var}(Y) + (EY)^2$.

Moment Generating Functions

13. (S) A fair die is rolled twice, with outcomes X for the first roll and Y for the second roll. Find the moment generating function $M_{X+Y}(t)$ of $X + Y$ (your answer should be a function of t and can contain unsimplified finite sums).

14. (S) Let $U_1, U_2, \dots, U_{60}$ be i.i.d. $\text{Unif}(0, 1)$ and $X = U_1 + U_2 + \cdots + U_{60}$. Find the MGF of X .

15. Let $W = X^2 + Y^2$, with X, Y i.i.d. $\mathcal{N}(0, 1)$. The MGF of X^2 turns out to be $(1 - 2t)^{-1/2}$ for $t < 1/2$ (you can assume this).

(a) Find the MGF of W .

(b) What famous distribution that we have studied so far does W follow (be sure to state the parameters in addition to the name)? In fact, the distribution of W is also a special case of two more famous distributions that we will study in later chapters!

16. Let $X \sim \text{Expo}(\lambda)$. Find the skewness of X , and explain why it is positive and why it does not depend on λ .

Hint: Recall that $\lambda X \sim \text{Expo}(1)$ and the n th moment of an $\text{Expo}(1)$ r.v. is $n!$ for all n .

17. Let $X_1, \dots, X_n$ be i.i.d. with mean μ , variance σ^2 , and MGF M . Let

$$Z_n = \sqrt{n} \left(\frac{\bar{X}_n - \mu}{\sigma} \right)$$

- (a) Show that Z_n is a standardized quantity, i.e., it has mean 0 and variance 1 .
- (b) Find the MGF of Z_n in terms of M , the MGF of each X_j .
18. Use the MGF of the Geom(p) distribution to give another proof that the mean of this distribution is q/p and the variance is q/p^2 , with $q = 1 - p$.
19. Use MGFs to determine whether $X + 2Y$ is Poisson if X and Y are i.i.d. $\text{Pois}(\lambda)$.
20. (S) Let $X \sim \text{Pois}(\lambda)$, and let $M(t)$ be the MGF of X . The cumulant generating function is defined to be $g(t) = \log M(t)$. Expanding $g(t)$ as a Taylor series

$$g(t) = \sum_{j=1}^{\infty} \frac{c_j}{j!} t^j$$

(the sum starts at $j = 1$ because $g(0) = 0$), the coefficient c_j is called the j th cumulant of X . Find the j th cumulant of X , for all $j \geq 1$.

21. (S) Let $X_n \sim \text{Bin}(n, p_n)$ for all $n \geq 1$, where np_n is a constant $\lambda > 0$ for all n (so $p_n = \lambda/n$). Let $X \sim \text{Pois}(\lambda)$. Show that the MGF of X_n converges to the MGF of X (this gives another way to see that the $\text{Bin}(n, p)$ distribution can be well-approximated by the $\text{Pois}(\lambda)$ when n is large, p is small, and $\lambda = np$ is moderate).
22. Consider a setting where a Poisson approximation should work well: let $A_1, \dots, A_n$ be independent, rare events, with n large and $p_j = P(A_j)$ small for all j . Let $X = I(A_1) + \dots + I(A_n)$ count how many of the rare events occur, and let $\lambda = E(X)$.
- (a) Find the MGF of X .
- (b) If the approximation $1 + x \approx e^x$ (this is a good approximation when x is very close to 0 but terrible when x is not close to 0) is used to write each factor in the MGF of X as e to a power, what happens to the MGF? Explain why the result makes sense intuitively.
23. Let U_1, U_2 be i.i.d. $\text{Unif}(0, 1)$. Example 8.2.5 in Chapter 8 shows that $U_1 + U_2$ has a Triangle distribution, with PDF given by

$$f(t) = \begin{cases} t & \text{for } 0 < t \leq 1 \\ 2 - t & \text{for } 1 < t < 2 \end{cases}$$

The method in Example 8.2.5 is useful but it often leads to difficult integrals, so having alternative methods is important. Show that $U_1 + U_2$ has a Triangle distribution by showing that they have the same MGF.

24. Let X and Y be i.i.d. $\text{Expo}(1)$, and $L = X - Y$. The Laplace distribution has PDF

$$f(x) = \frac{1}{2} e^{-|x|}$$

for all real x . Use MGFs to show that the distribution of L is Laplace.

25. Let $Z \sim \mathcal{N}(0, 1)$, and $Y = |Z|$. So Y has the Folded Normal distribution, discussed in Example 5.4.7. Find two expressions for the MGF of Y as unsimplified integrals: one integral based on the PDF of Y , and one based on the PDF of Z .
26. Let $X, Y, Z, W \sim \mathcal{N}(0, 1)$ be i.i.d.
- (a) Find an expression for $E(\Phi(Z)e^Z)$ as an unsimplified integral.
- (b) Find $E(\Phi(Z))$ and $E(e^Z)$ as fully simplified numbers.

4.7 Joint Distributions

When we first introduced random variables and their distributions in Chapter 3, we noted that the individual distributions of two r.v.s do not tell us anything about whether the r.v.s are

independent or dependent. For example, two Bern(1/2) r.v.s X and Y could be independent if they indicate Heads on two different coin flips, or dependent if they indicate Heads and Tails, respectively, on the same coin flip. Thus, although the PMF of X is a complete blueprint for X and the PMF of Y is a complete blueprint for Y , these individual PMFs are missing important information about how the two r.v.s are related.

Of course, in real life, we usually care about the relationship between multiple r.v.s in the same experiment. To give just a few examples:

- **Medicine:** To evaluate the effectiveness of a treatment, we may take multiple measurements per patient; an ensemble of blood pressure, heart rate, and cholesterol readings can be more informative than any of these measurements considered separately.
- **Genetics:** To study the relationships between various genetic markers and a particular disease, if we only looked separately at distributions for each genetic marker, we could fail to learn about whether an interaction between markers is related to the disease.
- **Time series:** To study how something evolves over time, we can often make a series of measurements over time, and then study the series jointly. There are many applications of such series, such as global temperatures, stock prices, or national unemployment rates. The series of measurements considered jointly can help us deduce trends for the purpose of forecasting future measurements.

This chapter considers joint distributions, also called multivariate distributions, which capture the previously missing information about how multiple r.v.s interact. We introduce multivariate analogs of the CDF, PMF, and PDF in order to provide a complete specification of the relationship between multiple r.v.s. After this groundwork is in place, we'll study a couple of famous named multivariate distributions, generalizing the Binomial and Normal distributions to higher dimensions.

4.7.1 7.1 Joint, Marginal, and Conditional

The three key concepts for this section are joint, marginal, and conditional distributions. Recall that the distribution of a single r.v. X provides complete information about the probability of X falling into any subset of the real line. Analogously, the joint distribution of two r.v.s X and Y provides complete information about the probability of the vector (X, Y) falling into any subset of the plane. The marginal distribution of X is the individual distribution of X , ignoring the value of Y , and the conditional distribution of X given $Y = y$ is the updated distribution for X after observing $Y = y$. We'll look at these concepts in the discrete case first, then extend them to the continuous case.

7.1.1 Discrete

The most general description of the joint distribution of two r.v.s is the joint CDF, which applies to discrete and continuous r.v.s alike.

Definition 7.1.1 (Joint CDF). The joint CDF of r.v.s X and Y is the function $F_{X,Y}$ given by

$$F_{X,Y}(x, y) = P(X \leq x, Y \leq y).$$

The joint CDF of n r.v.s is defined analogously.

Unfortunately, the joint CDF of discrete r.v.s is not a well-behaved function; as in the univariate case, it consists of jumps and flat regions. For this reason, with discrete r.v.s we usually work with the joint *PMF*, which also determines the joint distribution and is much easier to visualize.

Definition 7.1.2 (Joint PMF). The joint PMF of discrete r.v.s X and Y is the function $p_{X,Y}$ given by

$$p_{X,Y}(x, y) = P(X = x, Y = y).$$

The joint PMF of n discrete r.v.s is defined analogously.

Just as univariate PMFs must be nonnegative and sum to 1, we require valid joint PMFs to be nonnegative and sum to 1, where the sum is taken over all possible values of X and Y :

$$\sum_x \sum_y P(X = x, Y = y) = 1.$$

The joint PMF determines the distribution because we can use it to find the probability of the event $(X, Y) \in A$ for any set A of points in the support of (X, Y) . All we have to do is sum the joint PMF over A :

$$P((X, Y) \in A) = \sum_{(x,y) \in A} P(X = x, Y = y).$$

Figure 7.1 shows a sketch of what the joint PMF of two discrete r.v.s could look like. The height of a vertical bar at (x, y) represents the probability $P(X = x, Y = y)$. For the joint PMF to be valid, the total height of the vertical bars must be 1.




FIGURE 7.1 Joint PMF of discrete r.v.s X and Y .

From the joint distribution of X and Y , we can get the distribution of X alone by summing over the possible values of Y . This gives us the familiar PMF of X that we have seen in previous chapters. In the context of joint distributions, we will call it the marginal or unconditional distribution of X , to make it clear that we are referring to the distribution of X alone, without regard for the value of Y .

Definition 7.1.3 (Marginal PMF). For discrete r.v.s X and Y , the marginal PMF of X is

$$P(X = x) = \sum_y P(X = x, Y = y).$$

The marginal PMF of X is the PMF of X , viewing X individually rather than jointly with Y . The above equation follows from the axioms of probability (we are summing over disjoint cases). The operation of summing over the possible values of Y in order to convert the joint PMF into the marginal PMF of X is known as marginalizing out Y .

The process of obtaining the marginal PMF from the joint PMF is illustrated in Figure 7.2. Here we take a bird's-eye view of the joint PMF for a clearer perspective; each column of the joint PMF corresponds to a fixed x and each row corresponds to a fixed y . For any x , the probability $P(X = x)$ is the total height of the bars in the corresponding column of the joint PMF: we can imagine taking all the bars in that column and stacking them on top of each

other to get the marginal probability. Repeating this for all x , we arrive at the marginal PMF, depicted in bold.




FIGURE 7.2

Bird's-eye view of the joint PMF from Figure 7.1. The marginal PMF $P(X = x)$ is obtained by summing over the joint PMF in the y -direction.

Similarly, the marginal PMF of Y is obtained by summing over all possible values of X . So given the joint PMF, we can marginalize out X to get the PMF of Y , or marginalize out Y to get the PMF of X . But if we only know the marginal PMFs of X and Y , there is no way to recover the joint PMF without further assumptions. It is clear how to stack the bars in Figure 7.2, but very unclear how to unstack the bars after they have been stacked!

Another way to go from joint to marginal distributions is via the joint CDF. In that case, we take a limit rather than a sum: the marginal CDF of X is

$$F_X(x) = P(X \leq x) = \lim_{y \rightarrow \infty} P(X \leq x, Y \leq y) = \lim_{y \rightarrow \infty} F_{X,Y}(x, y).$$

However, as mentioned above it is usually easier to work with joint PMFs.

Now suppose that we observe the value of X and want to update our distribution of Y to reflect this information. Instead of using the marginal PMF $P(Y = y)$, which does not take into account any information about X , we should use a PMF that conditions on the event $X = x$, where x is the value we observed for X . This naturally leads us to consider conditional PMFs.

Definition 7.1.4 (Conditional PMF). For discrete r.v.s X and Y , the conditional PMF of Y given $X = x$ is

$$P(Y = y \mid X = x) = \frac{P(X = x, Y = y)}{P(X = x)}.$$

This is viewed as a function of y for fixed x .

Note that the conditional PMF (for fixed x) is a valid PMF. So we can define the conditional expectation of Y given $X = x$, denoted by $E(Y | X = x)$, in the same way that we defined $E(Y)$ except that we replace the PMF of Y with the conditional PMF of Y . Chapter 9 is devoted to conditional expectation.

Figure 7.3 illustrates the definition of conditional PMF. To condition on the event $X = x$, we first take the joint PMF and focus in on the vertical bars where X takes on the value x ; in the figure, these are shown in bold. All of the other vertical bars are irrelevant because they are inconsistent with the knowledge that $X = x$ occurred. Since the total height of the bold bars is the marginal probability $P(X = x)$, we then renormalize the conditional PMF by dividing by $P(X = x)$; this ensures that the conditional PMF will sum to 1. Therefore conditional PMFs are PMFs, just as conditional probabilities are probabilities. Notice that there is a different conditional PMF of Y for every possible value of X ; Figure 7.3 highlights just one of these conditional PMFs.




FIGURE 7.3 Conditional PMF of Y given $X = x$. The conditional PMF $P(Y = y | X = x)$ is obtained by renormalizing the column of the joint PMF that is compatible with the event $X = x$.

We can also relate the conditional distribution of Y given $X = x$ to that of X given $Y = y$, using Bayes' rule:

$$P(Y = y | X = x) = \frac{P(X = x | Y = y)P(Y = y)}{P(X = x)}.$$

And using LOTP, we have another way of getting the marginal PMF: the marginal PMF of X is a weighted average of the conditional PMFs $P(X = x | Y = y)$, where the weights are the probabilities $P(Y = y)$:

$$P(X = x) = \sum_y P(X = x | Y = y)P(Y = y).$$

Let's work through a numerical example to complement the graphs we've been looking at.

Example 7.1.5 (2×2 table). The simplest example of a discrete joint distribution is when X and Y are both Bernoulli r.v.s. In this case, the joint PMF is fully specified by the four values $P(X = 1, Y = 1)$, $P(X = 0, Y = 1)$, $P(X = 1, Y = 0)$, and $P(X = 0, Y = 0)$, so we can represent the joint PMF of X and Y using a 2×2 table.

This very simple scenario actually has an important place in statistics, as these so-called contingency tables are often used to study whether a treatment is associated with a particular outcome. In such scenarios, X may be the indicator of receiving the treatment, and Y may be the indicator of the outcome of interest.

For example, suppose we randomly sample an adult male from the United States population. Let X be the indicator of the sampled individual being a current smoker, and let Y be the indicator of his developing lung cancer at some point in his life. Suppose the joint PMF is as follows (these numbers are for illustrative purposes; they are not estimated from real data).

	$Y = 1$	$Y = 0$
$X = 1$	$\frac{5}{100}$	$\frac{20}{100}$
$X = 0$	$\frac{3}{100}$	$\frac{72}{100}$

To get the marginal probability $P(Y = 1)$, we add the probabilities in the two cells of the table where $Y = 1$. We do the same for $P(Y = 0)$, $P(X = 1)$, and $P(X = 0)$ and write these probabilities in the margins of the table (making "marginal" an appropriate name!).

	$Y = 1$	$Y = 0$	Total
$X = 1$	$\frac{5}{100}$	$\frac{20}{100}$	$\frac{25}{100}$
$X = 0$	$\frac{3}{100}$	$\frac{72}{100}$	$\frac{75}{100}$
Total	$\frac{8}{100}$	$\frac{92}{100}$	$\frac{100}{100}$

This shows that the marginal distribution of X is $\text{Bern}(0.25)$ and the marginal distribution of Y is $\text{Bern}(0.08)$. In words, the unconditional probability that the individual is a current smoker is 0.25, and the unconditional probability of his developing lung cancer is 0.08.

Now suppose we observe $X = 1$, i.e., the individual is a current smoker. We can then update our beliefs about his risk for lung cancer.

$$P(Y = 1 | X = 1) = \frac{P(X = 1, Y = 1)}{P(X = 1)} = \frac{5/100}{25/100} = 0.2,$$

so the conditional distribution of Y given $X = 1$ is $\text{Bern}(0.2)$. By a similar calculation, the conditional distribution of Y given $X = 0$ is $\text{Bern}(0.04)$. This tells us that the probability of developing lung cancer is 0.2 for current smokers but only 0.04 for non-smokers. $\square$

7.1.6. The word "marginal" has opposite meanings in economics and statistics. In economics it refers to a derivative, e.g., marginal revenue is the derivative of revenue with respect to quantity sold. In statistics it refers to an integral or sum, which can be thought of intuitively by writing totals in the margins of a table, as in the above example.

Armed with an understanding of joint, marginal, and conditional distributions, we can revisit the definition of independence that we introduced in Chapter 3.

Definition 7.1.7 (Independence of discrete r.v.s). Random variables X and Y are independent if for all x and y ,

$$F_{X,Y}(x, y) = F_X(x)F_Y(y).$$

If X and Y are discrete, this is equivalent to the condition

$$P(X = x, Y = y) = P(X = x)P(Y = y)$$

for all x, y , and it is also equivalent to the condition

$$P(Y = y | X = x) = P(Y = y)$$

for all x, y such that $P(X = x) > 0$.

Using the terminology from this chapter, the definition says that for independent r.v.s, the joint *CDF* factors into the product of the marginal *CDFs*, or that the joint PMF factors into the product of the marginal PMFs. Remember that in general, the marginal distributions do not determine the joint distribution: this is the reason why we wanted to study joint distributions in the first place! But in the special case of independence, the marginal distributions are all we need in order to specify the joint distribution; we can get the joint PMF by multiplying the marginal PMFs.

Another way of looking at independence is that all the conditional PMFs are the same as the marginal PMF. That is, starting with the marginal PMF of Y , no updating is necessary when we condition on $X = x$, regardless of what x is.

Example 7.1.8 (Independence in the 2×2 table). Returning to the table from the previous example, we can use these two views of independence to see why X and Y are not independent.

	$Y = 1$	$Y = 0$	Total
$X = 1$	$\frac{5}{100}$	$\frac{20}{100}$	$\frac{25}{100}$
$X = 0$	$\frac{3}{100}$	$\frac{72}{100}$	$\frac{75}{100}$
Total	$\frac{8}{100}$	$\frac{92}{100}$	$\frac{100}{100}$

First, the joint PMF is not the product of the marginal PMFs. For example,

$$P(X = 1, Y = 1) \neq P(X = 1)P(Y = 1)$$

Finding even one pair of values x and y such that

$$P(X = x, Y = y) \neq P(X = x)P(Y = y)$$

is enough to rule out independence.

Second, we found that the marginal distribution of Y is $\text{Bern}(0.08)$, whereas the conditional distribution of Y given $X = 1$ is $\text{Bern}(0.2)$ and the conditional distribution of Y given $X = 0$ is $\text{Bern}(0.04)$. Since conditioning on the value of X alters the distribution of Y , X and Y are not independent: learning whether or not the sampled individual is a current smoker gives us information about the probability that he will develop lung cancer.

Although we have found that X and Y are dependent, we cannot make conclusions about whether smoking causes lung cancer based on this association alone. As we learned from Simpson's paradox, misleading associations can arise when we fail to account for confounding variables.

We'll do one more example of a discrete joint distribution to round out this section. We've named it the chicken-egg story; in it, we use wishful thinking to find a joint PMF, and our efforts land us a surprising independence result.

Story 7.1.9 (Chicken-egg). Suppose a chicken lays a random number of eggs, N , where $N \sim \text{Pois}(\lambda)$. Each egg independently hatches with probability p and fails to hatch with probability $q = 1 - p$. Let X be the number of eggs that hatch and Y the number that do not hatch, so $X + Y = N$. What is the joint PMF of X and Y ?

Solution:

We seek the joint PMF $P(X = i, Y = j)$ for nonnegative integers i and j . Conditional on the total number of eggs N , the eggs are independent Bernoulli trials with probability of success p , so by the story of the Binomial, the conditional distributions of X and Y are $X | N = n \sim \text{Bin}(n, p)$ and $Y | N = n \sim \text{Bin}(n, q)$. Since our lives would be easier if only we knew the total number of eggs, let's use wishful thinking: condition on N and apply the law of total probability. This gives

$$P(X = i, Y = j) = \sum_{n=0}^{\infty} P(X = i, Y = j | N = n)P(N = n).$$

The sum is over all possible values of n , holding i and j fixed. But unless $n = i + j$, it is impossible for X to equal i and Y to equal j . For example, the only way there can be 5 hatched eggs and 6 unhatched eggs is if there are 11 eggs in total. So

$$P(X = i, Y = j | N = n) = 0$$

unless $n = i + j$, which means all other terms in the sum can be dropped:

$$P(X = i, Y = j) = P(X = i, Y = j | N = i + j)P(N = i + j)$$

Conditional on $N = i + j$, the events $X = i$ and $Y = j$ are exactly the same event, so keeping both is redundant. We'll keep $X = i$; the rest is a matter of plugging in the Binomial PMF to get $P(X = i | N = i + j)$ and the Poisson PMF to get $P(N = i + j)$. Thus,

$$\begin{aligned} P(X = i, Y = j) &= P(X = i | N = i + j)P(N = i + j) \\ &= \binom{i+j}{i} p^i q^j \cdot \frac{e^{-\lambda} \lambda^{i+j}}{(i+j)!} \\ &= \frac{e^{-\lambda p} (\lambda p)^i}{i!} \cdot \frac{e^{-\lambda q} (\lambda q)^j}{j!} \end{aligned}$$

The joint PMF factors into the product of the $\text{Pois}(\lambda p)$ PMF (as a function of i) and the $\text{Pois}(\lambda q)$ PMF (as a function of j). This tells us two elegant facts: (1) X and Y are independent, since their joint PMF is the product of their marginal PMFs, and (2) $X \sim \text{Pois}(\lambda p)$ and $Y \sim \text{Pois}(\lambda q)$.

At first it may seem deeply counterintuitive that X is independent of Y . Doesn't knowing that a lot of eggs hatched mean that there are probably not so many that didn't hatch? For a fixed number of eggs, this independence would be impossible: knowing the number of hatched eggs would perfectly determine the number of unhatched eggs. But in this example, the number of eggs is random, following a Poisson distribution, and this happens to be the right kind of randomness to make X and Y unconditionally independent. This is a very special property of the Poisson.

The chicken-egg story supplements the following result from Chapter 4:

Theorem 7.1.10. If $X \sim \text{Pois}(\lambda p)$, $Y \sim \text{Pois}(\lambda q)$, and X and Y are independent, then $N = X + Y \sim \text{Pois}(\lambda)$ and $X | N = n \sim \text{Bin}(n, p)$.

By the chicken-egg story, we now have the converse to this theorem.

Theorem 7.1.11. If $N \sim \text{Pois}(\lambda)$ and $X | N = n \sim \text{Bin}(n, p)$, then $X \sim \text{Pois}(\lambda p)$, $Y = N - X \sim \text{Pois}(\lambda q)$, and X and Y are independent.

7.1.12. In the chicken-egg story, it is not valid to say " $P(X = x | N = n) = P(X = x | X + Y = n) = P(X = x | Y = n - x)$ ", since in $P(X = x | N = n)$ we are not conditioning on $X = x$. Indeed, $P(X = x | N = n)$ is the $\text{Bin}(n, p)$ PMF, whereas $P(X = x | Y = n - x) = P(X = x)$ is the $\text{Pois}(\lambda p)$ PMF. This blunder of plugging in x for X on the right side of the conditioning bar illustrates how crucial it is to carefully distinguish between X and x , and between the left side and the right side of the conditioning bar.

7.1.2 Continuous

Once we have a handle on discrete joint distributions, it isn't much harder to consider continuous joint distributions. We simply make the now-familiar substitutions of integrals for sums and PDFs for PMFs, remembering that the probability of any individual point is now 0.

Formally, in order for X and Y to have a continuous joint distribution, we require that the joint CDF

$$F_{X,Y}(x, y) = P(X \leq x, Y \leq y)$$

be differentiable with respect to x and y . The partial derivative with respect to x and y is called the joint PDF. The joint PDF determines the joint distribution, as does the joint CDF.

Definition 7.1.13 (Joint PDF). If X and Y are continuous with joint CDF $F_{X,Y}$, their joint PDF is the derivative of the joint CDF with respect to x and y :

$$f_{X,Y}(x, y) = \frac{\partial^2}{\partial x \partial y} F_{X,Y}(x, y)$$

We require valid joint PDFs to be nonnegative and integrate to 1:

$$f_{X,Y}(x,y) \geq 0, \text{ and } \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_{X,Y}(x,y) dx dy = 1$$

In the univariate case, the PDF was the function we integrated to get the probability of an interval. Similarly, the joint PDF of two r.v.s is the function we integrate to get the probability of a two-dimensional region. For example,

$$P(X < 3, 1 < Y < 4) = \int_1^4 \int_{-\infty}^3 f_{X,Y}(x,y) dx dy$$

For a general region $A \subseteq \mathbb{R}^2$,

$$P((X,Y) \in A) = \iint_A f_{X,Y}(x,y) dx dy$$

Figure 7.4 shows a sketch of what a joint PDF of two r.v.s could look like. As usual with continuous r.v.s, we need to keep in mind that the height of the surface $f_{X,Y}(x,y)$ at a single point does not represent a probability. The probability of any specific point in the plane is 0. Now that we've gone up a dimension, the probability of any line or curve in the plane is also 0. The only way we can get nonzero probability is by integrating over a region of positive area in the xy -plane.

When we integrate the joint PDF over a region A , we are calculating the volume under the surface of the joint PDF and above A . Thus, probability is represented by volume under the joint PDF. The total volume under a valid joint PDF is 1.




FIGURE 7.4
Joint PDF of continuous r.v.s X and Y .

In the discrete case, we get the marginal PMF of X by summing over all possible values of Y in the joint PMF. In the continuous case, we get the marginal PDF of X by integrating over all possible values of Y in the joint PDF.

Definition 7.1.14 (Marginal PDF). For continuous r.v.s X and Y with joint PDF $f_{X,Y}$, the marginal PDF of X is

$$f_X(x) = \int_{-\infty}^{\infty} f_{X,Y}(x,y) dy$$

This is the PDF of X , viewing X individually rather than jointly with Y . To simplify notation, we have mainly been looking at the joint distribution of two r.v.s rather than n r.v.s, but marginalization works analogously with any number of variables. For example,

if we have the joint PDF of X, Y, Z, W but want the joint PDF of X, W , we just have to integrate over all possible values of Y and Z :

$$f_{X,W}(x, w) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_{X,Y,Z,W}(x, y, z, w) dy dz$$

Conceptually this is easy-just integrate over the unwanted variables to get the joint PDF of the wanted variables-but computing it may or may not be easy.

Returning to the case of the joint distribution of two r.v.s X and Y , let's consider how to update our distribution for Y after observing the value of X , using the conditional PDF.

Definition 7.1.15 (Conditional PDF). For continuous r.v.s X and Y with joint PDF $f_{X,Y}$, the conditional PDF of Y given $X = x$ is

$$f_{Y|X}(y | x) = \frac{f_{X,Y}(x, y)}{f_X(x)},$$

for all x with $f_X(x) > 0$. This is considered as a function of y for fixed x . As a convention, in order to make $f_{Y|X}(y | x)$ well-defined for all real x , let $f_{Y|X}(y | x) = 0$ for all x with $f_X(x) = 0$.

Notation 7.1.16. The subscripts that we place on all the f 's are just to remind us that we have three different functions on our plate. We could just as well write $g(y | x) = f(x, y)/h(x)$, where f is the joint PDF, h is the marginal PDF of X , and g is the conditional PDF of Y given $X = x$, but that makes it more difficult to remember which letter stands for which function.

Figure 7.5 illustrates the definition of conditional PDF. We take a vertical slice of the joint PDF corresponding to the observed value of X . Since the total area under this slice is $f_X(x)$, we then divide by $f_X(x)$ to ensure that the conditional PDF will have an area of 1. So the conditional PDF of Y given $X = x$ satisfies the properties of a valid PDF, for any x in the support of X .




FIGURE 7.5 Conditional PDF of Y given $X = x$. The conditional PDF $f_{Y|X}(y | x)$ is obtained by renormalizing the slice of the joint PDF at the fixed value x .

7.1.17. How can we speak of conditioning on $X = x$ for X a continuous r.v., considering that this event has probability 0? Rigorously speaking, we are actually conditioning on the event that X falls within a small interval containing x , say $X \in (x - \epsilon, x + \epsilon)$, and then taking a limit as ϵ approaches 0 from the right. We will not fuss over this technicality; fortunately, many important results such as Bayes' rule work in the continuous case exactly as one would hope.

Note that we can recover the joint PDF $f_{X,Y}$ if we have the conditional PDF $f_{Y|X}$ and the corresponding marginal f_X :

$$f_{X,Y}(x, y) = f_{Y|X}(y | x) f_X(x).$$

Similarly, we can recover the joint PDF if we have $f_{X|Y}$ and f_Y :

$$f_{X,Y}(x, y) = f_{X|Y}(x | y) f_Y(y).$$

This allows us to develop continuous versions of Bayes' rule and LOTP. The continuous versions are analogous to the discrete versions, with probability density functions in place of probabilities and integrals in place of sums.

Theorem 7.1.18 (Continuous form of Bayes' rule and LOTP). For continuous r.v.s X and Y , we have the following continuous form of Bayes' rule:

$$f_{Y|X}(y | x) = \frac{f_{X|Y}(x | y)f_Y(y)}{f_X(x)}, \text{ for } f_X(x) > 0$$

And we have the following continuous form of the law of total probability:

$$f_X(x) = \int_{-\infty}^{\infty} f_{X|Y}(x | y)f_Y(y)dy$$

Proof. By definition of conditional PDFs, we have

$$f_{Y|X}(y | x)f_X(x) = f_{X,Y}(x, y) = f_{X|Y}(x | y)f_Y(y)$$

The continuous version of Bayes' rule follows immediately from dividing by $f_X(x)$. The continuous version of LOTP follows immediately from integrating with respect to y :

$$f_X(x) = \int_{-\infty}^{\infty} f_{X,Y}(x, y)dy = \int_{-\infty}^{\infty} f_{X|Y}(x | y)f_Y(y)dy$$

Out of curiosity, let's see what would have happened if we had plugged in the other expression for $f_{X,Y}(x, y)$ instead in the proof of LOTP:

$$f_X(x) = \int_{-\infty}^{\infty} f_{X,Y}(x, y)dy = \int_{-\infty}^{\infty} f_{Y|X}(y | x)f_X(x)dy = f_X(x) \int_{-\infty}^{\infty} f_{Y|X}(y | x)dy$$

This just says that, for any x with $f_X(x) > 0$,

$$\int_{-\infty}^{\infty} f_{Y|X}(y | x)dy = 1$$

confirming the fact that conditional PDFs must integrate to 1 .

We now have versions of Bayes' rule and LOTP for two discrete r.v.s and for two continuous r.v.s. Better yet, there are also versions when we have one discrete r.v. and one continuous r.v. After understanding the discrete versions, it is easy to remember and use the other versions since they are analogous, replacing probabilities by PDFs when appropriate. For example, for X discrete and Y continuous, we have the following version of LOTP:

$$P(X = x) = \int_{-\infty}^{\infty} P(X = x | Y = y)f_Y(y)dy$$

Taking X to be the indicator r.v. of an event A and $x = 1$, we have an expression for a general probability $P(A)$ based on conditioning on a continuous r.v. Y :

$$P(A) = \int_{-\infty}^{\infty} P(A | Y = y)f_Y(y)dy$$

Here are the four versions of Bayes' rule, summarized in a table.

	Y discrete	Y continuous
X discrete	$P(Y = y X = x) = \frac{P(X=x Y=y)P(Y=y)}{P(X=x)}$	$f_Y(y X = x) = \frac{P(X=x Y=y)f_Y(y)}{P(X=x)}$
X continuous	$P(Y = y X = x) = \frac{f_{X Y}(x y)P(Y=y)}{f_X(x)}$	$f_{Y X}(y x) = \frac{f_{X Y}(x y)f_Y(y)}{f_X(x)}$

And here are the four versions of LOTP, summarized in a table. The top row gives expressions for $P(X = x)$, while the bottom row gives expressions for $f_X(x)$.

	Y discrete	Y continuous
X discrete	$\sum_y P(X = x Y = y)P(Y = y)$	$\int_{-\infty}^{\infty} P(X = x Y = y)f_Y(y)dy$
X continuous	$\sum_y f_X(x Y = y)P(Y = y)$	$\int_{-\infty}^{\infty} f_{X Y}(x y)f_Y(y)dy$

Finally, let's discuss the definition of independence for continuous r.v.s; then we'll turn to concrete examples. As in the discrete case, we can view independence of continuous r.v.s in two ways. One is that the joint CDF factors into the product of the marginal CDFs, or the joint PDF factors into the product of the marginal PDFs. The other is that the conditional PDF of Y given $X = x$ is the same as the marginal PDF of Y , so conditioning on X provides no information about Y .

Definition 7.1.19 (Independence of continuous r.v.s). Random variables X and Y are independent if for all x and y ,

$$F_{X,Y}(x, y) = F_X(x)F_Y(y).$$

If X and Y are continuous with joint PDF $f_{X,Y}$, this is equivalent to the condition

$$f_{X,Y}(x, y) = f_X(x)f_Y(y)$$

for all x, y , and it is also equivalent to the condition

$$f_{Y|X}(y | x) = f_Y(y)$$

for all x, y such that $f_X(x) > 0$.

7.1.20. The marginal PDF of Y , $f_Y(y)$, is a function of y only; it cannot depend on x in any way. The conditional PDF $f_{Y|X}(y | x)$ can depend on x in general. Only in the special case of independence is $f_{Y|X}(y | x)$ free of x .

Sometimes we have a joint PDF for X and Y that factors as a function of x times a function of y , without knowing in advance whether these functions are the marginal PDFs, or even whether they are valid PDFs. The next result addresses this situation.

Proposition 7.1.21. Suppose that the joint PDF $f_{X,Y}$ of X and Y factors as

$$f_{X,Y}(x, y) = g(x)h(y)$$

for all x and y , where g and h are nonnegative functions. Then X and Y are independent. Also, if either g or h is a valid PDF, then the other one is a valid PDF too and g and h are the marginal PDFs of X and Y , respectively. (The analogous result in the discrete case also holds.)

Proof. Let $c = \int_{-\infty}^{\infty} h(y)dy$. Multiplying and dividing by c , we can write

$$f_{X,Y}(x, y) = cg(x) \cdot \frac{h(y)}{c}.$$

(The point of this is that $h(y)/c$ is a valid PDF.) Then the marginal PDF of X is

$$f_X(x) = \int_{-\infty}^{\infty} f_{X,Y}(x, y)dy = cg(x) \int_{-\infty}^{\infty} \frac{h(y)}{c}dy = cg(x)$$

It follows that $\int_{-\infty}^{\infty} cg(x)dx = 1$ since a marginal PDF is a valid PDF (knowing the integral of h gave us the integral of g for free!). Then the marginal PDF of Y is

$$f_Y(y) = \int_{-\infty}^{\infty} f_{X,Y}(x, y)dx = \frac{h(y)}{c} \int_{-\infty}^{\infty} cg(x)dx = \frac{h(y)}{c}$$

Thus, X and Y are independent with PDFs $cg(x)$ and $h(y)/c$, respectively. If g or h is already a valid PDF, then $c = 1$, so the other one is also a valid PDF. $\square$

7.1.22. In the above proposition, we need the joint PDF to factor as a function of x times a function of y for all (x, y) in the plane $\mathbb{R}^2$, not just for (x, y) with $f_{X,Y}(x, y) > 0$. The reason for this is illustrated in the next example.

A simple case of a continuous joint distribution is when the joint PDF is constant over some region in the plane. In the following example, we'll compare a joint PDF that is constant on a square to a joint PDF that is constant on a disk.

Example 7.1.23 (Uniform on a region in the plane). Let (X, Y) be a completely random point in the square $\{(x, y) : x, y \in [0, 1]\}$, in the sense that the joint PDF of X and Y is constant over the square and 0 outside of it:

$$f_{X,Y}(x, y) = \begin{cases} 1 & \text{if } x, y \in [0, 1] \\ 0 & \text{otherwise} \end{cases}$$

The constant 1 is chosen so that the joint PDF will integrate to 1. This distribution is called the Uniform distribution on the square.

Intuitively, it makes sense that X and Y should be $\text{Unif}(0, 1)$ marginally. We can check this by computing

$$f_X(x) = \int_0^1 f_{X,Y}(x, y) dy = \int_0^1 1 dy = 1$$

and similarly for f_Y . Furthermore, X and Y are independent, since the joint PDF factors into the product of the marginal PDFs (this just reduces to $1 = 1 \cdot 1$, but it's important to note that the value of X does not constrain the possible values of Y). So the conditional distribution of Y given $X = x$ is $\text{Unif}(0, 1)$, regardless of x .

Now let (X, Y) be a completely random point in the unit disk $\{(x, y) : x^2 + y^2 \leq 1\}$, with joint PDF

$$f_{X,Y}(x, y) = \begin{cases} \frac{1}{\pi} & \text{if } x^2 + y^2 \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

Again, the constant $1/\pi$ is chosen to make the joint PDF integrate to 1; the value follows from the fact that the integral of 1 over some region in the plane is the area of that region.

Note that X and Y are not independent, since in general, knowing the value of X constrains the possible values of Y : larger values of $|X|$ restrict Y to be in a smaller range. It would fall into 7.1.22 disastrously to conclude independence from the fact that $f_{X,Y}(x, y) = g(x)h(y)$ for all (x, y) in the disk, where $g(x) = 1/\pi$ and $h(y) = 1$ are constant functions. To see from the definition that X and Y are not independent, note that, for example, $f_{X,Y}(0.9, 0.9) = 0$ since $(0.9, 0.9)$ is not in the unit disk, but $f_X(0.9)f_Y(0.9) \neq 0$ since 0.9 is in the supports of both X and Y .

The marginal distribution of X is now

$$f_X(x) = \int_{-\sqrt{1-x^2}}^{\sqrt{1-x^2}} \frac{1}{\pi} dy = \frac{2}{\pi} \sqrt{1-x^2}, \quad -1 \leq x \leq 1.$$

By symmetry, $f_Y(y) = \frac{2}{\pi} \sqrt{1-y^2}$. Note that the marginal distributions of X and Y are not Uniform on $[-1, 1]$; rather, X and Y are more likely to fall near 0 than near ± 1 .




FIGURE 7.6 Bird's-eye view of the Uniform joint PDF on the unit disk. Conditional on $X = x$, Y is restricted to the interval $[-\sqrt{1-x^2}, \sqrt{1-x^2}]$.

Suppose we observe $X = x$. As illustrated in Figure 7.6, this constrains Y to lie in the interval $[-\sqrt{1-x^2}, \sqrt{1-x^2}]$. Specifically, the conditional distribution of Y given $X = x$ is

$$f_{Y|X}(y | x) = \frac{f_{X,Y}(x, y)}{f_X(x)} = \frac{\frac{1}{\pi}}{\frac{2}{\pi}\sqrt{1-x^2}} = \frac{1}{2\sqrt{1-x^2}}$$

for $-\sqrt{1-x^2} \leq y \leq \sqrt{1-x^2}$, and 0 otherwise. This conditional PDF is constant as a function of y , which tells us that the conditional distribution of Y is Uniform on the interval $[-\sqrt{1-x^2}, \sqrt{1-x^2}]$. The fact that the conditional PDF is not free of x confirms the fact that X and Y are not independent.

In general, for a region R in the plane, the Uniform distribution on R is defined to have joint PDF that is constant inside R and 0 outside R . The constant is the reciprocal of the area of R . If R is the rectangle $\{(x, y) : a \leq x \leq b, c \leq y \leq d\}$, then X and Y will be independent; unlike for a disk, the vertical slices of a rectangle all look the same. But for any region where the value of X constrains the possible values of Y or vice versa, X and Y will not be independent.

As another example of working with joint PDFs, let's consider a question that comes up often when dealing with Exponentials of different rates.

Example 7.1.24 (Comparing Exponentials of different rates). Let $T_1 \sim \text{Expo}(\lambda_1)$ and $T_2 \sim \text{Expo}(\lambda_2)$ be independent. Find $P(T_1 < T_2)$. For example, T_1 could be the lifetime of a refrigerator and T_2 could be the lifetime of a stove (if we are willing to assume Exponential distributions for these), and then $P(T_1 < T_2)$ is the probability that the refrigerator fails before the stove. We know from Chapter 5 that $\min(T_1, T_2) \sim \text{Expo}(\lambda_1 + \lambda_2)$, which tells us about when the first appliance failure will occur, but we may also want to know about which appliance will fail first.

Solution:

We just need to integrate the joint PDF of T_1 and T_2 over the appropriate region, which is all (t_1, t_2) with $t_1 > 0, t_2 > 0$, and $t_1 < t_2$. This yields

$$\begin{aligned}
P(T_1 < T_2) &= \int_0^\infty \int_0^{t_2} \lambda_1 e^{-\lambda_1 t_1} \lambda_2 e^{-\lambda_2 t_2} dt_1 dt_2 \\
&= \int_0^\infty \left(\int_0^{t_2} \lambda_1 e^{-\lambda_1 t_1} dt_1 \right) \lambda_2 e^{-\lambda_2 t_2} dt_2 \\
&= \int_0^\infty (1 - e^{-\lambda_1 t_2}) \lambda_2 e^{-\lambda_2 t_2} dt_2 \\
&= 1 - \int_0^\infty \lambda_2 e^{-(\lambda_1 + \lambda_2)t_2} dt_2 \\
&= 1 - \frac{\lambda_2}{\lambda_1 + \lambda_2} \\
&= \frac{\lambda_1}{\lambda_1 + \lambda_2}
\end{aligned}$$

This result makes sense intuitively if we interpret λ_1 and λ_2 as rates. For example, if refrigerators have twice the failure rate of stoves, then it says that the odds are 2 to 1 in favor of the refrigerator failing first. As a simple check, note that the answer reduces to $1/2$ when $\lambda_1 = \lambda_2$, which must be true by symmetry.

An alternative method of getting the same result is to use LOTP to condition on T_1 (or condition on T_2). A third approach, using a story about Poisson processes, is in Chapter 13.

Our last example in this section demonstrates how we can use the joint distribution of X and Y to derive the distribution of a function of X and Y .

Example 7.1.25 (Cauchy PDF). Let X and Y be i.i.d. $\mathcal{N}(0, 1)$, and let $T = X/Y$. (We can define T arbitrarily in the case $Y = 0$; the choice of how to define T in that case has no effect on the distribution of T , since $P(Y = 0) = 0$.) The distribution of T is a famous named distribution called the Cauchy distribution, and we will encounter it again in later chapters. Meanwhile, find the PDF of T .

Solution:

We'll find an expression for the CDF of T first, and then differentiate to get the PDF. We can write

$$F_T(t) = P(T \leq t) = P\left(\frac{X}{Y} \leq t\right) = P\left(\frac{X}{|Y|} \leq t\right),$$

since the r.v.s $\frac{X}{Y}$ and $\frac{X}{|Y|}$ are identically distributed by the symmetry of the standard Normal distribution. Now $|Y|$ is nonnegative, so we can multiply by it on both sides without reversing the direction of the inequality. Thus, we are interested in finding

$$F_T(t) = P(X \leq t|Y|).$$

We calculate this probability by integrating the joint PDF of X and Y over the region where $X \leq t|Y|$ holds. The joint PDF of X and Y is just the product of the marginal PDFs, by independence. So

$$\begin{aligned}
F_T(t) &= P(X \leq t|Y|) \\
&= \int_{-\infty}^{\infty} \int_{-\infty}^{t|y|} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} dx dy
\end{aligned}$$

Note that the inner limits of integration (the limits for x) depend on y , whereas the outer limits of integration (the limits for y) can't depend on x (see the math appendix for more

about limits of integration with multiple integrals). With some manipulations, we can get the double integral down to a single integral:

$$\begin{aligned} F_T(t) &= \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} \left(\int_{-\infty}^{t|y|} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \right) dy \\ &= \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} \Phi(t|y|) dy \\ &= \sqrt{\frac{2}{\pi}} \int_0^{\infty} e^{-y^2/2} \Phi(ty) dy \end{aligned}$$

Alternatively, we can get the same result without doing a double integral by using one of the versions of LOTP. Letting I be the indicator r.v. for the event $X \leq t|Y|$, we again have

$$\begin{aligned} P(I = 1) &= \int_{-\infty}^{\infty} P(I = 1 \mid Y = y) f_Y(y) dy \\ &= \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} \Phi(t|y|) dy \end{aligned}$$

At this point we seem to be stuck with an integral we don't know how to do. Fortunately, we were asked to find the PDF, not the CDF, so instead of evaluating the integral, we can just differentiate it with respect to t (not with respect to y , which is a dummy variable). We are permitted to interchange the order of integration and differentiation under mild technical conditions, which are met here. (This technique is known as differentiation under the integral sign, or DUThis for short.) Then

$$\begin{aligned} f_T(t) &= F'_T(t) = \sqrt{\frac{2}{\pi}} \int_0^{\infty} \frac{\partial}{\partial t} \left(e^{-y^2/2} \Phi(ty) \right) dy \\ &= \sqrt{\frac{2}{\pi}} \int_0^{\infty} y e^{-y^2/2} \varphi(ty) dy \\ &= \frac{1}{\pi} \int_0^{\infty} y e^{-\frac{(1+t^2)y^2}{2}} dy \\ &= \frac{1}{\pi(1+t^2)} \end{aligned}$$

using the substitution $u = (1+t^2)y^2/2$, $du = (1+t^2)y dy$ for the final step. So the PDF of T is

$$f_T(t) = \frac{1}{\pi(1+t^2)}, \quad t \in \mathbb{R}$$

Since

$$\int_{-\infty}^{\infty} \frac{1}{1+t^2} dt = \arctan(\infty) - \arctan(-\infty) = \pi$$

we have obtained a valid PDF. If we want to obtain the CDF too, we can integrate the PDF over the appropriate interval:

$$F_T(t) = \int_{-\infty}^t \frac{1}{\pi(1+u^2)} du = \frac{1}{\pi} \arctan(t) + \frac{1}{2}$$

As we mentioned, the distribution of T is called the Cauchy distribution. The Cauchy PDF is similar in shape to the Normal bell curve, but with tails that decay less quickly to 0. Figure

7.7 superimposes the Cauchy and standard Normal PDFs; the heavier tails of the Cauchy PDF are evident.

An interesting fact about the Cauchy distribution is that although the PDF is symmetric about 0, its expected value does not exist, since the integral $\int_{-\infty}^{\infty} \frac{t}{\pi(1+t^2)} dt$ diverges: note that for large t , $\frac{t}{1+t^2} \approx \frac{1}{t}$, and $\int_1^{\infty} \frac{1}{t} dt = \infty$. It would be a blunder to write " $E\left(\frac{X}{Y}\right) = E(X)E\left(\frac{1}{Y}\right) = 0 \cdot E\left(\frac{1}{Y}\right) = 0$ ", since $E\left(\frac{1}{Y}\right)$ also does not exist.




FIGURE 7.7 Cauchy PDF (dark) and $\mathcal{N}(0,1)$ PDF (light). The Cauchy distribution has much heavier tails than the Normal distribution.

7.1.3 Hybrid

It is also possible that we could be interested in the joint distribution of a discrete r.v. and a continuous r.v. This case was mentioned when discussing the four forms of Bayes' rule and LOTP. Conceptually it is analogous to the other cases, but since the notation can be tricky, we'll work through an example that sheds light on it.

Example 7.1.26 (Which company made the lightbulb?). A lightbulb was manufactured by one of two companies. Bulbs that are made by Company 0 last an $\text{Expo}(\lambda_0)$ amount of time, and bulbs made by Company 1 last an $\text{Expo}(\lambda_1)$ amount of time, with $\lambda_0 < \lambda_1$. The bulb of interest here was made by Company 0 with probability p_0 and by Company 1 with probability $p_1 = 1 - p_0$, but from inspecting the bulb we don't know which company made it.

Let T be how long the bulb lasts, and I be the indicator of it having been made by Company 1.

- Find the CDF and PDF of T .
- Does T have the memoryless property?
- Find the conditional distribution of I given $T = t$. What happens to this as $t \rightarrow \infty$?

Solution:

Since T is a continuous r.v. and I is discrete, the joint distribution of T and I is a hybrid, as illustrated in Figure 7.8. In a joint PDF of two continuous r.v.s, there are infinitely many vertical slices of the joint PDF that we can take, each corresponding to a different conditional PDF. Here there are only two conditional PDFs of T , one for $I = 0$ and one for $I = 1$. As stated in the problem, the conditional distribution of T given $I = 0$ is $\text{Expo}(\lambda_0)$ and given $I = 1$ is $\text{Expo}(\lambda_1)$. The marginal distribution of I is $\text{Bern}(p_1)$.




FIGURE 7.8
Hybrid joint distribution of T and I .

Thus, we are given the joint distribution of T and I in terms of (1) the marginal distribution of I and (2) the conditional distribution of T given I . The problem then asks us to flip it around and find (1) the marginal distribution of T and (2) the conditional distribution of I given T . Phrased in this way, it becomes clear that LOTP and Bayes' rule will be our friends.

(a) In this part we are asked to derive the marginal distribution of T . For the CDF, we use the law of total probability, conditioning on I :

$$\begin{aligned} F_T(t) &= P(T \leq t) = P(T \leq t \mid I = 0)p_0 + P(T \leq t \mid I = 1)p_1 \\ &= (1 - e^{-\lambda_0 t})p_0 + (1 - e^{-\lambda_1 t})p_1 \\ &= 1 - p_0 e^{-\lambda_0 t} - p_1 e^{-\lambda_1 t} \end{aligned}$$

for $t > 0$. The marginal PDF is the derivative of the CDF:

$$f_T(t) = p_0 \lambda_0 e^{-\lambda_0 t} + p_1 \lambda_1 e^{-\lambda_1 t}, \text{ for } t > 0.$$

We could also have gotten this directly from the " X continuous, Y discrete" version of LOTP, but we did not write out a proof of that version of LOTP, and working through this example helps show why that version of LOTP works.

(b) Since $\lambda_0 \neq \lambda_1$, the above expression for the PDF does not reduce to the form $\lambda e^{-\lambda t}$. So the distribution of T is not Exponential, which implies that it does not have the memoryless property. (The distribution of T is called a mixture of two Exponentials.)

(c) Using a hybrid form of Bayes' rule, which is the " X continuous, Y discrete" version in the table of versions of Bayes' rule, we have

$$P(I = 1 \mid T = t) = \frac{f_T(t \mid I = 1)P(I = 1)}{f_T(t)},$$

where $f_T(t \mid I = 1)$ is the conditional PDF of T given $I = 1$, evaluated at t . Using the fact that $T \mid I = 1 \sim \text{Expo}(\lambda_1)$ and the marginal PDF derived in (a),

$$P(I = 1 \mid T = t) = \frac{p_1 \lambda_1 e^{-\lambda_1 t}}{p_0 \lambda_0 e^{-\lambda_0 t} + p_1 \lambda_1 e^{-\lambda_1 t}} = \frac{p_1 \lambda_1}{p_0 \lambda_0 e^{(\lambda_1 - \lambda_0)t} + p_1 \lambda_1}.$$

Thus the conditional distribution of I given $T = t$ is Bernoulli with this probability of success. This probability goes to 0 as $t \rightarrow \infty$, which makes sense intuitively: the longer the

bulb lasts, the more confident we will be that it was made by Company 0, as their bulbs have a lower failure rate λ and a higher life expectancy $\frac{1}{\lambda}$.

4.7.2 7.2 2D Lotus

The two-dimensional version of LOTUS lets us calculate the expectation of a random variable that is a function of two random variables X and Y , using the joint distribution of X and Y .

Theorem 7.2.1 (2D LOTUS). Let g be a function from $\mathbb{R}^2$ to $\mathbb{R}$. If X and Y are discrete, then

$$E(g(X, Y)) = \sum_x \sum_y g(x, y) P(X = x, Y = y).$$

If X and Y are continuous with joint PDF $f_{X,Y}$, then

$$E(g(X, Y)) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(x, y) f_{X,Y}(x, y) dx dy$$

Like its 1D counterpart, 2D LOTUS saves us from having to find the distribution of $g(X, Y)$ in order to calculate its expectation. Instead, having the joint PMF or joint PDF of X and Y is enough.

One use of 2D LOTUS is to find the expected distance between two r.v.s.

Example 7.2.2 (Expected distance between two Uniforms). Let X and Y be i.i.d. $\text{Unif}(0, 1)$ r.v.s. Find $E(|X - Y|)$.

Solution:

Since the joint PDF is 1 on the unit square $\{(x, y) : x, y \in [0, 1]\}$, 2D LOTUS gives

$$\begin{aligned} E(|X - Y|) &= \int_0^1 \int_0^1 |x - y| dx dy \\ &= \int_0^1 \int_y^1 (x - y) dx dy + \int_0^1 \int_0^y (y - x) dx dy \\ &= 2 \int_0^1 \int_y^1 (x - y) dx dy = 1/3 \end{aligned}$$

First we broke up the integral into two parts so we could eliminate the absolute value; then we used symmetry.

Incidentally, by solving this problem, we have also figured out the expected value of $M = \max(X, Y)$ and $L = \min(X, Y)$. Since $M + L$ is the same r.v. as $X + Y$ and $M - L$ is the same r.v. as $|X - Y|$,

$$\begin{aligned} E(M + L) &= E(X + Y) = 1 \\ E(M - L) &= E(|X - Y|) = 1/3 \end{aligned}$$

This is a system of two equations and two unknowns, which we can solve to get $E(M) = 2/3$ and $E(L) = 1/3$. As a check, $E(M)$ exceeds $E(L)$, as it should, and $E(M)$ and $E(L)$ are equidistant from $1/2$, as should be the case by symmetry.

Example 7.2.3 (Expected distance between two Normals). For $X, Y \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$, find $E(|X - Y|)$.

Solution:

We could again use 2D LOTUS, giving

$$E(|X - Y|) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} |x - y| \frac{1}{\sqrt{2\pi}} e^{-x^2/2} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} dx dy$$

but an easier solution uses the fact that the sum or difference of independent Normals is Normal, as we proved using MGFs in Chapter 6. Then $X - Y \sim \mathcal{N}(0, 2)$, so we can write $X - Y = \sqrt{2}Z$ where $Z \sim \mathcal{N}(0, 1)$, and $E(|X - Y|) = \sqrt{2}E|Z|$. Thus, we have reduced a 2D LOTUS to a 1D LOTUS! It was shown in Example 5.4.7 that

$$E|Z| = \sqrt{\frac{2}{\pi}}$$

so

$$E(|X - Y|) = \frac{2}{\sqrt{\pi}}$$

We can also use 2D LOTUS to give another proof of linearity of expectation.

Example 7.2.4 (Linearity via 2D LOTUS). Let X and Y be continuous r.v.s (the analogous method also works in the discrete case). By 2D LOTUS,

$$\begin{aligned} E(X + Y) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} (x + y) f_{X,Y}(x, y) dx dy \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x f_{X,Y}(x, y) dx dy + \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} y f_{X,Y}(x, y) dx dy \\ &= E(X) + E(Y) \end{aligned}$$

This is a short proof of linearity of expectation. For the last step, we used 2D LOTUS and the fact that X is a function of X and Y (that happens to be degenerate in the sense that it doesn't involve Y), and similarly for Y . Another way to get the last step is to write

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} y f_{X,Y}(x, y) dx dy = \int_{-\infty}^{\infty} y \int_{-\infty}^{\infty} f_{X,Y}(x, y) dx dy = \int_{-\infty}^{\infty} y f_Y(y) dy = E(Y)$$

where we took y out from the inner integral (since y is held constant when integrating with respect to x) and then recognized the marginal PDF of Y . For the $E(X)$ term we can first swap the order of integration, from $dx dy$ to $dy dx$, and then the same argument that we used for the $E(Y)$ term can be applied.

4.7.3 7.3 Covariance and Correlation

Just as the mean and variance provided single-number summaries of the distribution of a single r.v., covariance is a single-number summary of the joint distribution of two r.v.s. Roughly speaking, covariance measures a tendency of two r.v.s to go up or down together, relative to their means: positive covariance between X and Y indicates that when X goes up, Y also tends to go up, and negative covariance indicates that when X goes up, Y tends to go down. Here is the precise definition.

Definition 7.3.1 (Covariance). The covariance between r.v.s X and Y is

$$\text{Cov}(X, Y) = E((X - EX)(Y - EY)).$$

Multiplying this out and using linearity, we have an equivalent expression:

$$\text{Cov}(X, Y) = E(XY) - E(X)E(Y).$$

Let's think about the definition intuitively. If X and Y tend to move in the same direction, then $X - EX$ and $Y - EY$ will tend to be either both positive or both negative, so $(X - EX)(Y - EY)$ will be positive on average, giving a positive covariance. If X and Y tend to move in opposite directions, then $X - EX$ and $Y - EY$ will tend to have opposite signs, giving a negative covariance.

If X and Y are independent, then their covariance is zero. We say that r.v.s with zero covariance are uncorrelated.

Theorem 7.3.2. If X and Y are independent, then they are uncorrelated.

Proof. We'll show this in the case where X and Y are continuous, with PDFs f_X and f_Y . Since X and Y are independent, their joint PDF is the product of the marginal PDFs. By 2D LOTUS,

$$\begin{aligned} E(XY) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} xy f_X(x) f_Y(y) dx dy \\ &= \int_{-\infty}^{\infty} y f_Y(y) \left(\int_{-\infty}^{\infty} x f_X(x) dx \right) dy \\ &= \int_{-\infty}^{\infty} x f_X(x) dx \int_{-\infty}^{\infty} y f_Y(y) dy \\ &= E(X)E(Y) \end{aligned}$$

The proof in the discrete case is the same, with PMFs instead of PDFs.

The converse of this theorem is false: just because X and Y are uncorrelated does not mean they are independent. For example, let $X \sim \mathcal{N}(0, 1)$, and let $Y = X^2$. Then $E(XY) = E(X^3) = 0$ because the odd moments of the standard Normal distribution are equal to 0 by symmetry. Thus X and Y are uncorrelated,

$$\text{Cov}(X, Y) = E(XY) - E(X)E(Y) = 0 - 0 = 0$$

but they are certainly not independent: Y is a function of X , so knowing X gives us perfect information about Y . Covariance is a measure of linear association, so r.v.s can be dependent in nonlinear ways and still have zero covariance, as this example demonstrates. The bottom right plot of Figure 7.9 shows draws from the joint distribution of X and Y in this example. The other three plots illustrate positive correlation, negative correlation, and independence.

Covariance has the following key properties.

1. $\text{Cov}(X, X) = \text{Var}(X)$.
2. $\text{Cov}(X, Y) = \text{Cov}(Y, X)$.
3. $\text{Cov}(X, c) = 0$ for any constant c .
4. $\text{Cov}(aX, Y) = a \text{Cov}(X, Y)$ for any constant a .
5. $\text{Cov}(X + Y, Z) = \text{Cov}(X, Z) + \text{Cov}(Y, Z)$.
6. $\text{Cov}(X + Y, Z + W) = \text{Cov}(X, Z) + \text{Cov}(X, W) + \text{Cov}(Y, Z) + \text{Cov}(Y, W)$.
7. $\text{Var}(X + Y) = \text{Var}(X) + \text{Var}(Y) + 2 \text{Cov}(X, Y)$. For n r.v.s $X_1, \dots, X_n$,

$$\text{Var}(X_1 + \cdots + X_n) = \text{Var}(X_1) + \cdots + \text{Var}(X_n) + 2 \sum_{i < j} \text{Cov}(X_i, X_j)$$




FIGURE 7.9

Draws from the joint distribution of (X, Y) under various dependence structures. Top left: X and Y are positively correlated. Top right: X and Y are negatively correlated. Bottom left: X and Y are independent, hence uncorrelated. Bottom right: Y is a deterministic function of X , but X and Y are uncorrelated.

The first five properties follow readily from the definition and basic properties of expectation. Property 6 follows from Property 2 and Property 5, by expanding

$$\begin{aligned} \text{Cov}(X + Y, Z + W) &= \text{Cov}(X, Z + W) + \text{Cov}(Y, Z + W) \\ &= \text{Cov}(Z + W, X) + \text{Cov}(Z + W, Y) \\ &= \text{Cov}(Z, X) + \text{Cov}(W, X) + \text{Cov}(Z, Y) + \text{Cov}(W, Y) \\ &= \text{Cov}(X, Z) + \text{Cov}(X, W) + \text{Cov}(Y, Z) + \text{Cov}(Y, W). \end{aligned}$$

Property 7 follows from writing the variance of an r.v. as its covariance with itself (by Property 1) and then using Property 6 repeatedly.

We have now fulfilled our promise from Chapter 4 that for independent r.v.s, the variance of the sum is the sum of the variances:

$$\text{Var}\left(\sum_{j=1}^n X_j\right) = \sum_{j=1}^n \text{Var}(X_j) \text{ if } X_1, \dots, X_n \text{ are independent.}$$

By Theorem 7.3.2, independent r.v.s are uncorrelated, so in that case all the covariance terms drop out of the expression in Property 7.

7.3.3. If X and Y are independent, then the properties of covariance give

$$\text{Var}(X - Y) = \text{Var}(X) + \text{Var}(-Y) = \text{Var}(X) + \text{Var}(Y)$$

It is a common mistake to claim " $\text{Var}(X - Y) = \text{Var}(X) - \text{Var}(Y)$ "; this is a category error since $\text{Var}(X) - \text{Var}(Y)$ could be negative. For general X and Y , we have

$$\text{Var}(X - Y) = \text{Var}(X) + \text{Var}(Y) - 2 \text{Cov}(X, Y)$$

Since covariance depends on the units in which X and Y are measured-if we decide to measure X in centimeters rather than meters, the covariance is multiplied by 100 -it is easier to interpret a unitless version of covariance called correlation.

Definition 7.3.4 (Correlation). The correlation between r.v.s X and Y is

$$\text{Corr}(X, Y) = \frac{\text{Cov}(X, Y)}{\sqrt{\text{Var}(X) \text{Var}(Y)}}$$

(This is undefined in the degenerate cases $\text{Var}(X) = 0$ or $\text{Var}(Y) = 0$.)

Notice that shifting and scaling X and Y has no effect on their correlation. Shifting does not affect $\text{Cov}(X, Y)$, $\text{Var}(X)$, or $\text{Var}(Y)$, so the correlation is unchanged. As for scaling, the fact that we divide by the standard deviations of X and Y ensures that the scale factor cancels out:

$$\text{Corr}(cX, Y) = \frac{\text{Cov}(cX, Y)}{\sqrt{\text{Var}(cX) \text{Var}(Y)}} = \frac{c \text{Cov}(X, Y)}{\sqrt{c^2 \text{Var}(X) \text{Var}(Y)}} = \text{Corr}(X, Y).$$

Correlation is convenient to interpret because it does not depend on the units of measurement and is always between -1 and 1 .

Theorem 7.3.5 (Correlation bounds). For any r.v.s X and Y ,

$$-1 \leq \text{Corr}(X, Y) \leq 1.$$

Proof. Without loss of generality we can assume X and Y have variance 1, since scaling does not change the correlation. Let $\rho = \text{Corr}(X, Y) = \text{Cov}(X, Y)$. Using the fact that variance is nonnegative, along with Property 7 of covariance, we have

$$\begin{aligned} \text{Var}(X + Y) &= \text{Var}(X) + \text{Var}(Y) + 2 \text{Cov}(X, Y) = 2 + 2\rho \geq 0, \\ \text{Var}(X - Y) &= \text{Var}(X) + \text{Var}(Y) - 2 \text{Cov}(X, Y) = 2 - 2\rho \geq 0. \end{aligned}$$

Thus, $-1 \leq \rho \leq 1$.

Here is an example of how to calculate covariance and correlation.

Example 7.3.6 (Exponential max and min). Let X and Y be i.i.d. $\text{Expo}(1)$ r.v.s. Find the correlation between $\max(X, Y)$ and $\min(X, Y)$.

Solution:

Let $M = \max(X, Y)$ and $L = \min(X, Y)$. By the memoryless property and results from Chapter 5, we know that $L \sim \text{Expo}(2)$, $M - L \sim \text{Expo}(1)$, and $M - L$ is independent of L (see Example 5.6.5). Therefore,

$$\begin{aligned} \text{Cov}(M, L) &= \text{Cov}(M - L + L, L) = \text{Cov}(M - L, L) + \text{Cov}(L, L) = 0 + \text{Var}(L) = \frac{1}{4} \\ \text{Var}(M) &= \text{Var}(M - L + L) = \text{Var}(M - L) + \text{Var}(L) = 1 + \frac{1}{4} = \frac{5}{4} \end{aligned}$$

and

$$\text{Corr}(M, L) = \frac{\text{Cov}(M, L)}{\sqrt{\text{Var}(M) \text{Var}(L)}} = \frac{\frac{1}{4}}{\sqrt{\frac{5}{4} \cdot \frac{1}{4}}} = \frac{1}{\sqrt{5}}$$

7.3.7. In the above example, it makes sense that the correlation is positive because M is constrained to be at least as large as L . The following argument would be a blunder: "Either $M = X, L = Y$ or $M = Y, L = X$, so either $\text{Cov}(M, L) = \text{Cov}(X, Y)$ or $\text{Cov}(M, L) = \text{Cov}(Y, X)$. But $\text{Cov}(Y, X) = \text{Cov}(X, Y)$ so we always have $\text{Cov}(M, L) = \text{Cov}(X, Y) = 0$." It is true that either $M = X, L = Y$ will occur or $M = Y, L = X$ will occur, but these are events, not deterministic cases. The argument would fall apart if it were written carefully using conditional probability.

Covariance properties can also be a helpful tool for finding variances, especially when the r.v. of interest is a sum of dependent random variables. The next example uses properties of covariance to derive the variance of the Hypergeometric distribution. If you did Exercise 48 from Chapter 4, you can compare the two derivations.

Example 7.3.8 (Hypergeometric variance). Find $\text{Var}(X)$ for $X \sim \text{HGeom}(w, b, n)$.

Solution:

Interpret X as the number of white balls in a sample of size n from an urn with w white and b black balls. We can represent X as a sum of indicator random variables, $X = I_1 + \cdots + I_n$, where I_j is the indicator of the j th ball in the sample being white. Each I_j has mean $p = w/(w + b)$ and variance $p(1 - p)$, but because the I_j are dependent, we cannot simply add their variances. Instead, we apply properties of covariance:

$$\begin{aligned}\text{Var}(X) &= \text{Var}\left(\sum_{j=1}^n I_j\right) \\ &= \text{Var}(I_1) + \cdots + \text{Var}(I_n) + 2 \sum_{i < j} \text{Cov}(I_i, I_j) \\ &= np(1 - p) + 2 \binom{n}{2} \text{Cov}(I_1, I_2),\end{aligned}$$

since all $\binom{n}{2}$ pairs of indicators have the same covariance by symmetry. Now we just need to find $\text{Cov}(I_1, I_2)$. By the fundamental bridge,

$$\begin{aligned}\text{Cov}(I_1, I_2) &= E(I_1 I_2) - E(I_1) E(I_2) \\ &= P(\text{1 st and 2 nd balls both white}) - P(\text{1 st ball white}) P(\text{2 nd ball white}) \\ &= \frac{w}{w + b} \cdot \frac{w - 1}{w + b - 1} - p^2.\end{aligned}$$

Plugging this into the above formula and simplifying, we eventually obtain

$$\text{Var}(X) = \frac{N - n}{N - 1} np(1 - p),$$

where $N = w + b$. This differs from the Binomial variance of $np(1 - p)$ by a factor of $\frac{N - n}{N - 1}$, which is known as the finite population correction. The discrepancy arises from the fact that in the Binomial story, we sample with replacement, so the same ball can be drawn multiple times; in the Hypergeometric story, we sample without replacement, so each ball appears in the sample at most once.

If we consider N to be the "population size" of the urn, then as N grows very large relative to the sample size n , it becomes extremely unlikely that in sampling with replacement, we would draw the same ball more than once. Thus sampling with replacement and sampling without replacement become equivalent in the limit as $N \rightarrow \infty$ with n fixed, and the finite population correction approaches 1.

The other case where sampling with and without replacement are equivalent is the simple case where we only draw one ball from the urn, and indeed, the finite population correction also equals 1 when $n = 1$.

The last two sections in this chapter introduce the Multinomial and Multivariate Normal distributions. The Multinomial is the most famous discrete multivariate distribution, and the Multivariate Normal is the most famous continuous multivariate distribution.

4.7.4 7.4 Multinomial

The Multinomial distribution is a generalization of the Binomial. Whereas the Binomial distribution counts the successes in a fixed number of trials that can only be categorized as success or failure, the Multinomial distribution keeps track of trials whose outcomes can fall into multiple categories, such as excellent, adequate, poor; or red, yellow, green, blue.

Story 7.4.1 (Multinomial distribution). Each of n objects is independently placed into one of k categories. An object is placed into category j with probability p_j , where the p_j are nonnegative and $\sum_{j=1}^k p_j = 1$. Let X_1 be the number of objects in category 1, X_2 the number of objects in category 2, etc., so that $X_1 + \cdots + X_k = n$. Then $\mathbf{X} = (X_1, \dots, X_k)$ is said to have the Multinomial distribution with parameters n and $\mathbf{p} = (p_1, \dots, p_k)$. We write this as $\mathbf{X} \sim \text{Mult}_k(n, \mathbf{p})$.

We call $\mathbf{X}$ a random vector because it is a vector of random variables. The joint PMF of $\mathbf{X}$ can be derived from the story.

Theorem 7.4.2 (Multinomial joint PMF). If $\mathbf{X} \sim \text{Mult}_k(n, \mathbf{p})$, then the joint PMF of $\mathbf{X}$ is

$$P(X_1 = n_1, \dots, X_k = n_k) = \frac{n!}{n_1!n_2!\dots n_k!} \cdot p_1^{n_1} p_2^{n_2} \dots p_k^{n_k}$$

for $n_1, \dots, n_k$ satisfying $n_1 + \cdots + n_k = n$.

Proof. If $n_1, \dots, n_k$ don't add up to n , then the event $\{X_1 = n_1, \dots, X_k = n_k\}$ is impossible: every object has to go somewhere, and new objects can't appear out of nowhere. If $n_1, \dots, n_k$ do add up to n , then any particular way of putting n_1 objects into category 1, n_2 objects into category 2, etc., has probability $p_1^{n_1} p_2^{n_2} \dots p_k^{n_k}$, and there are

$$\frac{n!}{n_1!n_2!\dots n_k!}$$

ways to do this, as discussed in Example 1.4.18 in the context of rearranging the letters in STATISTICS. So the joint PMF is as claimed.

Since we've specified the joint distribution of $\mathbf{X}$, we have enough information to determine the marginal and conditional distributions, as well as the covariance between any two components of $\mathbf{X}$.

Let's take these one by one, starting with the marginal distribution of X_j , which is the j th component of $\mathbf{X}$. Were we to blindly apply the definition, we would have to sum the joint PMF over all components of $\mathbf{X}$ other than X_j . The prospect of $k - 1$ summations is an unpleasant one, to say the least. Fortunately, we can avoid tedious calculations if we instead use the story of the Multinomial distribution: X_j is the number of objects in category j , where each of the n objects independently belongs to category j with probability p_j . Define success as landing in category j . Then we just have n independent Bernoulli trials, so the marginal distribution of X_j is $\text{Bin}(n, p_j)$.

Theorem 7.4.3 (Multinomial marginals). The marginals of a Multinomial are Binomial. Specifically, if $\mathbf{X} \sim \text{Mult}_k(n, \mathbf{p})$, then $X_j \sim \text{Bin}(n, p_j)$.

More generally, whenever we merge multiple categories together in a Multinomial random vector, we get another Multinomial random vector. For example, suppose we randomly sample

n people in a country with 5 political parties. (If the sampling is done without replacement, the n trials are not independent, but independence is a good approximation as long as the population is large relative to the sample, as we discussed in Theorem 3.9.3 and Example 7.3.8.) Let

$$\mathbf{X} = (X_1, \dots, X_5) \sim \text{Mult}_5(n, (p_1, \dots, p_5))$$

represent the political party affiliations of the sample, i.e., X_j is the number of people in the sample who support party j .

Suppose that parties 1 and 2 are the dominant parties, while parties 3 through 5 are minor third parties. If we decide that instead of keeping track of all 5 parties, we only want to count the number of people in party 1, party 2, or "other", then we can define a new random vector that lumps all the third parties into one category:

$$\mathbf{Y} = (X_1, X_2, X_3 + X_4 + X_5).$$

By the story of the Multinomial,

$$\mathbf{Y} \sim \text{Mult}_3(n, (p_1, p_2, p_3 + p_4 + p_5)).$$

Of course, this idea applies to merging categories in any Multinomial, not just in the context of political parties.

Theorem 7.4.4 (Multinomial lumping). If $\mathbf{X} \sim \text{Mult}_k(n, \mathbf{p})$, then for any distinct i and j , $X_i + X_j \sim \text{Bin}(n, p_i + p_j)$. The random vector of counts obtained from merging categories i and j is still Multinomial. For example, merging categories 1 and 2 gives

$$(X_1 + X_2, X_3, \dots, X_k) \sim \text{Mult}_{k-1}(n, (p_1 + p_2, p_3, \dots, p_k)).$$

Now for the conditional distributions. Suppose we get to observe X_1 , the number of objects in category 1, and we wish to update our distribution for the other categories $(X_2, \dots, X_k)$. One way to do this is with the definition of conditional PMF:

$$P(X_2 = n_2, \dots, X_k = n_k \mid X_1 = n_1) = \frac{P(X_1 = n_1, X_2 = n_2, \dots, X_k = n_k)}{P(X_1 = n_1)}.$$

The numerator is the joint PMF of the Multinomial, and the denominator is the marginal PMF of X_1 , both of which we have already derived. However, we prefer to use the Multinomial story to deduce the conditional distribution of $(X_2, \dots, X_k)$ without algebra. Given that there are n_1 objects in category 1, the remaining $n - n_1$ objects fall into categories 2 through k , independently of one another. By Bayes' rule, the conditional probability of falling into category j is

$$P(\text{in category } j \mid \text{not in category 1}) = \frac{P(\text{in category } j)}{P(\text{not in category 1})} = \frac{p_j}{p_2 + \dots + p_k},$$

for $j = 2, \dots, k$. This makes intuitive sense: the updated probabilities are proportional to the original probabilities $(p_2, \dots, p_k)$, but these must be renormalized to yield a valid probability vector. Putting it all together, we have the following result.

Theorem 7.4.5 (Multinomial conditioning). If $\mathbf{X} \sim \text{Mult}_k(n, \mathbf{p})$, then

$$(X_2, \dots, X_k) \mid X_1 = n_1 \sim \text{Mult}_{k-1}(n - n_1, (p'_2, \dots, p'_k)),$$

where $p'_j = p_j / (p_2 + \dots + p_k)$.

Finally, we know that components within a Multinomial random vector are dependent since

they are constrained by $X_1 + \cdots + X_k = n$. To find the covariance between X_i and X_j , we can use the marginal and lumping properties we have just discussed.

Theorem 7.4.6 (Covariance in a Multinomial). Let $(X_1, \dots, X_k) \sim \text{Mult}_k(n, \mathbf{p})$, where $\mathbf{p} = (p_1, \dots, p_k)$. For $i \neq j$, $\text{Cov}(X_i, X_j) = -np_i p_j$.

Proof. For concreteness, let $i = 1$ and $j = 2$. Using the lumping property and the marginal distributions of a Multinomial, we know $X_1 + X_2 \sim \text{Bin}(n, p_1 + p_2)$, $X_1 \sim \text{Bin}(n, p_1)$, and $X_2 \sim \text{Bin}(n, p_2)$. Therefore

$$\text{Var}(X_1 + X_2) = \text{Var}(X_1) + \text{Var}(X_2) + 2\text{Cov}(X_1, X_2)$$

becomes

$$n(p_1 + p_2)(1 - (p_1 + p_2)) = np_1(1 - p_1) + np_2(1 - p_2) + 2\text{Cov}(X_1, X_2).$$

Solving for $\text{Cov}(X_1, X_2)$ gives $\text{Cov}(X_1, X_2) = -np_1 p_2$. By the same logic, for $i \neq j$, we have $\text{Cov}(X_i, X_j) = -np_i p_j$.

The components are negatively correlated, as we would expect: if we know there are a lot of objects in category i , then there aren't as many objects left over that could possibly be in category j . Exercise 65 asks for a different proof of this result, using indicators.

7.4.7 (Independent trials but dependent components). The k components of a Multinomial are dependent, but the n objects in the story of the Multinomial are independently categorized. In the extreme case $k = 2$, a $\text{Mult}_k(n, \mathbf{p})$ random vector looks like $(X, n - X)$ with $X \sim \text{Bin}(n, p_1)$, which we can think of as (number of successes, number of failures), where "success" is defined as getting category 1. The number of successes is perfectly negatively correlated with the number of failures, even though the trials are independent.

We conclude this section with a heroic example that ties together many of the most important concepts for the Multinomial.

Example 7.4.8 (Statwoman). The superhero Statwoman uses probability and statistics to fight crime. She has battled with countless foes, sometimes even fighting several at the same time. For simplicity though, assume that each of her battles is with exactly one of the following adversaries: the Confounder, the Extrapolator, and the Overfitter.

Suppose that Statwoman will have n battles next year (with n a positive integer), and that each battle is with the Confounder with probability p_1 , the Extrapolator with probability p_2 , and the Overfitter with probability p_3 , independently. Here p_1, p_2, p_3 are nonnegative and sum to 1. Let X_1, X_2, X_3 be the numbers of battles Statwoman will have with the Confounder, the Extrapolator, and the Overfitter next year, respectively.

(a) Find the joint distribution of X_1, X_2, X_3 .

(b) Find the correlation between X_1 and X_2 .

(c) Suppose for this part only that it turns out that the Extrapolator and the Overfitter have been devising evil plots together, so it is of interest to study their combined number of skirmishes with Statwoman. Let $X_{23} = X_2 + X_3$. Find the joint distribution of X_1, X_{23} .

(d) Suppose for this part only that the parameters p_1, p_2, p_3 are unknown, $n = 360$, and it is observed that exactly 36 of Statwoman's battles are with the Overfitter. A natural way to estimate p_3 would be to use $36/360 = 0.1$. The maximum likelihood estimate (MLE) of p_3 is the value of p_3 that makes the observed data, $X_3 = 36$, as likely as possible. That is, the MLE is the value of p_3 that maximizes $P(X_3 = 36)$. Show that the MLE is the natural estimate, 0.1.

(e) Assume for this part only that the Overfitter has been captured, through heroic efforts! So assume that all of Statwoman's battles next year will be with one of her other two adversaries. Find the joint PMF of X_1, X_2 , given that $X_3 = 0$.

(f) Now suppose that, instead of the number of battles being a constant n , the number of battles is $N \sim \text{Pois}(\lambda)$. Find the joint distribution of X_1, X_2, X_3 .

Solution:

- (a) By the story of the Multinomial, $(X_1, X_2, X_3) \sim \text{Mult}_3(n, (p_1, p_2, p_3))$.
 (b) We have $X_1 \sim \text{Bin}(n, p_1)$, $X_2 \sim \text{Bin}(n, p_2)$, and $\text{Cov}(X_1, X_2) = -np_1p_2$, so
 $\text{Corr}(X_1, X_2) = \frac{\text{Cov}(X_1, X_2)}{\text{SD}(X_1)\text{SD}(X_2)} = \frac{-np_1p_2}{\sqrt{np_1(1-p_1)np_2(1-p_2)}} = -\sqrt{\frac{p_1p_2}{(1-p_1)(1-p_2)}}$.
 (c) By the lumping property of the Multinomial,

$$(X_1, X_{23}) \sim \text{Mult}_2(n, (p_1, p_2 + p_3)).$$

- (d) In general for $X \sim \text{Bin}(n, p)$, if $X = x$ is observed then the MLE of p is the value of p that maximizes the function

$$L(p) = \binom{n}{x} p^x (1-p)^{n-x}$$

(The function L is called the likelihood function in statistics. It is the probability of the data, regarded as a function of the parameter with the data treated as fixed.)

As if often the case when dealing with a product of positive numbers, it is helpful to take the log. It is equivalent to find the value $\hat{p}$ that maximizes

$$\log L(p) = \log \binom{n}{x} + x \log p + (n-x) \log(1-p)$$

since log is a continuous, strictly increasing function. Setting the derivative of $\log L(p)$ (with respect to p , holding x constant) equal to 0, we have

$$\frac{x}{\hat{p}} - \frac{n-x}{1-\hat{p}} = 0,$$

which rearranges to $\hat{p} = x/n$. We have found a maximum since the second derivative of $\log L(p)$ is

$$-\frac{x}{p^2} - \frac{n-x}{(1-p)^2} < 0.$$

So the MLE of p_3 is $\hat{p}_3 = 36/360 = 0.1$.

- (e) By the result on Multinomial conditioning,

$$(X_1, X_2) \Big| (X_3 = 0) \sim \text{Mult}_2 \left(n, \left(\frac{p_1}{p_1 + p_2}, \frac{p_2}{p_1 + p_2} \right) \right).$$

- (f) Reasoning as in the chicken-egg story,

$$P(X_1 = x_1, X_2 = x_2, X_3 = x_3) = \sum_{n=0}^{\infty} P(X_1 = x_1, X_2 = x_2, X_3 = x_3 \mid N = n) P(N = n),$$

where all terms of the sum are 0 except the term with $n = x_1 + x_2 + x_3$. For this value of n ,

$$\begin{aligned} P(X_1 = x_1, X_2 = x_2, X_3 = x_3) &= P(X_1 = x_1, X_2 = x_2, X_3 = x_3 \mid N = n) P(N = n) \\ &= \frac{n!}{x_1! x_2! x_3!} p_1^{x_1} p_2^{x_2} p_3^{x_3} \cdot e^{-\lambda} \frac{\lambda^n}{n!} \\ &= \frac{e^{-\lambda p_1} (\lambda p_1)^{x_1}}{x_1!} \cdot \frac{e^{-\lambda p_2} (\lambda p_2)^{x_2}}{x_2!} \cdot \frac{e^{-\lambda p_3} (\lambda p_3)^{x_3}}{x_3!}, \end{aligned}$$

for all nonnegative integers x_1, x_2, x_3 . Therefore, X_1, X_2, X_3 are independent, with $X_j \sim \text{Pois}(\lambda p_j)$. This result is a Multinomial extension of the chicken-egg story. $\square$

4.7.5 7.5 Multivariate Normal

The Multivariate Normal is a continuous multivariate distribution that generalizes the Normal distribution into higher dimensions. We will not work with the rather unwieldy joint PDF of the Multivariate Normal. Instead we define the Multivariate Normal by its relationship to the ordinary Normal.

Definition 7.5.1 (Multivariate Normal distribution). A k -dimensional random vector $\mathbf{X} = (X_1, \dots, X_k)$ is said to have a Multivariate Normal (MVN) distribution if every linear combination of the X_j has a Normal distribution. That is, we require

$$t_1 X_1 + \dots + t_k X_k$$

to have a Normal distribution for any constants $t_1, \dots, t_k$. If $t_1 X_1 + \dots + t_k X_k$ is a constant (such as when all $t_i = 0$), we consider it to have a Normal distribution, albeit a degenerate Normal with variance 0. An important special case is $k = 2$; this distribution is called the Bivariate Normal (BVN).

If $(X_1, \dots, X_k)$ is MVN, then the marginal distribution of X_1 is Normal, since we can take t_1 to be 1 and all other t_j to be 0. Similarly, the marginal distribution of each X_j is Normal. However, the converse is false: it is possible to have Normally distributed r.v.s $X_1, \dots, X_k$ such that $(X_1, \dots, X_k)$ is not Multivariate Normal.

Example 7.5.2 (Non-example of MVN). Here is an example of two r.v.s whose marginal distributions are Normal but whose joint distribution is not Bivariate Normal. Let $X \sim \mathcal{N}(0, 1)$, and let

$$S = \begin{cases} 1 & \text{with probability } 1/2 \\ -1 & \text{with probability } 1/2 \end{cases}$$

be a random sign independent of X . Then $Y = SX$ is a standard Normal r.v., due to the symmetry of the Normal distribution (see Exercise 30 from Chapter 5). However, (X, Y) is not Bivariate Normal because $P(X + Y = 0) = P(S = -1) = 1/2$, which implies that $X + Y$ can't be Normal (or, for that matter, have any continuous distribution). Since $X + Y$ is a linear combination of X and Y that is not Normally distributed, (X, Y) is not Bivariate Normal. $\square$

Example 7.5.3 (Actual MVN). For $Z, W \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$, (Z, W) is Bivariate Normal because the sum of independent Normals is Normal. Also, $(Z + 2W, 3Z + 5W)$ is Bivariate Normal, since an arbitrary linear combination

$$t_1(Z + 2W) + t_2(3Z + 5W)$$

can also be written as a linear combination of Z and W ,

$$(t_1 + 3t_2)Z + (2t_1 + 5t_2)W,$$

which is Normal. $\square$

The above example showed that if we start with a Multivariate Normal and take linear combinations of the components, we form a new Multivariate Normal. The next two theorems state that we can also produce new MVNs from old MVNs with the operations of subsetting and concatenation.

Theorem 7.5.4. If (X_1, X_2, X_3) is Multivariate Normal, then so is the subvector (X_1, X_2) .

Proof. Any linear combination $t_1 X_1 + t_2 X_2$ can be thought of as a linear combination of X_1, X_2, X_3 where the coefficient of X_3 is 0. So $t_1 X_1 + t_2 X_2$ is Normal for all t_1, t_2 , which shows that (X_1, X_2) is MVN. $\square$

Theorem 7.5.5. If $\mathbf{X} = (X_1, \dots, X_n)$ and $\mathbf{Y} = (Y_1, \dots, Y_m)$ are Multivariate Normal random vectors with $\mathbf{X}$ independent of $\mathbf{Y}$, then the concatenated random vector $\mathbf{W} = (X_1, \dots, X_n, Y_1, \dots, Y_m)$ is Multivariate Normal.

Proof. Any linear combination $s_1X_1 + \dots + s_nX_n + t_1Y_1 + \dots + t_mY_m$ is Normal since $s_1X_1 + \dots + s_nX_n$ and $t_1Y_1 + \dots + t_mY_m$ are Normal (by definition of MVN) and are independent, so their sum is Normal (as shown in Chapter 6 using MGFs). $\square$

A Multivariate Normal distribution is fully specified by knowing the mean of each component, the variance of each component, and the covariance or correlation between any two components. Another way to say this is that the parameters of an MVN random vector $(X_1, \dots, X_k)$ are as follows:

- the mean vector $(\mu_1, \dots, \mu_k)$, where $E(X_j) = \mu_j$;
- the covariance matrix, which is the $k \times k$ matrix of covariances between components, arranged so that the row i , column j entry is $\text{Cov}(X_i, X_j)$.

For example, in order to fully specify a Bivariate Normal distribution for (X, Y) , we need to know five parameters:

- the means $E(X), E(Y)$;
- the variances $\text{Var}(X), \text{Var}(Y)$;
- the correlation $\text{Corr}(X, Y)$.

We will show in Example 8.1.10 that the joint PDF of a Bivariate Normal (X, Y) with $\mathcal{N}(0, 1)$ marginal distributions and correlation $\rho \in (-1, 1)$ is

$$f_{X,Y}(x, y) = \frac{1}{2\pi\tau} \exp\left(-\frac{1}{2\tau^2}(x^2 + y^2 - 2\rho xy)\right)$$

with $\tau = \sqrt{1 - \rho^2}$. Figure 7.10 plots the joint PDFs for two different Bivariate Normal distributions with $\mathcal{N}(0, 1)$ marginals, along with the corresponding contour plots. On the left, X and Y are uncorrelated, so the level curves of the joint PDF are circles. On the right, X and Y have a correlation of 0.75, so the level curves are ellipsoidal, reflecting the fact that Y tends to be large when X is large.

Just as the distribution of an r.v. is determined by its CDF, PMF/PDF, or MGF, the joint distribution of a random vector is determined by its joint CDF, joint PMF/PDF, or joint MGF, which we now define.

Definition 7.5.6 (Joint MGF). The joint moment generating function (joint MGF) of a random vector $\mathbf{X} = (X_1, \dots, X_k)$ is the function M defined by

$$M(\mathbf{t}) = E\left(e^{\mathbf{t}'\mathbf{X}}\right) = E\left(e^{t_1X_1 + \dots + t_kX_k}\right)$$

for $\mathbf{t} = (t_1, \dots, t_k) \in \mathbb{R}^k$. We require this expectation to be finite in a box containing the origin in $\mathbb{R}^k$; otherwise we say the joint MGF does not exist.

For a Multivariate Normal random vector, the joint MGF is particularly nice because the term in the exponent, $t_1X_1 + \dots + t_kX_k$, is a Normal r.v. by definition. This means we can use what we know about the univariate Normal MGF to find the Multivariate Normal joint MGF! Recall that for any Normal r.v. W ,

$$E(e^W) = e^{E(W) + \frac{1}{2}\text{Var}(W)}.$$

Therefore the joint MGF of a Multivariate Normal $(X_1, \dots, X_k)$ is

$$E(e^{t_1 X_1 + \dots + t_k X_k}) = \exp\left(t_1 E(X_1) + \dots + t_k E(X_k) + \frac{1}{2} \text{Var}(t_1 X_1 + \dots + t_k X_k)\right)$$

The variance term can be expanded using properties of covariance.

We know that in general, independence is a stronger condition than zero correlation; r.v.s can be uncorrelated but not independent. A special property of the Multivariate Normal distribution is that for r.v.s whose joint distribution is MVN, independence and zero correlation are equivalent conditions.

Theorem 7.5.7. Within an MVN random vector, uncorrelated implies independent. That is, if $\mathbf{X} \sim \text{MVN}$ can be written as $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2)$, where $\mathbf{X}_1$ and $\mathbf{X}_2$ are




FIGURE 7.10

Joint PDFs of two Bivariate Normal distributions, and the corresponding contour plots. On the left, X and Y are marginally $\mathcal{N}(0,1)$ and have zero correlation. On the right, X and Y are marginally $\mathcal{N}(0,1)$ and have correlation 0.75.

subvectors, and every component of $\mathbf{X}_1$ is uncorrelated with every component of $\mathbf{X}_2$, then $\mathbf{X}_1$ and $\mathbf{X}_2$ are independent. ¹

In particular, if (X, Y) is Bivariate Normal and $\text{Corr}(X, Y) = 0$, then X and Y are independent.

Proof. We will prove this theorem in the case of a Bivariate Normal; the proof in higher dimensions is analogous. Let (X, Y) be Bivariate Normal with $E(X) = \mu_1$, $E(Y) = \mu_2$, $\text{Var}(X) = \sigma_1^2$, $\text{Var}(Y) = \sigma_2^2$, and $\text{Corr}(X, Y) = \rho$. The joint MGF is

$$\begin{aligned} M_{X,Y}(s, t) &= E(e^{sX+tY}) = \exp\left(s\mu_1 + t\mu_2 + \frac{1}{2} \text{Var}(sX + tY)\right) \\ &= \exp\left(s\mu_1 + t\mu_2 + \frac{1}{2} (s^2\sigma_1^2 + t^2\sigma_2^2 + 2st\sigma_1\sigma_2\rho)\right) \end{aligned}$$

If $\rho = 0$, the joint MGF reduces to

$$M_{X,Y}(s, t) = \exp\left(s\mu_1 + t\mu_2 + \frac{1}{2} (s^2\sigma_1^2 + t^2\sigma_2^2)\right).$$

But this is also the joint MGF of (Z, W) where $Z \sim \mathcal{N}(\mu_1, \sigma_1^2)$ and $W \sim \mathcal{N}(\mu_2, \sigma_2^2)$ and Z is independent of W . Since the joint MGF determines the joint distribution, it must be that (X, Y) has the same joint distribution as (Z, W) . Therefore X and Y are independent. $\square$

This theorem does not apply to Example 7.5.2. In that example, as you can verify, X and Y are uncorrelated and not independent, but this does not contradict the theorem because (X, Y) is not BVN. The next two examples show situations where the theorem does apply.

Example 7.5.8 (Independence of sum and difference). Let $X, Y \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$. Find the joint distribution of $(X + Y, X - Y)$.

Solution:

Since $(X + Y, X - Y)$ is Bivariate Normal and

$$\text{Cov}(X + Y, X - Y) = \text{Var}(X) - \text{Cov}(X, Y) + \text{Cov}(Y, X) - \text{Var}(Y) = 0,$$

$X + Y$ is independent of $X - Y$. Furthermore, they are i.i.d. $\mathcal{N}(0, 2)$. By the same method, we have that if $X \sim \mathcal{N}(\mu_1, \sigma^2)$ and $Y \sim \mathcal{N}(\mu_2, \sigma^2)$ are independent (with the same variance), then $X + Y$ is independent of $X - Y$.

It can be shown that the independence of the sum and difference is a unique characteristic of the Normal! That is, if X and Y are i.i.d. and $X + Y$ is independent of $X - Y$, then X and Y must have Normal distributions.

In Exercise 72, you'll extend this example to the case where X and Y are Bivariate Normal with general correlation ρ .

Example 7.5.9 (Independence of sample mean and sample variance). Let $X_1, \dots, X_n$ be i.i.d. $\mathcal{N}(\mu, \sigma^2)$, with $n \geq 2$. Define

$$\begin{aligned}\bar{X}_n &= \frac{1}{n} (X_1 + \dots + X_n), \\ S_n^2 &= \frac{1}{n-1} \sum_{j=1}^n (X_j - \bar{X}_n)^2.\end{aligned}$$

As shown in Chapter 6, the sample mean $\bar{X}_n$ has expectation μ (the true mean) and the sample variance S_n^2 has expectation σ^2 (the true variance). Show that $\bar{X}_n$ and S_n^2 are independent by applying MVN ideas to $(\bar{X}_n, X_1 - \bar{X}_n, \dots, X_n - \bar{X}_n)$.

Solution:

The vector $(\bar{X}_n, X_1 - \bar{X}_n, \dots, X_n - \bar{X}_n)$ is MVN since any linear combination of its components can be written as a linear combination of $X_1, \dots, X_n$. Furthermore, $E(X_j - \bar{X}_n) = 0$ by linearity. Now we compute the covariance of $\bar{X}_n$ with $X_j - \bar{X}_n$:

$$\text{Cov}(\bar{X}_n, X_j - \bar{X}_n) = \text{Cov}(\bar{X}_n, X_j) - \text{Cov}(\bar{X}_n, \bar{X}_n)$$

For $\text{Cov}(\bar{X}_n, X_j)$, we can expand out $\bar{X}_n$, and most of the terms cancel due to independence: $\text{Cov}(\bar{X}_n, X_j) = \text{Cov}(\frac{1}{n}X_1 + \dots + \frac{1}{n}X_n, X_j) = \text{Cov}(\frac{1}{n}X_j, X_j) = \frac{1}{n} \text{Var}(X_j) = \frac{\sigma^2}{n}$. For $\text{Cov}(\bar{X}_n, \bar{X}_n)$, we use properties of variance:

$$\text{Cov}(\bar{X}_n, \bar{X}_n) = \text{Var}(\bar{X}_n) = \frac{1}{n^2} (\text{Var}(X_1) + \dots + \text{Var}(X_n)) = \frac{\sigma^2}{n}.$$

⁰¹ Independence of random vectors is defined analogously to independence of random variables. In particular, if $\mathbf{X}_i$ has joint PDF $f_{\mathbf{X}_i}$, then it says that $f_{\mathbf{X}}(\mathbf{x}_1, \mathbf{x}_2) = f_{\mathbf{X}_1}(\mathbf{x}_1) f_{\mathbf{X}_2}(\mathbf{x}_2)$ for all $\mathbf{x}_1, \mathbf{x}_2$.

Therefore $\text{Cov}(\bar{X}_n, X_j - \bar{X}_n) = 0$, which means $\bar{X}_n$ is uncorrelated with every component of $(X_1 - \bar{X}_n, \dots, X_n - \bar{X}_n)$. Since uncorrelated implies independent within an MVN vector, we have that $\bar{X}_n$ is independent of the vector $(X_1 - \bar{X}_n, \dots, X_n - \bar{X}_n)$. But S_n^2 is a function of $(X_1 - \bar{X}_n, \dots, X_n - \bar{X}_n)$, so $\bar{X}_n$ is also independent of S_n^2 .

It can be shown that the independence of the sample mean and variance is another unique characteristic of the Normal! If the X_j followed any other distribution, then $\bar{X}_n$ and S_n^2 would be dependent.

Example 7.5.10 (Bivariate Normal generation). Suppose that we have access to i.i.d. random variables $X, Y \sim \mathcal{N}(0, 1)$, but want to generate a Bivariate Normal random vector (Z, W) with $\text{Corr}(Z, W) = \rho$ and Z, W marginally $\mathcal{N}(0, 1)$, for the purpose of running a simulation. How can we construct Z and W from linear combinations of X and Y ?

Solution:

By definition of Multivariate Normal, any (Z, W) of the form

$$Z = aX + bY$$

$$W = cX + dY$$

will be Bivariate Normal. So let's try to find suitable a, b, c, d . The means are already 0. Setting the variances equal to 1 gives

$$a^2 + b^2 = 1, c^2 + d^2 = 1$$

Setting the covariance of Z and W equal to ρ gives

$$ac + bd = \rho$$

There are more unknowns than equations here, and we just need one solution. To simplify, let's look for a solution with $b = 0$. Then $a^2 = 1$, so let's take $a = 1$. Now $ac + bd = \rho$ reduces to $c = \rho$, and then we can use $c^2 + d^2 = 1$ to find a suitable d . Putting everything together, we can generate (Z, W) as

$$Z = X$$

$$W = \rho X + \sqrt{1 - \rho^2} Y.$$

Note that in the extreme case $\rho = 1$ (known as perfect positive correlation) this says to let $W = Z \sim \mathcal{N}(0, 1)$, in the extreme case $\rho = -1$ (known as perfect negative correlation) it says to let $W = -Z$ with $Z \sim \mathcal{N}(0, 1)$, and in the simple case $\rho = 0$ it says to just let $(Z, W) = (X, Y)$.

4.7.6 7.6 Recap

Joint distributions allow us to describe the behavior of multiple random variables that arise from the same experiment. Important functions associated with a joint distribution are the joint CDF, joint PMF/PDF, marginal PMF/PDF, and conditional PMF/PDF. The table on the next page summarizes these definitions for two discrete r.v.s and two continuous r.v.s. Joint distributions can also be a hybrid of discrete and continuous, in which case we mix and match PMFs and PDFs.

	Two discrete r.v.s
Joint CDF	$F_{X,Y}(x,y) = P(X \leq x, Y \leq y)$
Joint PMF/PDF	$P(X = x, Y = y)$ - Joint PMF is nonnegative. - Joint PMF sums to 1. - $P((X,Y) \in A) = \sum_{(x,y) \in A} \sum_A P(X = x, Y = y)$.
Marginal PMF/PDF	$P(X = x) = \sum_y P(X = x, Y = y)$ $= \sum_y P(X = x Y = y)P(Y = y)$
Conditional PMF/PDF	$P(Y = y X = x) = \frac{P(X = x, Y = y)}{P(X = x)}$ $= \frac{P(X = x Y = y)P(Y = y)}{P(X = x)}$
Independence	$P(X \leq x, Y \leq y) = P(X \leq x)P(Y \leq y)$ $P(X = x, Y = y) = P(X = x)P(Y = y)$ for all x and y .
LOTUS	$E(g(X, Y)) = \sum_y \sum_x g(x, y)P(X = x, Y = y)$

Covariance is a single-number summary of the tendency of two r.v.s to move in the same direction. If two r.v.s are independent, then they are uncorrelated, but the converse does not hold. Correlation is a unitless, standardized version of covariance that is always between -1 and 1 .

Two important named multivariate distributions are the Multinomial and Multivariate Normal. The Multinomial distribution is a generalization of the Binomial; a $\text{Mult}_k(n, \mathbf{p})$ random vector counts the number of objects, out of n , that fall into each of k categories, where $\mathbf{p}$ is the vector of probabilities for the k categories.

The Multivariate Normal distribution is a generalization of the Normal; a random vector is defined to be MVN if any linear combination of its components has a Normal distribution. A key property of the MVN distribution is that within an MVN random vector, uncorrelated implies independent.




FIGURE 7.11

Fundamental objects of probability for multivariate distributions. A joint distribution is determined by a joint CDF , joint PMF/PDF, or joint MGF . A random vector (X, Y) gives

rise to many useful joint, marginal, and conditional quantities. Using 2D LOTUS, we can find the expected value of a function of X and Y . Summing the joint PMF or integrating the joint PDF over all y gives the marginal distribution of X , bringing us back to the case of a one-dimensional distribution.

Figure 7.11 extends our diagram of the fundamental objects of probability to the multivariate setting (taken as bivariate to simplify notation). A joint distribution can be used to generate random vectors (X, Y) . Various joint, marginal, and conditional quantities can then be studied. Summing the joint PMF or integrating the joint PDF over all y gives the marginal distribution of X , bringing us back to the one-dimensional realm.

4.7.7 7.7r

Multinomial

The functions for the Multinomial distribution are `dmultinom` (which is the joint PMF of the Multinomial distribution) and `rmultinom` (which generates realizations of Multinomial random vectors). The joint CDF of the Multinomial is a pain to work with, so it is not built into R.

To use `dmultinom`, we have to input the value at which to evaluate the joint PMF, as well as the parameters of the distribution. For example,

```
x <- c(2,0,3)
n <- 5
p <- c(1/3,1/3,1/3)
dmultinom(x,n,p)
```

returns the probability $P(X_1 = 2, X_2 = 0, X_3 = 3)$, where

$$\mathbf{X} = (X_1, X_2, X_3) \sim \text{Mult}_3(5, (1/3, 1/3, 1/3)).$$

Of course, `n` has to equal `sum(x)`; if we attempted to do `dmultinom(x, 7, p)`, R would report an error.

For `rmultinom`, the first input is the number of Multinomial random vectors to generate, and the other inputs are the same. When we typed `rmultinom(10, n, p)` with n and p as above, R gave us the following matrix:

0	2	1	3	2	3	1	2	3	4
2	2	2	2	3	0	1	2	0	0
3	1	2	0	0	2	3	1	2	1

Each column of the matrix corresponds to a draw from the $\text{Mult}_3(5, (1/3, 1/3, 1/3))$ distribution. In particular, the sum of each column is 5.

Multivariate Normal

Functions for the Multivariate Normal distribution are located in the package `mvtnorm`. Online resources can teach you how to install packages in R for your system, but for many systems an easy way is to use the `install.packages` command, e.g., by typing `install.packages("mvtnorm")` to install the `mvtnorm` package. After installing it, load the package with `library(mvtnorm)`. Then `dmvnorm` can be used for calculating the joint PDF, and `rmvnorm` can be used for generating random vectors. For example, suppose that we want to generate 1000 independent Bivariate Normal pairs (Z, W) , with correlation $\rho = 0.7$ and $\mathcal{N}(0, 1)$ marginals. To do this, we can enter the following:


```
meanvector <- c(0,0)
rho <- 0.7
covmatrix <- matrix(c(1,rho,rho,1), nrow = 2, ncol = 2)
r <- rmvnorm(n = 10^3, mean = meanvector, sigma = covmatrix)
```

The covariance matrix here is

$$\begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$

because

- $\text{Cov}(Z, Z) = \text{Var}(Z) = 1$ (this is the upper left entry),
- $\text{Cov}(W, W) = \text{Var}(W) = 1$ (this is the lower right entry),
- $\text{Cov}(Z, W) = \text{Corr}(Z, W) \text{SD}(Z) \text{SD}(W) = \rho$ (this is the other two entries).

Now r is a 1000×2 matrix, with each row a BVN random vector. To see these as points in the plane, we can use `plot(r)` to make a scatter plot, from which the strong positive correlation should be clear. To estimate the covariance of Z and W , we can use `cov(r)`, which the true covariance matrix.

Example 7.5.10 gives another approach to the BVN generation problem:

```
rho <- 0.7
tau <- sqrt(1-rho^2)
x <- rnorm(10^3)
y <- rnorm(10^3)
z <- x
w <- rho*x + tau*y
```

This gives the Z -coordinates in a vector z and the W -coordinates in a vector w . If we want to put them into one 1000×2 matrix as we had above, we can type `cbind(z, w)` to bind the vectors together as columns.

Cauchy

We can work with the Cauchy distribution introduced in Example 7.1.25 using the three functions `dcauchy`, `pcauchy`, and `rcauchy`. Only one input is needed; for example, `dcauchy(0)` is the Cauchy PDF evaluated at 0.

For an amusing demonstration of the very heavy tails of the Cauchy distribution, try creating a histogram of 1000 simulated values of the Cauchy distribution:

```
hist(rcauchy(1000))
```

Due to extreme values in the tails of the distribution, this histogram looks nothing like the PDF of the distribution from which it was generated.

4.7.8 7.8 Exercises

Exercises marked with (S) have detailed solutions at <http://stat110.net>.

Joint, Marginal, and Conditional Distributions

1. Alice and Bob arrange to meet for lunch on a certain day at noon. However, neither is known for punctuality. They both arrive independently at uniformly distributed times between noon and 1 pm on that day. Each is willing to wait up to 15 minutes for the other to show up. What is the probability they will meet for lunch that day?
2. Alice, Bob, and Carl arrange to meet for lunch on a certain day. They arrive independently at uniformly distributed times between 1 pm and 1:30 pm on that day.
 - (a) What is the probability that Carl arrives first?

For the rest of this problem, assume that Carl arrives first at 1:10 pm, and condition on this fact.

- (b) What is the probability that Carl will be waiting alone for more than 10 minutes?
- (c) What is the probability that Carl will have to wait more than 10 minutes until his party is complete?
- (d) What is the probability that the person who arrives second will have to wait more than 5 minutes for the third person to show up?
3. One of two doctors, Dr. Hibbert and Dr. Nick, is called upon to perform a series of n surgeries. Let H be the indicator r.v. for Dr. Hibbert performing the surgeries, and suppose that $E(H) = p$. Given that Dr. Hibbert is performing the surgeries, each surgery is successful with probability a , independently. Given that Dr. Nick is performing the surgeries, each surgery is successful with probability b , independently. Let X be the number of successful surgeries.
 - (a) Find the joint PMF of H and X .
 - (b) Find the marginal PMF of X .
 - (c) Find the conditional PMF of H given $X = k$.
4. A fair coin is flipped twice. Let X be the number of Heads in the two tosses, and Y be the indicator r.v. for the tosses landing the same way.
 - (a) Find the joint PMF of X and Y .
 - (b) Find the marginal PMFs of X and Y .
 - (c) Are X and Y independent?
 - (d) Find the conditional PMFs of Y given $X = x$ and of X given $Y = y$.
5. A fair die is rolled, and then a coin with probability p of Heads is flipped as many times as the die roll says, e.g., if the result of the die roll is a 3, then the coin is flipped 3 times. Let X be the result of the die roll and Y be the number of times the coin lands Heads.
 - (a) Find the joint PMF of X and Y . Are they independent?
 - (b) Find the marginal PMFs of X and Y .
 - (c) Find the conditional PMFs of Y given $X = x$ and of X given $Y = y$.
6. A committee of size k is chosen from a group of n women and m men. All possible committees of size k are equally likely. Let X and Y be the numbers of women and men on the committee, respectively.
 - (a) Find the joint PMF of X and Y . Be sure to specify the support.
 - (b) Find the marginal PMF of X in two different ways: by doing a computation using the joint PMF, and using a story.
 - (c) Find the conditional PMF of Y given that $X = x$.
7. A stick of length L (a positive constant) is broken at a uniformly random point X . Given that $X = x$, another breakpoint Y is chosen uniformly on the interval $[0, x]$.
 - (a) Find the joint PDF of X and Y . Be sure to specify the support.
 - (b) We already know that the marginal distribution of X is $\text{Unif}(0, L)$. Check that marginalizing out Y from the joint PDF agrees that this is the marginal distribution of X .
 - (c) We already know that the conditional distribution of Y given $X = x$ is $\text{Unif}(0, x)$. Check that using the definition of conditional PDFs (in terms of joint and marginal PDFs) agrees that

this is the conditional distribution of Y given $X = x$.

(d) Find the marginal PDF of Y .

(e) Find the conditional PDF of X given $Y = y$.

8. (a) Five cards are randomly chosen from a standard deck, one at a time with replacement. Let X, Y, Z be the numbers of chosen queens, kings, and other cards. Find the joint PMF of X, Y, Z .

(b) Find the joint PMF of X and Y .

Hint: In summing the joint PMF of X, Y, Z over the possible values of Z , note that most terms are 0 because of the constraint that the number of chosen cards is five.

(c) Now assume instead that the sampling is without replacement (all 5-card hands are equally likely). Find the joint PMF of X, Y, Z .

Hint: Use the naive definition of probability.

9. Let X and Y be i.i.d. $\text{Geom}(p)$, and $N = X + Y$.

(a) Find the joint PMF of X, Y, N .

(b) Find the joint PMF of X and N .

(c) Find the conditional PMF of X given $N = n$, and give a simple description in words of what the result says.

10. Let X and Y be i.i.d. $\text{Expo}(\lambda)$, and $T = X + Y$.

(a) Find the conditional CDF of T given $X = x$. Be sure to specify where it is zero.

(b) Find the conditional PDF $f_{T|X}(t | x)$, and verify that it is a valid PDF.

(c) Find the conditional PDF $f_{X|T}(x | t)$, and verify that it is a valid PDF.

Hint: This can be done using Bayes' rule without having to know the marginal PDF of T , by recognizing what the conditional PDF is up to a normalizing constant—then the normalizing constant must be whatever is needed to make the conditional PDF valid.

(d) In Example 8.2.4, we will show that the marginal PDF of T is $f_T(t) = \lambda^2 t e^{-\lambda t}$, for $t > 0$. Give a short alternative proof of this fact, based on the previous parts and Bayes' rule.

11. Let X, Y, Z be r.v.s such that $X \sim \mathcal{N}(0, 1)$ and conditional on $X = x$, Y and Z are i.i.d. $\mathcal{N}(x, 1)$.

(a) Find the joint PDF of X, Y, Z .

(b) By definition, Y and Z are conditionally independent given X . Discuss intuitively whether or not Y and Z are also unconditionally independent.

(c) Find the joint PDF of Y and Z . You can leave your answer as an integral, though the integral can be done with some algebra (such as completing the square) and facts about the Normal distribution.

12. Let $X \sim \text{Expo}(\lambda)$, and let c be a positive constant.

(a) If you remember the memoryless property, you already know that the conditional distribution of X given $X > c$ is the same as the distribution of $c + X$ (think of waiting c minutes for a "success" and then having a fresh $\text{Expo}(\lambda)$ additional waiting time). Derive this in another way, by finding the conditional CDF of X given $X > c$ and the conditional PDF of X given $X > c$.

(b) Find the conditional CDF of X given $X < c$ and the conditional PDF of X given $X < c$.

13. Let X and Y be i.i.d. $\text{Expo}(\lambda)$. Find the conditional distribution of X given $X < Y$ in two different ways:

(a) by using calculus to find the conditional PDF.

(b) without using calculus, by arguing that the conditional distribution of X given $X < Y$ is the same distribution as the unconditional distribution of $\min(X, Y)$, and then applying an earlier result about the minimum of independent Exponentials.

14. (S) (a) A stick is broken into three pieces by picking two points independently and uniformly along the stick, and breaking the stick at those two points. What is the probability that the three pieces can be assembled into a triangle?

Hint: A triangle can be formed from 3 line segments of lengths a, b, c if and only if $a, b, c \in (0, 1/2)$. The probability can be interpreted geometrically as proportional to an area in the plane, avoiding all calculus, but make sure for that approach that the distribution of the random point in the plane is Uniform over some region.

(b) Three legs are positioned uniformly and independently on the perimeter of a round table. What is the probability that the table will stand?

15. Let X and Y be continuous r.v.s., with joint CDF $F(x, y)$. Show that the probability that (X, Y) falls into the rectangle $[a_1, a_2] \times [b_1, b_2]$ is

$$F(a_2, b_2) - F(a_1, b_2) + F(a_1, b_1) - F(a_2, b_1).$$

16. Let X and Y have joint PDF

$$f_{X,Y}(x, y) = x + y, \text{ for } 0 < x < 1 \text{ and } 0 < y < 1.$$

(a) Check that this is a valid joint PDF.

(b) Are X and Y independent?

(c) Find the marginal PDFs of X and Y .

(d) Find the conditional PDF of Y given $X = x$.

17. Let X and Y have joint PDF

$$f_{X,Y}(x, y) = cxy, \text{ for } 0 < x < y < 1.$$

(a) Find c to make this a valid joint PDF.

(b) Are X and Y independent?

(c) Find the marginal PDFs of X and Y .

(d) Find the conditional PDF of Y given $X = x$.

18. (S) Let (X, Y) be a uniformly random point in the triangle in the plane with vertices $(0, 0)$, $(0, 1)$, $(1, 0)$. Find the joint PDF of X and Y , the marginal PDF of X , and the conditional PDF of X given Y .

19. (S) A random point (X, Y, Z) is chosen uniformly in the ball

$$B = \{(x, y, z) : x^2 + y^2 + z^2 \leq 1\}.$$

(a) Find the joint PDF of X, Y, Z .

(b) Find the joint PDF of X, Y .

(c) Find an expression for the marginal PDF of X , as an integral.

20. (S) Let U_1, U_2, U_3 be i.i.d. $\text{Unif}(0, 1)$, and let $L = \min(U_1, U_2, U_3)$, $M = \max(U_1, U_2, U_3)$.

(a) Find the marginal CDF and marginal PDF of M , and the joint CDF and joint PDF of L, M .

Hint: For the latter, start by considering $P(L \geq l, M \leq m)$.

(b) Find the conditional PDF of M given L .

21. Find the probability that the quadratic polynomial $Ax^2 + Bx + 1$, where the coefficients A and B are determined by drawing i.i.d. $\text{Unif}(0, 1)$ random variables, has at least one real root.

Hint: By the quadratic formula, the polynomial $ax^2 + bx + c$ has a real root if and only if $b^2 - 4ac \geq 0$.

22. Let X and Y each have support $(0, \infty)$ marginally, and suppose that the joint PDF $f_{X,Y}$ of X and Y is positive for $0 < x < y$ and 0 otherwise.

(a) What is the support of the conditional PDF of Y given $X = x$?

(b) Show that X and Y can't be independent.

23. The volume of a region in n -dimensional Euclidean space $\mathbb{R}^n$ is the integral of 1 over that

region. The unit ball in $\mathbb{R}^n$ is $\{(x_1, \dots, x_n) : x_1^2 + \dots + x_n^2 \leq 1\}$, the ball of radius 1 centered at 0. As mentioned in Section A. 7 of the math appendix, the volume of the unit ball in n dimensions is

$$v_n = \frac{\pi^{n/2}}{\Gamma(n/2 + 1)}$$

where Γ is the gamma function, a very famous function which is defined by

$$\Gamma(a) = \int_0^\infty x^a e^{-x} \frac{dx}{x}$$

for all $a > 0$, and which will play an important role in the next chapter. A few useful facts about the gamma function (which you can assume) are that $\Gamma(a + 1) = a\Gamma(a)$ for any $a > 0$, and that $\Gamma(1) = 1$ and $\Gamma(\frac{1}{2}) = \sqrt{\pi}$. Using these facts, it follows that $\Gamma(n) = (n - 1)!$ for n a positive integer, and we can also find $\Gamma(n + \frac{1}{2})$ when n is a nonnegative integer. For practice, verify that $v_2 = \pi$ (the area of the unit disk in 2 dimensions) and $v_3 = \frac{4}{3}\pi$ (the volume of the unit ball in 3 dimensions).

Let $U_1, U_2, \dots, U_n \sim \text{Unif}(-1, 1)$ be i.i.d.

- Find the probability that $(U_1, U_2, \dots, U_n)$ is in the unit ball in $\mathbb{R}^n$.
- Evaluate the result from (a) numerically for $n = 1, 2, \dots, 10$, and plot the results (using a computer unless you are extremely good at making hand-drawn graphs). The facts above about the gamma function are sufficient so that you can do this without doing any integrals, but you can also use the command `gamma` in R to compute the gamma function.
- Let c be a constant with $0 < c < 1$, and let X_n count how many of the U_j satisfy $|U_j| > c$. What is the distribution of X_n ?
- For $c = 1/\sqrt{2}$, use the result of Part (c) to give a simple, short derivation of what happens to the probability from (a) as $n \rightarrow \infty$.

24. (S) Two students, A and B , are working independently on homework assignments. Student A takes $Y_1 \sim \text{Expo}(\lambda_1)$ hours to finish their homework, while B takes $Y_2 \sim \text{Expo}(\lambda_2)$ hours.

- Find the CDF and PDF of Y_1/Y_2 , the ratio of their problem-solving times.
- Find the probability that A finishes their homework before B does.

25. Two companies, Company 1 and Company 2, have just been founded. Stock market crashes occur according to a Poisson process with rate λ_0 . Such a crash would put both companies out of business. For $j \in \{1, 2\}$, there may be an adverse event of type j , which puts Company j out of business (if it is not already out of business) but does not affect the other company; such events occur according to a Poisson process with rate λ_j . If there has not been a stock market crash or an adverse event of type j , then Company j remains in business. The three Poisson processes are independent of each other. Let X_1 and X_2 be how long Company 1 and Company 2 stay in business, respectively.

- Find the marginal distributions of X_1 and X_2 .
- Find $P(X_1 > x_1, X_2 > x_2)$, and use this to find the joint CDF of X_1 and X_2 .

26. (S) The bus company from Blissville decides to start service in Blotchville, sensing a promising business opportunity. Meanwhile, Fred has moved back to Blotchville. Now when Fred arrives at the bus stop, either of two independent bus lines may come by (both of which take him home). The Blissville company's bus arrival times are exactly 10 minutes apart, whereas the time from one Blotchville company bus to the next is $\text{Expo}(\frac{1}{10})$. Fred arrives at a uniformly random time on a certain day.

- What is the probability that the Blotchville company bus arrives first?

Hint: One good way is to use the continuous law of total probability.

- What is the CDF of Fred's waiting time for a bus?

27. A longevity study is being conducted on n married hobbit couples. Let p be the probability

that an individual hobbit lives at least until their eleventy-first birthday, and assume that the lifespans of different hobbits are independent. Let N_0, N_1, N_2 be the number of couples in which neither hobbit reaches age eleventy-one, one hobbit does but not the other, and both hobbits reach eleventy-one, respectively.

(a) Find the joint PMF of N_0, N_1, N_2 .

For the rest of this problem, suppose that it is observed that exactly h of the cohort of hobbits reach their eleventy-first birthdays.

(b) Using (a) and the definition of conditional probability, find the conditional PMF of N_2 given this information, up to a normalizing constant (that is, you do not need to find the normalizing constant in this part, but just to give a simplified expression that is proportional to the conditional PMF). For simplicity, you can and should ignore multiplicative constants in this part; this includes multiplicative factors that are functions of h , since h is now being treated as a known constant.

(c) Now obtain the conditional PMF of N_2 using a direct counting argument, now including any normalizing constants needed in order to have a valid conditional PMF.

(d) Discuss intuitively whether or not p should appear in the answer to (c).

(e) What is the conditional expectation of N_2 , given the above information (simplify fully)? This can be done without doing any messy sums.

28. There are n stores in a shopping center, labeled from 1 to n . Let X_i be the number of customers who visit store i in a particular month, and suppose that $X_1, X_2, \dots, X_n$ are i.i.d. with PMF $p(x) = P(X_i = x)$. Let $I \sim \text{DUnif}(1, 2, \dots, n)$ be the label of a randomly chosen store, so X_I is the number of customers at a randomly chosen store.

(a) For $i \neq j$, find $P(X_i = X_j)$ in terms of a sum involving the PMF $p(x)$.

(b) Find the joint PMF of I and X_I . Are they independent?

(c) Does X_I , the number of customers for a random store, have the same marginal distribution as X_1 , the number of customers for store 1?

(d) Let $J \sim \text{DUnif}(1, 2, \dots, n)$ also be the label of a randomly chosen store, with I and J independent. Find $P(X_I = X_J)$ in terms of a sum involving the PMF $p(x)$. How does $P(X_I = X_J)$ compare to $P(X_i = X_j)$ for fixed i, j with $i \neq j$?

29. Let X and Y be i.i.d. $\text{Geom}(p)$, $L = \min(X, Y)$, and $M = \max(X, Y)$.

(a) Find the joint PMF of L and M . Are they independent?

(b) Find the marginal distribution of L in two ways: using the joint PMF, and using a story.

(c) Find EM .

Hint: A quick way is to use (b) and the fact that $L + M = X + Y$.

(d) Find the joint PMF of L and $M - L$. Are they independent?

30. Let X, Y have the joint CDF

$$F(x, y) = 1 - e^{-x} - e^{-y} + e^{-(x+y+\theta xy)}$$

for $x > 0, y > 0$ (and $F(x, y) = 0$ otherwise), where the parameter θ is in $[0, 1]$.

(a) Find the joint PDF of X, Y . For which values of θ (if any) are they independent?

(b) Explain why we require θ to be in $[0, 1]$.

(c) Find the marginal PDFs of X and Y by working directly from the joint PDF from (a). When integrating, do not use integration by parts or computer assistance; rather, pattern-match to facts we know about moments of famous distributions.

(d) Find the marginal CDFs of X and Y by working directly from the joint CDF.

2D Lotus

31. (S) Let X and Y be i.i.d. $\text{Unif}(0, 1)$. Find the standard deviation of the distance between X and Y .

32. (S) Let X, Y be i.i.d. $\text{Expo}(\lambda)$. Find $E|X - Y|$ in two different ways: (a) using 2D LOTUS and (b) using the memoryless property without any calculus.
33. Alice walks into a post office with 2 clerks. Both clerks are in the midst of serving customers, but Alice is next in line. The clerk on the left takes an $\text{Expo}(\lambda_1)$ time to serve a customer, and the clerk on the right takes an $\text{Expo}(\lambda_2)$ time to serve a customer. Let T be the amount of time Alice has to wait until it is her turn.
- (a) Write down expressions for the mean and variance of T , in terms of double integrals (which you do not need to evaluate).
- (b) Find the distribution, mean, and variance of T , without using calculus.
34. Let (X, Y) be a uniformly random point in the triangle in the plane with vertices $(0, 0), (0, 1), (1, 0)$. Find $\text{Cov}(X, Y)$. (Exercise 18 is about joint, marginal, and conditional PDFs in this setting.)
35. A random point is chosen uniformly in the unit disk $\{(x, y) : x^2 + y^2 \leq 1\}$. Let R be its distance from the origin.
- (a) Find $E(R)$ using 2D LOTUS.

Hint: To do the integral, convert to polar coordinates (see the math appendix).

- (b) Find the CDFs of R^2 and of R without using calculus, using the fact that for a Uniform distribution on a region, probability within that region is proportional to area. Then get the PDFs of R^2 and of R , and find $E(R)$ in two more ways: using the definition of expectation, and using a 1D LOTUS by thinking of R as a function of R^2 .
36. Let X and Y be discrete r.v.s.
- (a) Use 2D LOTUS (without assuming linearity) to show that $E(X + Y) = E(X) + E(Y)$.
- (b) Now suppose that X and Y are independent. Use 2D LOTUS to show that they are uncorrelated, i.e., $E(XY) = E(X)E(Y)$.
37. Let X and Y be i.i.d. continuous random variables with PDF f , mean μ , and variance σ^2 . We know that the expected squared distance of X from its mean is σ^2 , and likewise for Y ; this problem is about the expected squared distance of X from Y .
- (a) Use 2D LOTUS to express $E(X - Y)^2$ as a double integral.
- (b) By expanding $(x - y)^2 = x^2 - 2xy + y^2$ and evaluating the double integral from (a), show that

$$E(X - Y)^2 = 2\sigma^2$$

- (c) Give an alternative proof of the result from (b), based on the trick of adding and subtracting μ :

$$(X - Y)^2 = (X - \mu + \mu - Y)^2 = (X - \mu)^2 - 2(X - \mu)(Y - \mu) + (Y - \mu)^2.$$

Covariance

38. (S) Let X and Y be r.v.s. Is it correct to say " $\max(X, Y) + \min(X, Y) = X + Y$ "? Is it correct to say $\text{Cov}(\max(X, Y), \min(X, Y)) = \text{Cov}(X, Y)$ since either the max is X and the min is Y or vice versa, and covariance is symmetric"? Explain.
39. (S) Two fair, six-sided dice are rolled (one green and one orange), with outcomes X and Y for the green die and the orange die, respectively.
- (a) Compute the covariance of $X + Y$ and $X - Y$.
- (b) Are $X + Y$ and $X - Y$ independent?

40. Let X and Y be i.i.d. $\text{Unif}(0, 1)$.
 (a) Compute the covariance of $X + Y$ and $X - Y$.
 (b) Are $X + Y$ and $X - Y$ independent?
41. (S) Let X and Y be standardized r.v.s (i.e., marginally they each have mean 0 and variance 1) with correlation $\rho \in (-1, 1)$. Find a, b, c, d (in terms of ρ) such that $Z = aX + bY$ and $W = cX + dY$ are uncorrelated but still standardized.
42. (S) Let X be the number of distinct birthdays in a group of 110 people (i.e., the number of days in a year such that at least one person in the group has that birthday). Under the usual assumptions (no February 29, all the other 365 days of the year are equally likely, and the day when one person is born is independent of the days when the other people are born), find the mean and variance of X .
43. (a) Let X and Y be Bernoulli r.v.s, possibly with different parameters. Show that if X and Y are uncorrelated, then they are independent.
 (b) Give an example of three Bernoulli r.v.s such that each pair of them is uncorrelated, yet the three r.v.s are dependent.
44. Find the variance of the number of toys needed until you have a complete set in Example 4.3.12 (the coupon collector problem), as a sum.
45. A random triangle is formed in some way, such that all pairs of angles have the same joint distribution. What is the correlation between two of the angles (assuming that the variance of the angles is nonzero)?
46. Each of $n \geq 2$ people puts their name on a slip of paper (no two have the same name). The slips of paper are shuffled in a hat, and then each person draws one (uniformly at random at each stage, without replacement). Find the standard deviation of the number of people who draw their own names.
47. As in Example 4.4.7, an urn contains w white balls and b black balls. The balls are randomly drawn one by one without replacement until r white balls have been drawn. Let $X \sim \text{NHGeom}(w, b, r)$ be the number of black balls drawn before drawing the r th white ball. In this exercise, you will derive $\text{Var}(X)$.
 As explained in Example 4.4.7, we can assume that we continue drawing balls until the urn has been emptied out. Label the black balls as $1, 2, \dots, b$, and write $X = \sum_{j=1}^b I_j$, where I_j is the indicator of black ball j being drawn before the r th white ball is drawn.
 (a) Show that

$$E(I_j) = \frac{r}{w+1}.$$

- (b) Give an intuitive explanation of whether I_i and I_j are positively correlated, uncorrelated, or negatively correlated, for $i \neq j$.
 (c) Show that for $i \neq j$,

$$E(I_i I_j) = \frac{\binom{r+1}{2}}{\binom{w+2}{2}} = \frac{(r+1)r}{(w+2)(w+1)}$$

Hint: Imagine $w+2$ slots, into which black balls i and j and the w white balls will be placed. All orderings for these $w+2$ balls are equally likely.

- (d) Find an expression for $\text{Var}(X)$. With some algebra (which you don't have to do), your expression should simplify to

$$\text{Var}(X) = \frac{rb(w+b+1)(w-r+1)}{(w+1)^2(w+2)}$$

48. (S) Athletes compete one at a time at the high jump. Let X_j be how high the j th jumper jumped, with $X_1, X_2, \dots$ i.i.d. with a continuous distribution. We say that the j th jumper sets a record if X_j is greater than all of $X_{j-1}, \dots, X_1$.

Find the variance of the number of records among the first n jumpers (as a sum). What happens to the variance as $n \rightarrow \infty$?

49. (S) A chicken lays a $\text{Pois}(\lambda)$ number N of eggs. Each egg hatches a chick with probability p , independently. Let X be the number which hatch, so $X | N = n \sim \text{Bin}(n, p)$.

Find the correlation between N (the number of eggs) and X (the number of eggs which hatch). Simplify; your final answer should work out to a simple function of p (the λ should cancel out).

50. Let $X_1, \dots, X_n$ be random variables such that $\text{Corr}(X_i, X_j) = \rho$ for all $i \neq j$. Show that $\rho \geq -\frac{1}{n-1}$. This is a bound on how negatively correlated a collection of r.v.s can all be with each other.

Hint: Assume $\text{Var}(X_i) = 1$ for all i ; this can be done without loss of generality, since rescaling two r.v.s does not affect the correlation between them. Then use the fact that $\text{Var}(X_1 + \dots + X_n) \geq 0$.

51. Let X and Y be independent r.v.s. Show that

$$\text{Var}(XY) = \text{Var}(X) \text{Var}(Y) + (EX)^2 \text{Var}(Y) + (EY)^2 \text{Var}(X)$$

Hint: It is often useful when working with a second moment $E(T^2)$ to express it as $\text{Var}(T) + (ET)^2$.

52. Stat 110 shirts come in 3 sizes: small, medium, and large. There are n shirts of each size (where $n \geq 2$). There are $3n$ students. For each size, n of the students have that size as the best fit. This seems ideal. But suppose that instead of giving each student the right size shirt, each student is given a shirt completely randomly (all allocations of the shirts to the students, with one shirt per student, are equally likely). Let X be the number of students who get their right size shirt.

(a) Find $E(X)$.

(b) Give each student an ID number from 1 to $3n$, such that the right size shirt is small for students 1 through n , medium for students $n+1$ through $2n$, and large for students $2n+1$ through $3n$. Let A_j be the event that student j gets their right size shirt. Find $P(A_1, A_2)$ and $P(A_1, A_{n+1})$.

(c) Find $\text{Var}(X)$.

53. (S) A drunken man wanders around randomly in a large space. At each step, he moves one unit of distance North, South, East, or West, with equal probabilities. Choose coordinates such that his initial position is $(0, 0)$ and if he is at (x, y) at some time, then one step later he is at $(x, y+1)$, $(x, y-1)$, $(x+1, y)$, or $(x-1, y)$. Let (X_n, Y_n) and R_n be his position and distance from the origin after n steps, respectively.

General hint: Note that X_n is a sum of r.v.s with possible values $-1, 0, 1$, and likewise for Y_n , but be careful throughout the problem about independence.

(a) Determine whether or not X_n is independent of Y_n .

(b) Find $\text{Cov}(X_n, Y_n)$.

(c) Find $E(R_n^2)$.

54. (S) A scientist makes two measurements, considered to be independent standard Normal r.v.s. Find the correlation between the larger and smaller of the values.

Hint: Note that $\max(x, y) + \min(x, y) = x + y$ and $\max(x, y) - \min(x, y) = |x - y|$.

55. Let $U \sim \text{Unif}(-1, 1)$ and $V = 2|U| - 1$.

(a) Find the distribution of V (give the PDF and, if it is a named distribution we have studied, its name and parameters).

Hint: Find the support of V , and then find the CDF of V by reducing $P(V \leq v)$ to probability calculations about U .

(b) Show that U and V are uncorrelated, but not independent. This is also another example illustrating the fact that knowing the marginal distributions of two r.v.s does not determine the joint distribution.

56. (S) Consider the following method for creating a bivariate Poisson (a joint distribution for two r.v.s such that both marginals are Poissons). Let $X = V + W, Y = V + Z$ where V, W, Z are i.i.d. $\text{Pois}(\lambda)$ (the idea is to have something borrowed and something new but not something old or something blue).

(a) Find $\text{Cov}(X, Y)$.

(b) Are X and Y independent? Are they conditionally independent given V ?

(c) Find the joint PMF of X, Y (as a sum).

57. You are playing an exciting game of Battleship. Your opponent secretly positions ships on a 10 by 10 grid and you try to guess where the ships are. Each of your guesses is a hit if there is a ship there and a miss otherwise.

The game has just started and your opponent has 3 ships: a battleship (length 4), a submarine (length 3), and a destroyer (length 2). (Usually there are 5 ships to start, but to simplify the calculations we are considering 3 here.) You are playing a variation in which you unleash a salvo, making 5 simultaneous guesses. Assume that your 5 guesses are a simple random sample drawn from the 100 grid positions, i.e., all sets of 5 grid positions are equally likely.

Find the mean and variance of the number of distinct ships you will hit in your salvo. (Give exact answers in terms of binomial coefficients or factorials, and also numerical values computed using a computer.)

Hint: First work in terms of the number of ships missed, expressing this as a sum of indicator r.v.s. Then use the fundamental bridge and naive definition of probability.

58. This problem explores a visual interpretation of covariance. Data are collected for n individuals, where for each individual two variables are measured (e.g., height and weight). Assume independence across individuals (e.g., person 1's variables gives no information about the other people), but not within individuals (e.g., a person's height and weight may be correlated).

Let $(x_1, y_1), \dots, (x_n, y_n)$ be the n data points, with $n \geq 2$. The data are considered here as fixed, known numbers—they are the observed values after performing an experiment. Imagine plotting all the points (x_i, y_i) in the plane, and drawing the rectangle determined by each pair of points. For example, the points $(1, 3)$ and $(4, 6)$ determine the rectangle with vertices $(1, 3), (1, 6), (4, 6), (4, 3)$.

The signed area contributed by (x_i, y_i) and (x_j, y_j) is the area of the rectangle they determine if the slope of the line between them is positive, and is the negative of the area of the rectangle they determine if the slope of the line between them is negative. (Define the signed area to be 0 if $x_i = x_j$ or $y_i = y_j$, since then the rectangle is degenerate.) So the signed area is positive if a higher x value goes with a higher y value for the pair of points, and negative otherwise. Assume that the x_i are all distinct and the y_i are all distinct.

(a) The sample covariance of the data is defined to be

$$r = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y}),$$

where

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \text{ and } \bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$$

are the sample means. (There are differing conventions about whether to divide by $n - 1$ or n in the definition of sample covariance, but that need not concern us here.)

Let (X, Y) be one of the (x_i, y_i) pairs, chosen uniformly at random. Determine precisely how $\text{Cov}(X, Y)$ is related to the sample covariance.

(b) Let (X, Y) be as in (a), and $(\tilde{X}, \tilde{Y})$ be an independent draw from the same distribution. That is, (X, Y) and $(\tilde{X}, \tilde{Y})$ are randomly chosen from the n points, independently (so it is possible for the same point to be chosen twice).

Express the total signed area of the rectangles as a constant times $E((X - \tilde{X})(Y - \tilde{Y}))$. Then show that the sample covariance of the data is a constant times the total signed area of the rectangles.

Hint: Consider $E((X - \tilde{X})(Y - \tilde{Y}))$ in two ways: as the average signed area of the random rectangle formed by (X, Y) and $(\tilde{X}, \tilde{Y})$, and using properties of expectation to relate it to $\text{Cov}(X, Y)$. For the former, consider the n^2 possibilities for which point (X, Y) is and which point $(\tilde{X}, \tilde{Y})$; note that n such choices result in degenerate rectangles.

(c) Based on the interpretation from (b), give intuitive explanations of why for any r.v.s W_1, W_2, W_3 and constants a_1, a_2 , covariance has the following properties:

- (i) $\text{Cov}(W_1, W_2) = \text{Cov}(W_2, W_1)$;
- (ii) $\text{Cov}(a_1 W_1, a_2 W_2) = a_1 a_2 \text{Cov}(W_1, W_2)$;
- (iii) $\text{Cov}(W_1 + a_1, W_2 + a_2) = \text{Cov}(W_1, W_2)$;
- (iv) $\text{Cov}(W_1, W_2 + W_3) = \text{Cov}(W_1, W_2) + \text{Cov}(W_1, W_3)$.

59. A statistician is trying to estimate an unknown parameter θ based on some data. She has available two independent estimators $\hat{\theta}_1$ and $\hat{\theta}_2$ (an estimator is a function of the data, used to estimate a parameter). For example, $\hat{\theta}_1$ could be the sample mean of a subset of the data and $\hat{\theta}_2$ could be the sample mean of another subset of the data, disjoint from the subset used to calculate $\hat{\theta}_1$. Assume that both of these estimators are unbiased, i.e., $E(\hat{\theta}_j) = \theta$.

Rather than having a bunch of separate estimators, the statistician wants one combined estimator. It may not make sense to give equal weights to $\hat{\theta}_1$ and $\hat{\theta}_2$ since one could be much more reliable than the other, so she decides to consider combined estimators of the form

$$\hat{\theta} = w_1 \hat{\theta}_1 + w_2 \hat{\theta}_2,$$

a weighted combination of the two estimators. The weights w_1 and w_2 are nonnegative and satisfy $w_1 + w_2 = 1$.

(a) Check that $\hat{\theta}$ is also unbiased, i.e., $E(\hat{\theta}) = \theta$.

(b) Determine the optimal weights w_1, w_2 , in the sense of minimizing the mean squared error $E(\hat{\theta} - \theta)^2$. Express your answer in terms of the variances of $\hat{\theta}_1$ and $\hat{\theta}_2$. The optimal weights are known as Fisher weights.

Hint: As discussed in Exercise 53 from Chapter 5, mean squared error is variance plus squared bias, so in this case the mean squared error of $\hat{\theta}$ is $\text{Var}(\hat{\theta})$. Note that there is no need for multivariable calculus here, since $w_2 = 1 - w_1$.

(c) Give a simple description of what the estimator found in (b) amounts to if the data are i.i.d. random variables $X_1, \dots, X_n, Y_1, \dots, Y_m$, $\hat{\theta}_1$ is the sample mean of $X_1, \dots, X_n$, and $\hat{\theta}_2$ is the sample mean of $Y_1, \dots, Y_m$.

\{\}paragraphChicken-egg

60. (S) A $\text{Pois}(\lambda)$ number of people vote in a certain election. Each voter votes for candidate A with probability p and for candidate B with probability $q = 1 - p$, independently of all the other voters. Let V be the difference in votes, defined as the number of votes for A minus the number for B .

- (a) Find $E(V)$.
- (b) Find $\text{Var}(V)$.

61. A traveler gets lost $N \sim \text{Pois}(\lambda)$ times on a long journey. When lost, the traveler asks someone for directions with probability p . Let X be the number of times that the traveler is lost and asks for directions, and Y be the number of times that the traveler is lost and does not ask for directions.
- Find the joint PMF of N, X, Y . Are they independent?
 - Find the joint PMF of N, X . Are they independent?
 - Find the joint PMF of X, Y . Are they independent?
62. The number of people who visit the Leftorium store in a day is $\text{Pois}(100)$. Suppose that 10% of customers are sinister (left-handed), and 90% are dexterous (right-handed). Half of the sinister customers make purchases, but only a third of the dexterous customers make purchases. The characteristics and behavior of people are independent, with probabilities as described in the previous two sentences. On a certain day, there are 42 people who arrive at the store but leave without making a purchase. Given this information, what is the conditional PMF of the number of customers on that day who make a purchase?
63. A chicken lays n eggs. Each egg independently does or doesn't hatch, with probability p of hatching. For each egg that hatches, the chick does or doesn't survive (independently of the other eggs), with probability s of survival. Let $N \sim \text{Bin}(n, p)$ be the number of eggs which hatch, X be the number of chicks which survive, and Y be the number of chicks which hatch but don't survive (so $X + Y = N$). Find the marginal PMF of X , and the joint PMF of X and Y . Are X and Y independent?
64. There will be $X \sim \text{Pois}(\lambda)$ courses offered at a certain school next year.
- Find the expected number of choices of 4 courses (in terms of λ , fully simplified), assuming that simultaneous enrollment is allowed if there are time conflicts.
 - Now suppose that simultaneous enrollment is not allowed. Suppose that most faculty only want to teach on Tuesdays and Thursdays, and most students only want to take courses that start at 10 am or later, and as a result there are only four possible time slots: 10am, 11 : 30am, 1pm, 2 : 30pm (each course meets Tuesday-Thursday for an hour and a half, starting at one of these times). Rather than trying to avoid major conflicts, the school schedules the courses completely randomly: after the list of courses for next year is determined, they randomly get assigned to time slots, independently and with probability $1/4$ for each time slot.

Let X_{am} and X_{pm} be the number of morning and afternoon courses for next year, respectively (where "morning" means starting before noon). Find the joint PMF of X_{am} and X_{pm} , i.e., find $P(X_{\text{am}} = a, X_{\text{pm}} = b)$ for all a, b .

- (c) Continuing as in (b), let X_1, X_2, X_3, X_4 be the number of 10 am , 11:30 am, 1 pm , 2:30 pm courses for next year, respectively. What is the joint distribution of X_1, X_2, X_3, X_4 ? (The result is completely analogous to that of $X_{\text{am}}, X_{\text{pm}}$; you can derive it by thinking conditionally, but for this part you are also allowed to just use the fact that the result is analogous to that of (b).) Use this to find the expected number of choices of 4 nonconflicting courses (in terms of λ , fully simplified). What is the ratio of the expected value from (a) to this expected value?

Multinomial

65. (S) Let $(X_1, \dots, X_k)$ be Multinomial with parameters n and $(p_1, \dots, p_k)$. Use indicator r.v.s to show that $\text{Cov}(X_i, X_j) = -np_i p_j$ for $i \neq j$.
66. (S) Consider the birthdays of 100 people. Assume people's birthdays are independent, and the 365 days of the year (exclude the possibility of February 29) are equally likely. Find

- the covariance and correlation between how many of the people were born on January 1 and how many were born on January 2.
67. A certain course has a freshmen, b sophomores, c juniors, and d seniors. Let X be the number of freshmen and sophomores (total), Y be the number of juniors, and Z be the number of seniors in a random sample of size n , where for Part (a) the sampling is with replacement and for Part (b) the sampling is without replacement (for both parts, at each stage the allowed choices have equal probabilities).
- Find the joint PMF of X, Y, Z , for sampling with replacement.
 - Find the joint PMF of X, Y, Z , for sampling without replacement.
68. (S) A group of $n \geq 2$ people decide to play an exciting game of Rock-Paper-Scissors. As you may recall, Rock smashes Scissors, Scissors cuts Paper, and Paper covers Rock (despite Bart Simpson once saying "Good old rock, nothing beats that!"). Usually this game is played with 2 players, but it can be extended to more players as follows. If exactly 2 of the 3 choices appear when everyone reveals their choice, say $a, b \in \{\text{Rock, Paper, Scissors}\}$ where a beats b , the game is decisive: the players who chose a win, and the players who chose b lose. Otherwise, the game is indecisive and the players play again.
- For example, with 5 players, if one player picks Rock, two pick Scissors, and two pick Paper, the round is indecisive and they play again. But if 3 pick Rock and 2 pick Scissors, then the Rock players win and the Scissors players lose the game.
- Assume that the n players independently and randomly choose between Rock, Scissors, and Paper, with equal probabilities. Let X, Y, Z be the number of players who pick Rock, Scissors, Paper, respectively in one game.
- Find the joint PMF of X, Y, Z .
 - Find the probability that the game is decisive. Simplify your answer.
 - What is the probability that the game is decisive for $n = 5$? What is the limiting probability that a game is decisive as $n \rightarrow \infty$? Explain briefly why your answer makes sense.
69. (S) Emails arrive in an inbox according to a Poisson process with rate λ (so the number of emails in a time interval of length t is distributed as $\text{Pois}(\lambda t)$, and the numbers of emails arriving in disjoint time intervals are independent). Let X, Y, Z be the numbers of emails that arrive from 9 am to noon, noon to 6 pm, and 6 pm to midnight (respectively) on a certain day.
- Find the joint PMF of X, Y, Z .
 - Find the conditional joint PMF of X, Y, Z given that $X + Y + Z = 36$.
 - Find the conditional PMF of $X + Y$ given that $X + Y + Z = 36$.
 - Find $E(X + Y \mid X + Y + Z = 36)$ and $\text{Var}(X + Y \mid X + Y + Z = 36)$. (Conditional expectation and conditional variance given an event are defined in the same way as expectation and variance, using the conditional distribution given the event in place of the unconditional distribution; these concepts are explored more in Chapter 9.)
70. Let X be the number of statistics majors in a certain college in the Class of 2030, viewed as an r.v. Each statistics major chooses between two tracks: a general track in statistical principles and methods, and a track in quantitative finance. Suppose that each statistics major chooses randomly which of these two tracks to follow, independently, with probability p of choosing the general track. Let Y be the number of statistics majors who choose the general track, and Z be the number of statistics majors who choose the quantitative finance track.
- Suppose that $X \sim \text{Pois}(\lambda)$. (This isn't the exact distribution in reality since a Poisson

is unbounded, but it may be a very good approximation.) Find the correlation between X and Y .

(b) Let n be the size of the Class of 2030, where n is a known constant. For this part and the next, instead of assuming that X is Poisson, assume that each of the n students chooses to be a statistics major with probability r , independently. Find the joint distribution of Y, Z , and the number of non-statistics majors, and their marginal distributions.

(c) Continuing as in (b), find the correlation between X and Y .

71. In humans (and many other organisms), genes come in pairs. Consider a gene of interest, which comes in two types (alleles): type a and type A . The genotype of a person for that gene is the types of the two genes in the pair: AA, Aa , or aa (aA is equivalent to Aa). According to the Hardy-Weinberg law, for a population in equilibrium, the frequencies of AA, Aa, aa will be $p^2, 2p(1-p), (1-p)^2$, respectively, for some p with $0 < p < 1$. Suppose that the Hardy-Weinberg law holds, and that n people are drawn randomly from the population, independently. Let X_1, X_2, X_3 be the number of people in the sample with genotypes AA, Aa, aa , respectively.

(a) What is the joint PMF of X_1, X_2, X_3 ?

(b) What is the distribution of the number of people in the sample who have an A ?

(c) What is the distribution of how many of the $2n$ genes among the people are A 's?

(d) Now suppose that p is unknown, and must be estimated using the observed data X_1, X_2, X_3 . The maximum likelihood estimator (MLE) of p is the value of p for which the observed data are as likely as possible. Find the MLE of p .

(e) Now suppose that p is unknown, and that our observations can't distinguish between AA and Aa . So for each person in the sample, we just know whether or not that person is an aa (in genetics terms, AA and Aa have the same phenotype, and we only get to observe the phenotypes, not the genotypes). Find the MLE of p .

Multivariate Normal

72. (S) Let (X, Y) be Bivariate Normal, with X and Y marginally $\mathcal{N}(0, 1)$ and with correlation ρ between X and Y .

(a) Show that $(X + Y, X - Y)$ is also Bivariate Normal.

(b) Find the joint PDF of $X + Y$ and $X - Y$ (without using calculus), assuming that $-1 < \rho < 1$.

73. Let the joint PDF of X and Y be

$$f_{X,Y}(x, y) = c \exp\left(-\frac{x^2}{2} - \frac{y^2}{2}\right) \text{ for all } x \text{ and } y$$

where c is a constant.

(a) Find c to make this a valid joint PDF.

(b) What are the marginal distributions of X and Y ? Are X and Y independent?

(c) Is (X, Y) Bivariate Normal?

74. Let the joint PDF of X and Y be

$$f_{X,Y}(x, y) = c \exp\left(-\frac{x^2}{2} - \frac{y^2}{2}\right) \text{ for } xy > 0$$

where c is a constant (the joint PDF is 0 for $xy \leq 0$).

(a) Find c to make this a valid joint PDF.

(b) What are the marginal distributions of X and Y ? Are X and Y independent?

(c) Is (X, Y) Bivariate Normal?

75. Let X, Y, Z be i.i.d. $\mathcal{N}(0, 1)$. Find the joint MGF of $(X + 2Y, 3X + 4Z, 5Y + 6Z)$.

76. Let X and Y be i.i.d. $\mathcal{N}(0, 1)$, and let S be a random sign (1 or -1, with equal probabilities) independent of (X, Y) .

(a) Determine whether or not $(X, Y, X + Y)$ is Multivariate Normal.

(b) Determine whether or not $(X, Y, SX + SY)$ is Multivariate Normal.

(c) Determine whether or not (SX, SY) is Multivariate Normal.

77. Let (X, Y) be Bivariate Normal with $X \sim \mathcal{N}(0, \sigma_1^2)$ and $Y \sim \mathcal{N}(0, \sigma_2^2)$ marginally and with $\text{Corr}(X, Y) = \rho$. Find a constant c such that $Y - cX$ is independent of X .

Hint: First find c (in terms of ρ, σ_1, σ_2) such that $Y - cX$ and X are uncorrelated.

78. A mother and a father have 6 children. The 8 heights in the family (in inches) are $\mathcal{N}(\mu, \sigma^2)$ r.v.s (with the same distribution, but not necessarily independent).

(a) Assume for this part that the heights are all independent. On average, how many of the children are taller than both parents?

(b) Let X_1 be the height of the mother, X_2 be the height of the father, and $Y_1, \dots, Y_6$ be the heights of the children. Suppose that $(X_1, X_2, Y_1, \dots, Y_6)$ is Multivariate Normal, with $\mathcal{N}(\mu, \sigma^2)$ marginals and $\text{Corr}(X_1, Y_j) = \rho$ for $1 \leq j \leq 6$, with $\rho < 1$. On average, how many of the children are more than 1 inch taller than their mother?

Mixed Practice

79. Cars pass by a certain point on a road according to a Poisson process with rate λ cars/minute. Let $N_t \sim \text{Pois}(\lambda t)$ be the number of cars that pass by that point in the time interval $[0, t]$, with t measured in minutes.

(a) A certain device is able to count cars as they pass by, but it does not record the arrival times. At time 0, the counter on the device is reset to 0. At time 3 minutes, the device is observed and it is found that exactly 1 car had passed by. Given this information, find the conditional CDF of when that car arrived. Also describe in words what the result says.

(b) In the late afternoon, you are counting blue cars. Each car that passes by is blue with probability b , independently of all other cars. Find the joint PMF and marginal PMFs of the number of blue cars and number of non-blue cars that pass by the point in 10 minutes.

80. In a U.S. election, there will be $V \sim \text{Pois}(\lambda)$ registered voters. Suppose each registered voter is a registered Democrat with probability p and a registered Republican with probability $1 - p$, independent of other voters. Also, each registered voter shows up to the polls with probability s and stays home with probability $1 - s$, independent of other voters and independent of their own party affiliation. In this problem, we are interested in X , the number of registered Democrats who actually vote.

(a) What is the distribution of X , before we know anything about the number of registered voters?

(b) Suppose we learn that $V = v$; that is, v people registered to vote. What is the conditional distribution of X given this information?

(c) Suppose we learn there were d registered Democrats and r registered Republicans (where $d + r = v$). What is the conditional distribution of X given this information?

(d) Finally, we learn in addition to all of the above information that n people showed up at the polls on election day. What is the conditional distribution of X given this information?

81. A certain college has m freshmen, m sophomores, m juniors, and m seniors. A certain class at the college consists of a simple random sample of size n students, i.e., all sets of

n of the $4m$ students are equally likely. Let $X_1, \dots, X_4$ be the numbers of freshmen, ..., seniors in the class.

- (a) Find the joint PMF of X_1, X_2, X_3, X_4 .
- (b) Give both an intuitive explanation and a mathematical justification for whether or not the distribution from (a) is Multinomial.
- (c) Find $\text{Cov}(X_1, X_3)$, fully simplified.

Hint: Take the variance of both sides of $X_1 + X_2 + X_3 + X_4 = n$.

82. Let $X \sim \text{Expo}(\lambda)$ and let Y be a nonnegative random variable, discrete or continuous, whose MGF M is finite everywhere. Show that $P(Y < X) = M(c)$ for a certain value of c (which you should specify).

83. A publishing company employs two proofreaders, Prue and Frida. When Prue is proofreading a book, for each typo she has probability p of catching it and $q = 1 - p$ of missing it, independently. When Frida is proofreading a book, for each typo she has probability f of catching it and $g = 1 - f$ of missing it, independently.

- (a) A certain book draft has n typos. The company randomly assigns it to one of the two proofreaders, with equal probabilities. Find the distribution of the number of typos that get detected.
- (b) Another book is being written. When a draft of the book is complete, it will have a $\text{Pois}(\lambda)$ number of typos, and will be assigned to Prue to proofread. Find the probability that Prue catches exactly k typos, given that she misses exactly m typos.

84. Two authors, Bob and Martha, are about to begin writing an epic co-authored book, The Adventures of Aaron the Aardwolf. It will take them A years to write. When they finish this book, they will immediately begin work on new, individually authored books. Bob will spend X years writing The Bilinear Bonanza of Bonnie the Butterfly, and Martha will spend Y years writing Memoirs of Maude the Magnificent Mangabey, independently. Suppose that A, X, Y are i.i.d. $\text{Expo}(\lambda)$. On a timeline where time 0 is defined as the time when they begin their collaboration, consider the following quantities.

A: time at which The Adventures of Aaron the Aardwolf is completed;

B: time at which The Bilinear Bonanza of Bonnie the Butterfly is completed;

M : time at which Memoirs of Maude the Magnificent Mangabey is completed;

T : time at which the last to be completed of these three books is completed.

- (a) Find the distribution of B (which is also the distribution of M).
- (b) Find $\text{Cov}(A, B)$.
- (c) Find $E(T)$.

85. A DNA sequence can be represented as a sequence of letters, where the alphabet has 4 letters: A, C, G, T. Suppose that a random DNA sequence of length $n \geq 4$ is formed by independently generating letters one at a time, with p_A, p_C, p_G, p_T the probabilities of A, C, G, T, respectively, where $p_A + p_C + p_G + p_T = 1$.

- (a) Find the covariance between the number of A's and the number of C's in the sequence.
- (b) It is observed that the sequence contains exactly a A's, c C's, g G's, and t T's, where $a + c + g + t = n$ and $a \geq 2$. Given this information, find the probability that the first A in the sequence is followed immediately by another A.

Hint: How does this part relate to Exercise 74 in Chapter 2?

- (c) Given the information from (b) about how many times each letter occurs, find the expected number of occurrences of the expression CAT in the sequence.

86. To test for a certain disease, the level of a certain substance in the blood is measured. Let T be this measurement, considered as a continuous r.v. The patient tests positive (i.e., is declared to have the disease) if $T > t_0$ and tests negative if $T \leq t_0$, where t_0 is a threshold decided upon in advance. Let D be the indicator of having the disease. As discussed in Example 2.3.9, the sensitivity of the test is the probability of testing positive given that the patient has the disease,

and the specificity of the test is the probability of testing negative given that the patient does not have the disease.

(a) The ROC (receiver operator characteristic) curve of the test is the plot of sensitivity vs. 1 minus specificity, where sensitivity (the vertical axis) and 1 minus specificity (the horizontal axis) are viewed as functions of the threshold t_0 . ROC curves are widely used in medicine and engineering as a way to study the performance of procedures for classifying individuals into two groups (in this case, the two groups are "diseased people" and "non-diseased people").

Given that $D = 1, T$ has CDF G and PDF g ; given that $D = 0, T$ has CDF H and PDF h . Here g and h are positive on an interval $[a, b]$ and 0 outside this interval. Show that the area under the ROC curve is the probability that a randomly selected diseased person has a higher T value than a randomly selected non-diseased person.

(b) Explain why the result of (a) makes sense in two extreme cases: when $g = h$, and when there is a threshold t_0 such that $P(T > t_0 | D = 1)$ and $P(T \leq t_0 | D = 0)$ are very close to 1.

87. Let J be Discrete Uniform on $\{1, 2, \dots, n\}$.

(a) Find $E(J)$ and $\text{Var}(J)$, fully simplified, using results from Section A. 8 of the math appendix.

(b) Discuss intuitively whether the results in (a) should be approximately the same as the mean and variance (respectively) of a Uniform distribution on a certain interval.

(c) Let $X_1, \dots, X_n$ be i.i.d. $\mathcal{N}(0, 1)$ r.v.s, and let $R_1, \dots, R_n$ be their ranks (the smallest X_i has rank 1, the next has rank 2, ..., and the largest has rank n). Explain why

$$R_n = 1 + \sum_{j=1}^{n-1} I_j$$


where $I_j = I(X_n > X_j)$. Then use this to find $E(R_n)$ and $\text{Var}(R_n)$ directly using symmetry, linearity, the fundamental bridge, and properties of covariance.

(d) Explain how the results of (a) and (c) relate. Then prove the identities

$$\sum_{j=1}^n j = \frac{n(n+1)}{2} \text{ and } \sum_{j=1}^n j^2 = \frac{n(n+1)(2n+1)}{6}$$

by giving them probabilistic interpretations.

88. (S) A network consists of n nodes, each pair of which may or may not have an edge joining them. For example, a social network can be modeled as a group of n nodes (representing people), where an edge between i and j means they know each other. Assume the network is undirected and does not have edges from a node to itself (for a social network, this says that if i knows j , then j knows i and that, contrary to Socrates' advice, a person does not know himself or herself). A clique of size k is a set of k nodes where every node has an edge to every other node (i.e., within the clique, everyone knows everyone). An anticlique of size k is a set of k nodes where there are no edges between them (i.e., within the anticlique, no one knows anyone else). For example, the picture below shows a network with nodes labeled 1, 2, ..., 7, where $\{1, 2, 3, 4\}$ is a clique of size 4, and $\{3, 5, 7\}$ is an anticlique of size 3.



(a) Form a random network with n nodes by independently flipping fair coins to decide for each pair $\{x, y\}$ whether there is an edge joining them. Find the expected number of cliques of size k (in terms of n and k).

(b) A triangle is a clique of size 3. For a random network as in (a), find the variance of the

number of triangles (in terms of n).

Hint: Find the covariances of the indicator random variables for each possible clique. There are $\binom{n}{3}$ such indicator r.v.s, some pairs of which are dependent.

*(c) Suppose that $\binom{n}{k} < 2\binom{k}{2} - 1$. Show that there is a network with n nodes containing no cliques of size k or anticliques of size k .

Hint: Explain why it is enough to show that for a random network with n nodes, the probability of the desired property is positive; then consider the complement.

89. (S) Shakespeare wrote a total of 884647 words in his known works. Of course, many words are used more than once, and the number of distinct words in Shakespeare's known writings is 31534 (according to one computation). This puts a lower bound on the size of Shakespeare's vocabulary, but it is likely that Shakespeare knew words which he did not use in these known writings.

More specifically, suppose that a new poem of Shakespeare were uncovered, and consider the following (seemingly impossible) problem: give a good prediction of the number of words in the new poem that do not appear anywhere in Shakespeare's previously known works.

Ronald Thisted and Bradley Efron studied this problem in the papers [8] and [9], developing theory and methods and then applying the methods to try to determine whether Shakespeare was the author of a poem discovered by a Shakespearean scholar in 1985. A simplified version of their method is developed in the problem below. The method was originally invented by Alan Turing (the founder of computer science) and I.J. Good as part of the effort to break the German Enigma code during World War II.

Let N be the number of distinct words that Shakespeare knew, and assume these words are numbered from 1 to N . Suppose for simplicity that Shakespeare wrote only two plays, A and B . The plays are reasonably long and they are of the same length. Let X_j be the number of times that word j appears in play A , and Y_j be the number of times it appears in play B , for $1 \leq j \leq N$.

(a) Explain why it is reasonable to model X_j as being Poisson, and Y_j as being Poisson with the same parameter as X_j .

(b) Let the numbers of occurrences of the word "eyeball" (which was coined by Shakespeare) in the two plays be independent $\text{Pois}(\lambda)$ r.v.s. Show that the probability that "eyeball" is used in play B but not in play A is

$$e^{-\lambda} (\lambda - \lambda^2/2! + \lambda^3/3! - \lambda^4/4! + \dots).$$

(c) Now assume that λ from (b) is unknown and is itself taken to be a random variable to reflect this uncertainty. So let λ have a PDF f_0 . Let X be the number of times the word "eyeball" appears in play A and Y be the corresponding value for play B . Assume that the conditional distribution of X, Y given λ is that they are independent $\text{Pois}(\lambda)$ r.v.s. Show that the probability that "eyeball" is used in play B but not in play A is the alternating series

$$P(X = 1) - P(X = 2) + P(X = 3) - P(X = 4) + \dots$$

Hint: Condition on λ and use Part (b).

(d) Assume that every word's numbers of occurrences in A and B are distributed as in Part (c), where λ may be different for different words but f_0 is fixed. Let W_j be the number of words that appear exactly j times in play A . Show that the expected number of distinct words appearing in play B but not in play A is

$$E(W_1) - E(W_2) + E(W_3) - E(W_4) + \dots$$

(This shows that $W_1 - W_2 + W_3 - W_4 + \dots$ is an unbiased predictor of the number of distinct words appearing in play B but not in play A : on average it is correct. Moreover, it can be

computed just from having seen play A , without needing to know f_0 or any of the λ_j . This method can be extended in various ways to give predictions for unobserved plays based on observed plays.)

4.8 Transformations

The topic for this chapter is transformations of random variables and random vectors. After applying a function to a random variable X or random vector $\mathbf{X}$, the goal is to find the distribution of the transformed random variable or joint distribution of the transformed random vector.

Transformations of random variables appear all over the place in statistics. Here are a few examples, to preview the kinds of transformations we'll be looking at in this chapter.

- Unit conversion: In one dimension, we've already seen how standardization and location-scale transformations can be useful tools for learning about an entire family of distributions. A location-scale change is linear, converting an r.v. X to the r.v. $Y = aX + b$ where a and b are constants (with $a > 0$).

There are also many situations in which we may be interested in nonlinear transformations, e.g., converting from the dollar-yen exchange rate to the yen-dollar exchange rate, or converting information like "Janet's waking hours yesterday consisted of 8 hours of work, 4 hours visiting friends, and 4 hours surfing the web" to the format "Janet was awake for 16 hours yesterday; she spent $\frac{1}{2}$ of that time working, $\frac{1}{4}$ of that time visiting friends, and $\frac{1}{4}$ of that time surfing the web". The change of variables formula, which is the first result in this chapter, shows what happens to the distribution when a random vector is transformed.

- Sums and averages as summaries: It is common in statistics to summarize n observations by their sum or sample average. Turning $X_1, \dots, X_n$ into the sum $T = X_1 + \dots + X_n$ or sample mean $\bar{X}_n = T/n$ is a transformation from $\mathbb{R}^n$ to $\mathbb{R}$.

The term for a sum of independent random variables is convolution. We have already encountered stories and MGFs as two techniques for dealing with convolutions. In this chapter, convolution sums and integrals, which are based on the law of total probability, will give us another way of obtaining the distribution of a sum of r.v.s.

- Extreme values: In many contexts, we may be interested in the distribution of the most extreme observations. For disaster preparedness, government agencies may be concerned about the most extreme flood or earthquake in a 100-year period; in finance, a portfolio manager with an eye toward risk management will want to know the worst 1% or 5% of portfolio returns. In these applications, we are concerned with the maximum or minimum of a set of observations. The transformation that sorts observations, turning $X_1, \dots, X_n$ into the order statistics $\min(X_1, \dots, X_n), \dots, \max(X_1, \dots, X_n)$, is a transformation from $\mathbb{R}^n$ to $\mathbb{R}^n$ that is not invertible. Order statistics are addressed in the last section in this chapter.

Furthermore, it is especially important to us to understand transformations because of the approach we've taken to learning about the named distributions. Starting from a few basic distributions, we have defined other distributions as transformations of these elementary building blocks, in order to understand how the named distributions are related to one another. We'll continue in that spirit here as we introduce two new distributions, the Beta and Gamma, which generalize the Uniform and Exponential.

We already have quite a few tools in our toolbox for dealing with transformations, so let's review those briefly. First, if we are only looking for the expectation of $g(X)$, LOTUS shows us the way: it tells us that the PMF or PDF of X is enough for calculating $E(g(X))$. LOTUS also applies to functions of several r.v.s, as we learned in the previous chapter.

If we need the full distribution of $g(X)$, not just its expectation, our approach depends on whether X is discrete or continuous.

- In the discrete case, we get the PMF of $g(X)$ by translating the event $g(X) = y$ into an equivalent event involving X . To do so, we look for all values x such that $g(x) = y$; as long as X equals any of these x 's, the event $g(X) = y$ will occur. This gives the formula

$$P(g(X) = y) = \sum_{x:g(x)=y} P(X = x).$$

For a one-to-one g , the situation is particularly simple, because there is only one value of x such that $g(x) = y$, namely $g^{-1}(y)$. Then we can use

$$P(g(X) = y) = P(X = g^{-1}(y))$$

to convert between the PMFs of X and $g(X)$, as also discussed in Section 3.7. For example, it is extremely easy to convert between the Geometric and First Success distributions.

- In the continuous case, a universal approach is to start from the CDF of $g(X)$, and translate the event $g(X) \leq y$ into an equivalent event involving X . For general g , we may have to think carefully about how to express $g(X) \leq y$ in terms of X , and there is no easy formula we can plug into. But when g is continuous and strictly increasing, the translation is easy: $g(X) \leq y$ is the same as $X \leq g^{-1}(y)$, so

$$F_{g(X)}(y) = P(g(X) \leq y) = P(X \leq g^{-1}(y)) = F_X(g^{-1}(y)).$$

We can then differentiate with respect to y to get the PDF of $g(X)$. This gives a one-dimensional version of the change of variables formula, which generalizes to invertible transformations in multiple dimensions.

4.8.1 8.1 Change of Variables

Theorem 8.1.1 (Change of variables in one dimension). Let X be a continuous r.v. with PDF f_X , and let $Y = g(X)$, where g is differentiable and strictly increasing (or strictly decreasing). Then the PDF of Y is given by

$$f_Y(y) = f_X(x) \left| \frac{dx}{dy} \right|$$

where $x = g^{-1}(y)$. The support of Y is all $g(x)$ with x in the support of X .

Proof. Let g be strictly increasing. The CDF of Y is

$$F_Y(y) = P(Y \leq y) = P(g(X) \leq y) = P(X \leq g^{-1}(y)) = F_X(g^{-1}(y)) = F_X(x),$$

so by the chain rule, the PDF of Y is

$$f_Y(y) = f_X(x) \frac{dx}{dy}$$

The proof for g strictly decreasing is analogous. In that case the PDF ends up as $-f_X(x)\frac{dx}{dy}$, which is nonnegative since $\frac{dx}{dy} < 0$ if g is strictly decreasing. Using $\left|\frac{dx}{dy}\right|$, as in the statement of the theorem, covers both cases.

When applying the change of variables formula, we can choose whether to compute $\frac{dx}{dy}$, or compute $\frac{dy}{dx}$ and take the reciprocal. By the chain rule, these give the same result, so we can do whichever is easier.

8.1.2. When finding the distribution of Y , be sure to:

- Check the assumptions of the change of variables theorem carefully if you wish to apply it (if it doesn't apply, a good strategy is to start with the CDF of Y).
- Express your final answer for the PDF of Y as a function of y .
- Specify the support of Y .

The change of variables formula (in the strictly increasing g case) is easy to remember when written in the form

$$f_Y(y)dy = f_X(x)dx$$

which has an aesthetically pleasing symmetry to it. This formula also makes sense if we think about units. For example, let X be a measurement in inches and $Y = 2.54X$ be the conversion into centimeters (cm). Then the units of $f_X(x)$ are inches⁻¹ and the units of $f_Y(y)$ are cm⁻¹, so it would be absurd to say something like " $f_Y(y) = f_X(x)$ ". But dx is measured in inches and dy is measured in cm, so $f_Y(y)dy$ and $f_X(x)dx$ are unitless quantities, and it makes sense to equate them. Better yet,

$f_X(x)dx$ and $f_Y(y)dy$ have probability interpretations (recall from Chapter 5 that $f_X(x)dx$ is essentially the probability that X is in a tiny interval of length dx , centered at x), which makes it easier to think intuitively about what the change of variables formula is saying.

The next two examples derive the PDFs of two r.v.s that are defined as transformations of a standard Normal r.v. In the first example the change of variables formula applies; in the second example it does not.

Example 8.1.3 (Log-Normal PDF). Let $X \sim \mathcal{N}(0, 1)$, $Y = e^X$. In Chapter 6 we named the distribution of Y the Log-Normal, and we found all of its moments using the MGF of the Normal distribution. Now we can use the change of variables formula to find the PDF of Y , since $g(x) = e^x$ is strictly increasing. Let $y = e^x$, so $x = \log y$ and $dy/dx = e^x$. Then

$$f_Y(y) = f_X(x) \left| \frac{dx}{dy} \right| = \varphi(x) \frac{1}{e^x} = \varphi(\log y) \frac{1}{y}, \quad y > 0.$$

Note that after applying the change of variables formula, we write everything on the right-hand side in terms of y , and we specify the support of the distribution. To determine the support, we just observe that as x ranges from $-\infty$ to ∞ , e^x ranges from 0 to ∞ .

We can get the same result by working from the definition of the CDF, translating the event $Y \leq y$ into an equivalent event involving X . For $y > 0$,

$$F_Y(y) = P(Y \leq y) = P(e^X \leq y) = P(X \leq \log y) = \Phi(\log y),$$

so the PDF is again

$$f_Y(y) = \frac{d}{dy} \Phi(\log y) = \varphi(\log y) \frac{1}{y}, \quad y > 0$$

Example 8.1.4 (Chi-Square PDF). Let $X \sim \mathcal{N}(0, 1)$, $Y = X^2$. The distribution of Y is an example of a Chi-Square distribution, which is formally introduced in Chapter 10. To find the PDF of Y , we can no longer apply the change of variables formula because $g(x) = x^2$ is not one-to-one; instead we start from the CDF.

By drawing the graph of $y = x^2$, we can see that the event $X^2 \leq y$ is equivalent to the event $-\sqrt{y} \leq X \leq \sqrt{y}$. Then

$$F_Y(y) = P(X^2 \leq y) = P(-\sqrt{y} \leq X \leq \sqrt{y}) = \Phi(\sqrt{y}) - \Phi(-\sqrt{y}) = 2\Phi(\sqrt{y}) - 1,$$

so

$$f_Y(y) = 2\varphi(\sqrt{y}) \cdot \frac{1}{2}y^{-1/2} = \varphi(\sqrt{y})y^{-1/2}, \quad y > 0$$

The following example sheds light on an unexpected appearance of a decidedly non-Normal distribution.

Example 8.1.5 (Lighthouse). A lighthouse on a shore is shining light toward the ocean at a random angle U (measured in radians), where

$$U \sim \text{Unif}\left(\frac{-\pi}{2}, \frac{\pi}{2}\right).$$

Consider a line which is parallel to the shore and 1 mile away from the shore, as illustrated in Figure 8.1. An angle of 0 would mean the ray of light is perpendicular to the shore, while an angle of $\pi/2$ would mean the ray is along the shore, shining to the right from the perspective of the figure.

Let X be the point that the light hits on the line, where the line's origin is the point on the line that is closest to the lighthouse. Find the distribution of X .




FIGURE 8.1 A lighthouse shining light at a random angle U , viewed from above.

Solution: Looking at the right triangle in Figure 8.1, the length of the opposite side of U divided by the length of the adjacent side of U is $X/1 = X$, so

$$X = \tan(U).$$

(The figure illustrates a case where $U > 0$ and, correspondingly, $X > 0$, but the same relationship holds when $U \leq 0$.) Let x be a possible value of X and u be the corresponding possible value of U , so

$$x = \tan(u) \text{ and } u = \arctan(x).$$

By the change of variables formula, which applies since $\tan$ is a differentiable, strictly increasing function on $(-\pi/2, \pi/2)$,

$$f_X(x) = f_U(u) \frac{du}{dx} = \frac{1}{\pi} \cdot \frac{1}{1+x^2},$$

which shows that X is Cauchy. In particular, this implies that $E|X|$ is infinite (since the expected value of a Cauchy does not exist), so on average X is infinitely far from the origin of the line!

The fact that X is Cauchy also makes sense in light of universality of the Uniform. As shown in Example 7.1.25, the Cauchy CDF is

$$F(x) = \frac{1}{\pi} \arctan(x) + 0.5.$$

The inverse is $F^{-1}(v) = \tan(\pi(v - 0.5))$, so for $V \sim \text{Unif}(0, 1)$ we have

$$F^{-1}(V) = \tan(\pi(V - 0.5)) \sim \text{Cauchy}.$$

This agrees with our earlier result since $\pi(V - 0.5) \sim \text{Unif}(-\pi/2, \pi/2)$.

We can also use the change of variables formula to find the PDF of a location-scale transformation.

Example 8.1.6 (PDF of a location-scale transformation). Let X have PDF f_X , and let $Y = a + bX$, with $b \neq 0$. Let $y = a + bx$, to mirror the relationship between Y and X . Then $\frac{dy}{dx} = b$, so the PDF of Y is

$$f_Y(y) = f_X(x) \left| \frac{dx}{dy} \right| = f_X\left(\frac{y-a}{b}\right) \frac{1}{|b|}$$

The change of variables formula generalizes to n dimensions, where it tells us how to use the joint PDF of a random vector $\mathbf{X}$ to get the joint PDF of the transformed random vector $\mathbf{Y} = g(\mathbf{X})$. The formula is analogous to the one-dimensional version, but it involves a multivariate generalization of the derivative called a Jacobian matrix; see sections A. 6 and A. 7 of the math appendix for more about Jacobians.

Theorem 8.1.7 (Change of variables). Let $\mathbf{X} = (X_1, \dots, X_n)$ be a continuous random vector with joint PDF $f_{\mathbf{X}}$. Let $g : A_0 \rightarrow B_0$ be an invertible function, where A_0 and B_0 are open¹ subsets of $\mathbb{R}^n$, A_0 contains the support of $\mathbf{X}$, and B_0 is the range of g .

Let $\mathbf{Y} = g(\mathbf{X})$, and mirror this by letting $\mathbf{y} = g(\mathbf{x})$. Since g is invertible, we also have $\mathbf{X} = g^{-1}(\mathbf{Y})$ and $\mathbf{x} = g^{-1}(\mathbf{y})$.

Suppose that all the partial derivatives $\frac{\partial x_i}{\partial y_j}$ exist and are continuous, so we can form the Jacobian matrix

$$\frac{\partial \mathbf{x}}{\partial \mathbf{y}} = \begin{pmatrix} \frac{\partial x_1}{\partial y_1} & \frac{\partial x_1}{\partial y_2} & \dots & \frac{\partial x_1}{\partial y_n} \\ \vdots & & & \vdots \\ \frac{\partial x_n}{\partial y_1} & \frac{\partial x_n}{\partial y_2} & \dots & \frac{\partial x_n}{\partial y_n} \end{pmatrix}.$$

Also assume that the determinant of this Jacobian matrix is never 0. Then the joint PDF of $\mathbf{Y}$ is

$$f_{\mathbf{Y}}(\mathbf{y}) = f_{\mathbf{X}}(g^{-1}(\mathbf{y})) \cdot \left\| \frac{\partial \mathbf{x}}{\partial \mathbf{y}} \right\| \quad \text{for } \mathbf{y} \in B_0$$

and 0 otherwise. (The inner bars around the Jacobian say to take the determinant and the outer bars say to take the absolute value.)

That is, to convert $f_{\mathbf{X}}(\mathbf{x})$ to $f_{\mathbf{Y}}(\mathbf{y})$ we express the $\mathbf{x}$ in $f_{\mathbf{X}}(\mathbf{x})$ in terms of $\mathbf{y}$ and then multiply by the absolute value of the determinant of the Jacobian $\partial \mathbf{x} / \partial \mathbf{y}$.

⁰¹ A set $C \subset \mathbb{R}^n$ is open if for each $\mathbf{x} \in C$, there exists $\epsilon > 0$ such that all points with distance less than ϵ from $\mathbf{x}$ are contained in C . Sometimes we take $A_0 = B_0 = \mathbb{R}^n$, but often we would like more flexibility for the domain and range of g . For example, if $n = 2$, and X_1 and X_2 have support $(0, \infty)$, we may want to work with the open set $A_0 = (0, \infty) \times (0, \infty)$ rather than all of $\mathbb{R}^2$.

As in the 1D case,

$$\left| \frac{\partial \mathbf{x}}{\partial \mathbf{y}} \right| = \left| \frac{\partial \mathbf{y}}{\partial \mathbf{x}} \right|^{-1}$$

so we can compute whichever of the two Jacobians is easier, and then at the end express the joint PDF of $\mathbf{Y}$ as a function of $\mathbf{y}$.

We will not prove the change of variables formula here, but the idea is to apply the change of variables formula from multivariable calculus and the fact that if A is a region in A_0 and $B = \{g(\mathbf{x}) : \mathbf{x} \in A\}$ is the corresponding region in B_0 , then $\mathbf{X} \in A$ is equivalent to $\mathbf{Y} \in B$ —they are the same event. So $P(\mathbf{X} \in A) = P(\mathbf{Y} \in B)$, which shows that

$$\int_A f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} = \int_B f_{\mathbf{Y}}(\mathbf{y}) d\mathbf{y}$$

The change of variables formula from multivariable calculus (which is reviewed in the math appendix) can then be applied to the integral on the left-hand side, with the substitution $\mathbf{x} = g^{-1}(\mathbf{y})$.

8.1.8. A crucial conceptual difference between transformations of discrete r.v.s and transformations of continuous r.v.s is that with discrete r.v.s we don't need a Jacobian, while with continuous r.v.s we do need a Jacobian. For example, let X be a positive r.v. and $Y = X^3$. If X is discrete, then

$$P(Y = y) = P(X = y^{1/3})$$

converts between the PMFs. But if X is continuous, we need a Jacobian (which in one dimension is just a derivative) to convert between the PDFs:

$$f_Y(y) = f_X(x) \frac{dx}{dy} = f_X(y^{1/3}) \frac{1}{3y^{2/3}}$$

Exercise 23 is a cautionary tale about someone who failed to use a Jacobian when it was needed.

The next two examples apply the 2D change of variables formula.

Example 8.1.9 (Box-Muller). Let $U \sim \text{Unif}(0, 2\pi)$, and let $T \sim \text{Expo}(1)$ be independent of U . Define

$$X = \sqrt{2T} \cos U \text{ and } Y = \sqrt{2T} \sin U$$

Find the joint PDF of (X, Y) . Are they independent? What are their marginal distributions?

Solution:

The joint PDF of U and T is

$$f_{U,T}(u, t) = \frac{1}{2\pi} e^{-t},$$

for $u \in (0, 2\pi)$ and $t > 0$. Viewing (X, Y) as a point in the plane,

$$X^2 + Y^2 = 2T (\cos^2 U + \sin^2 U) = 2T$$

is the squared distance from the origin and U is the angle; that is, $(\sqrt{2T}, U)$ expresses (X, Y) in polar coordinates.

Since we can recover (U, T) from (X, Y) , the transformation is invertible. The Jacobian matrix

$$\frac{\partial(x, y)}{\partial(u, t)} = \begin{pmatrix} -\sqrt{2t} \sin u & \frac{1}{\sqrt{2t}} \cos u \\ \sqrt{2t} \cos u & \frac{1}{\sqrt{2t}} \sin u \end{pmatrix}$$

exists, has continuous entries, and has absolute determinant

$$|-\sin^2 u - \cos^2 u| = 1$$

(which is never 0). Then letting $x = \sqrt{2t} \cos u$, $y = \sqrt{2t} \sin u$ to mirror the transformation from (U, T) to (X, Y) , we have

$$\begin{aligned} f_{X,Y}(x, y) &= f_{U,T}(u, t) \cdot \left\| \frac{\partial(u, t)}{\partial(x, y)} \right\| \\ &= \frac{1}{2\pi} e^{-t} \cdot 1 \\ &= \frac{1}{2\pi} e^{-\frac{1}{2}(x^2+y^2)} \\ &= \frac{1}{\sqrt{2\pi}} e^{-x^2/2} \cdot \frac{1}{\sqrt{2\pi}} e^{-y^2/2}, \end{aligned}$$

for all real x and y .

The joint PDF $f_{X,Y}$ factors into a function of x times a function of y , so X and Y are independent. Furthermore, we recognize the joint PDF as the product of two standard Normal PDFs, so X and Y are i.i.d. $\mathcal{N}(0, 1)$ r.v.s! This result is called the Box-Muller method for generating Normal r.v.s.

Example 8.1.10 (Bivariate Normal joint PDF). In Chapter 7, we saw some properties of the Bivariate Normal distribution and found its joint MGF. Now let's find its joint PDF.

Let (Z, W) be BVN with $\mathcal{N}(0, 1)$ marginals and $\text{Corr}(Z, W) = \rho$. (If we want the joint PDF when the marginals are not standard Normal, we can standardize both components separately and use the result below.) Assume that $-1 < \rho < 1$ since otherwise the distribution is degenerate (with Z and W perfectly correlated).

As shown in Example 7.5.10, we can construct (Z, W) as

$$Z = X$$

$$W = \rho X + \tau Y,$$

with $\tau = \sqrt{1 - \rho^2}$ and X, Y i.i.d. $\mathcal{N}(0, 1)$. We also need the inverse transformation. Solving $Z = X$ for X , we have $X = Z$. Plugging this into $W = \rho X + \tau Y$ and solving for Y , we have

$$X = Z$$

$$Y = -\frac{\rho}{\tau} Z + \frac{1}{\tau} W$$

The Jacobian is

$$\frac{\partial(x, y)}{\partial(z, w)} = \begin{pmatrix} 1 & 0 \\ -\frac{\rho}{\tau} & \frac{1}{\tau} \end{pmatrix}$$

which has absolute determinant $1/\tau$. So by the change of variables formula,

$$\begin{aligned}
f_{Z,W}(z, w) &= f_{X,Y}(x, y) \cdot \left\| \frac{\partial(x, y)}{\partial(z, w)} \right\| \\
&= \frac{1}{2\pi\tau} \exp\left(-\frac{1}{2}(x^2 + y^2)\right) \\
&= \frac{1}{2\pi\tau} \exp\left(-\frac{1}{2}\left(z^2 + \left(-\frac{\rho}{\tau}z + \frac{1}{\tau}w\right)^2\right)\right) \\
&= \frac{1}{2\pi\tau} \exp\left(-\frac{1}{2\tau^2}(z^2 + w^2 - 2\rho zw)\right), \text{ for all real } z, w
\end{aligned}$$

In the last step we multiplied things out and used the fact that $\rho^2 + \tau^2 = 1$. $\square$

4.8.2 8.2 Convolutions

A convolution is a sum of independent random variables. As we mentioned earlier, we often add independent r.v.s because the sum is a useful summary of an experiment (in n Bernoulli trials, we may only care about the total number of successes), and because sums lead to averages, which are also useful (in n Bernoulli trials, the proportion of successes).

The main task in this section is to determine the distribution of $T = X + Y$, where X and Y are independent r.v.s whose distributions are known. In previous chapters, we've already seen how stories and MGFs can help us accomplish this task. For example, we used stories to show that the sum of independent Binomials with the same success probability is Binomial, and that the sum of i.i.d. Geometrics is Negative Binomial. We used MGFs to show that a sum of independent Normals is Normal.

A third method for obtaining the distribution of T is by using a convolution sum or integral. The formulas are given in the following theorem. As we'll see, a convolution sum is nothing more than the law of total probability, conditioning on the value of either X or Y ; a convolution integral is analogous.

Theorem 8.2.1 (Convolution sums and integrals). Let X and Y be independent r.v.s and $T = X + Y$ be their sum. If X and Y are discrete, then the PMF of T is

$$\begin{aligned}
P(T = t) &= \sum_x P(Y = t - x)P(X = x) \\
&= \sum_y P(X = t - y)P(Y = y)
\end{aligned}$$

If X and Y are continuous, then the PDF of T is

$$\begin{aligned}
f_T(t) &= \int_{-\infty}^{\infty} f_Y(t - x)f_X(x)dx \\
&= \int_{-\infty}^{\infty} f_X(t - y)f_Y(y)dy
\end{aligned}$$

Proof. For the discrete case, we use LOTP, conditioning on X :

$$\begin{aligned}
P(T = t) &= \sum_x P(X + Y = t \mid X = x)P(X = x) \\
&= \sum_x P(Y = t - x \mid X = x)P(X = x) \\
&= \sum_x P(Y = t - x)P(X = x)
\end{aligned}$$

Conditioning on Y instead, we obtain the second formula for the PMF of T .

8.2.2. We use the assumption that X and Y are independent in order to get from $P(Y = t - x \mid X = x)$ to $P(Y = t - x)$ in the last step. We are only justified in dropping the condition $X = x$ if the conditional distribution of Y given $X = x$ is the same as the marginal distribution of Y , i.e., X and Y are independent. A common mistake is to assume that after plugging in x for X , we've "already used the information" that $X = x$, when in fact we need an independence assumption to drop the condition. Otherwise we destroy information without justification.

In the continuous case, since the value of a PDF at a point is not a probability, we first find the CDF, and then differentiate to get the PDF. By LOTP,

$$\begin{aligned} F_T(t) &= P(X + Y \leq t) = \int_{-\infty}^{\infty} P(X + Y \leq t \mid X = x) f_X(x) dx \\ &= \int_{-\infty}^{\infty} P(Y \leq t - x) f_X(x) dx \\ &= \int_{-\infty}^{\infty} F_Y(t - x) f_X(x) dx \end{aligned}$$

Again, we need independence to drop the condition $X = x$. To get the PDF, we then differentiate with respect to t , interchanging the order of integration and differentiation. This gives

$$f_T(t) = \int_{-\infty}^{\infty} f_Y(t - x) f_X(x) dx$$

Conditioning on Y instead, we get the second formula for f_T .

An alternative derivation uses the change of variables formula in two dimensions. The only snag is that the change of variables formula requires an invertible transformation from $\mathbb{R}^2$ to $\mathbb{R}^2$, but $(X, Y) \mapsto X + Y$ maps $\mathbb{R}^2$ to $\mathbb{R}$ and is not invertible. We can get around this by adding a redundant component to the transformation, in order to make it invertible. Accordingly, we consider the invertible transformation $(X, Y) \mapsto (X + Y, X)$ (using $(X, Y) \mapsto (X + Y, Y)$ would be equally valid). Once we have the joint PDF of $X + Y$ and X , we integrate out X to get the marginal PDF of $X + Y$.

Let $T = X + Y$, $W = X$, and let $t = x + y$, $w = x$. It may seem redundant to give X the new name " W ", but doing this makes it easier to distinguish between pre-transformation variables and post-transformation variables: we are transforming $(X, Y) \mapsto (T, W)$. Then

$$\frac{\partial(t, w)}{\partial(x, y)} = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$$

has absolute determinant equal to 1, so $\left\| \frac{\partial(x, y)}{\partial(t, w)} \right\|$ is also 1. Thus, the joint PDF of T and W is

$$f_{T, W}(t, w) = f_{X, Y}(x, y) = f_X(x) f_Y(y) = f_X(w) f_Y(t - w)$$

and the marginal PDF of T is

$$f_T(t) = \int_{-\infty}^{\infty} f_{T, W}(t, w) dw = \int_{-\infty}^{\infty} f_X(x) f_Y(t - x) dx$$

in agreement with our result above.

8.2.3. It is not hard to remember the convolution integral formula by reasoning by analogy from

$$P(T = t) = \sum_x P(Y = t - x)P(X = x)$$

to

$$f_T(t) = \int_{-\infty}^{\infty} f_Y(t - x)f_X(x)dx$$

But care is still needed. For example, Exercise 23 shows that an analogous-looking formula for the PDF of the product of two independent continuous r.v.s is wrong: a Jacobian is needed (for convolutions, the absolute Jacobian determinant is 1 so it isn't noticeable in the convolution integral formula).

Since convolution sums are just the law of total probability, we have already used them in previous chapters without mentioning the word convolution; see, for example, the first and most tedious proof of Theorem 3.8.9 (sum of independent Binomials), as well as the proof of Theorem 4.8.1 (sum of independent Poissons). In the following examples, we find the distribution of a sum of Exponentials and a sum of Uniforms using a convolution integral.

Example 8.2.4 (Exponential convolution). Let $X, Y \stackrel{\text{i.i.d.}}{\sim} \text{Expo}(\lambda)$. Find the distribution of $T = X + Y$.

Solution:

For $t > 0$, the convolution formula gives

$$f_T(t) = \int_{-\infty}^{\infty} f_Y(t - x)f_X(x)dx = \int_0^t \lambda e^{-\lambda(t-x)}\lambda e^{-\lambda x}dx$$

where we restricted the integral to be from 0 to t since we need $t - x > 0$ and $x > 0$ for the PDFs inside the integral to be nonzero. Simplifying, we have

$$f_T(t) = \lambda^2 \int_0^t e^{-\lambda t}dx = \lambda^2 t e^{-\lambda t}, \text{ for } t > 0$$

This is known as the Gamma(2, λ) distribution. We will introduce the Gamma distribution in detail in Section 8.4.

Example 8.2.5 (Uniform convolution). Let $X, Y \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(0, 1)$. Find the distribution of $T = X + Y$.

Solution:

The PDF of X (and of Y) is

$$g(x) = \begin{cases} 1, & x \in (0, 1) \\ 0, & \text{otherwise} \end{cases}$$

The convolution formula gives

$$f_T(t) = \int_{-\infty}^{\infty} f_Y(t - x)f_X(x)dx = \int_{-\infty}^{\infty} g(t - x)g(x)dx$$

The integrand is 1 if and only if $0 < t - x < 1$ and $0 < x < 1$; this is a parallelogramshaped constraint. Equivalently, the constraint is $\max(0, t - 1) < x < \min(t, 1)$.




FIGURE 8.2

Region in the (t, x) -plane where $g(t-x)g(x)$ is 1 .

From Figure 8.2, we see that for $0 < t \leq 1$, x is constrained to be in $(0, t)$, and for $1 < t < 2$, x is constrained to be in $(t-1, 1)$. Therefore, the PDF of T is a piecewise linear function:

$$f_T(t) = \begin{cases} \int_0^t dx = t & \text{for } 0 < t \leq 1, \\ \int_{t-1}^1 dx = 2 - t & \text{for } 1 < t < 2. \end{cases}$$

Figure 8.3 plots the PDF of T . It is shaped like a triangle with vertices at 0, 1, and 2, so it is called the Triangle(0, 1, 2) distribution.

Heuristically, it makes sense that T is more likely to take on values near the middle than near the extremes: a value near 1 can be obtained if both X and Y are moderate, if X is large but Y is small, or if Y is large but X is small. In contrast, a value near 2 is only possible if both X and Y are large. Thinking back to Example 3.2.5, the PMF of the sum of two die rolls was also shaped like a triangle. A single die roll has a Discrete Uniform distribution on the integers 1 through 6, so in that problem we were looking at a convolution of two Discrete Uniforms. It makes sense that the PDF we obtained here is similar in shape. $\square$

4.8.3 8.3 Beta

In this section and the next, we will introduce two continuous distributions, the Beta and Gamma, which are related to several named distributions we have already




FIGURE 8.3

PDF of $T = X + Y$, where X and Y are i.i.d. $\text{Unif}(0,1)$.

studied and are also related to each other via a shared story. This is an interlude from the subject of transformations, but we'll eventually need to use a change of variables to tie the Beta and Gamma distributions together.

The Beta distribution is a continuous distribution on the interval $(0, 1)$. It is a generalization of the $\text{Unif}(0, 1)$ distribution, allowing the PDF to be non-constant on $(0, 1)$.

Definition 8.3.1 (Beta distribution). An r.v. X is said to have the Beta distribution with parameters a and b , where $a > 0$ and $b > 0$, if its PDF is

$$f(x) = \frac{1}{\beta(a, b)} x^{a-1} (1-x)^{b-1}, \quad 0 < x < 1,$$

where the constant $\beta(a, b)$ is chosen to make the PDF integrate to 1. We write this as $X \sim \text{Beta}(a, b)$.

Taking $a = b = 1$, the $\text{Beta}(1, 1)$ PDF is constant on $(0, 1)$, so the $\text{Beta}(1, 1)$ and $\text{Unif}(0, 1)$ distributions are the same. By varying the values of a and b , we get PDFs with a variety of shapes; Figure 8.4 shows four examples. Here are a couple of general patterns:

- If $a < 1$ and $b < 1$, the PDF is U-shaped and opens upward. If $a > 1$ and $b > 1$, the PDF opens down.
- If $a = b$, the PDF is symmetric about $1/2$. If $a > b$, the PDF favors values larger than $1/2$; if $a < b$, the PDF favors values smaller than $1/2$.

By definition, the constant $\beta(a, b)$ satisfies

$$\beta(a, b) = \int_0^1 x^{a-1} (1-x)^{b-1} dx$$




FIGURE 8.4

Beta PDFs for various values of a and b . Clockwise from top left: $\text{Beta}(0.5, 0.5)$, $\text{Beta}(2, 1)$, $\text{Beta}(5, 5)$, $\text{Beta}(2, 8)$.

An integral of this form is called a beta integral, and we will need to do some calculus to derive a general formula for $\beta(a, b)$. But in the special case where a and b are positive integers, Thomas Bayes figured out how to do the integral using a story proof rather than calculus!

Story 8.3.2 (Bayes' billiards). Show without using calculus that for any integers k and n with $0 \leq k \leq n$,

$$\int_0^1 \binom{n}{k} x^k (1-x)^{n-k} dx = \frac{1}{n+1}$$

Solution:

By telling two stories, we will show that the left-hand and right-hand sides are both equal to $P(X = k)$, where X is an r.v. that we will construct.

Story 1: Start with $n+1$ balls, n white and 1 gray. Randomly throw each ball onto the unit interval $[0, 1]$, such that the positions of the balls are i.i.d. $\text{Unif}(0, 1)$. Let X be the number of white balls to the left of the gray ball; X is a discrete r.v. with possible values $0, 1, \dots, n$. Figure 8.5 illustrates the setup of our experiment.

To get the probability of the event $X = k$, we use LOTP, conditioning on the position of the gray ball, which we'll call B . Conditional on $B = p$, the number of white balls landing to the left of p has a $\text{Bin}(n, p)$ distribution, since we can consider each of the white balls to be an independent Bernoulli trial, where success is defined as landing to the left of p . Let f be the PDF of B ; $f(p) = 1$ since $B \sim \text{Unif}(0, 1)$. So

$$P(X = k) = \int_0^1 P(X = k \mid B = p) f(p) dp = \int_0^1 \binom{n}{k} p^k (1-p)^{n-k} dp$$




FIGURE 8.5 Bayes' billiards. Here we throw $n = 6$ white balls and one gray ball onto the unit interval, and we observe $X = 2$ balls to the left of the gray ball.

Story 2: Start with $n+1$ balls, all white. Randomly throw each ball onto the unit interval; then choose one ball at random and paint it gray. Again, let X be the number of white balls to the left of the gray ball. By symmetry, any one of the $n+1$ balls is equally likely to be painted gray, so

$$P(X = k) = \frac{1}{n+1}$$

for $k = 0, 1, \dots, n$.

Here's the crux: X has the same distribution in the two stories! It does not matter whether we paint the gray ball first and then throw, or whether we throw first and then paint the gray ball. So $P(X = k)$ is the same in Story 1 and Story 2, and

$$\int_0^1 \binom{n}{k} p^k (1-p)^{n-k} dp = \frac{1}{n+1}$$

for $k = 0, 1, \dots, n$. Despite the k 's in the integrand, the value of the integral doesn't depend on k . Substituting $a-1$ for k and $b-1$ for $n-k$, this shows that for positive integer values of a and b ,

$$\beta(a, b) = \frac{1}{(a+b-1) \binom{a+b-2}{a-1}} = \frac{(a-1)!(b-1)!}{(a+b-1)!}.$$

Later in this chapter, we'll learn what $\beta(a, b)$ is for general a and b . The Beta is a flexible family of distributions on $(0, 1)$, and has many stories. One of these

stories is that a Beta r.v. is often used to represent an unknown probability. That is, we can use the Beta to put probabilities on unknown probabilities!

Story 8.3.3 (Beta-Binomial conjugacy). We have a coin that lands Heads with probability p , but we don't know what p is. Our goal is to infer the value of p after observing the outcomes of n tosses of the coin. The larger that n is, the more accurately we should be able to estimate p .

There are several ways to go about doing this. One major approach is Bayesian inference, which treats all unknown quantities as random variables. In the Bayesian approach, we would treat the unknown probability p as a random variable and give p a distribution. This is called a prior distribution, and it reflects our uncertainty about the true value of p before observing the coin tosses. After the experiment is performed and the data are gathered, the prior distribution is updated using Bayes' rule; this yields the posterior distribution, which reflects our new beliefs about p .

Let's see what happens if the prior distribution on p is a Beta distribution. Let $p \sim \text{Beta}(a, b)$ for known constants a and b , and let X be the number of Heads in n tosses of the coin. Conditional on knowing the true value of p , the tosses would just be independent Bernoulli trials with probability p of success, so

$$X \mid p \sim \text{Bin}(n, p).$$

Note that X is not marginally Binomial; it is conditionally Binomial, given p . The marginal distribution of X is called the Beta-Binomial distribution. To get the posterior distribution of p , we use Bayes' rule (in a hybrid form, since X is discrete and p is continuous). Letting $f(p)$ be the prior distribution and $f(p \mid X = k)$ be the posterior distribution after observing k Heads,

$$\begin{aligned} f(p \mid X = k) &= \frac{P(X = k \mid p)f(p)}{P(X = k)} \\ &= \frac{\binom{n}{k} p^k (1-p)^{n-k} \cdot \frac{1}{\beta(a,b)} p^{a-1} (1-p)^{b-1}}{P(X = k)} \end{aligned}$$

The denominator, which is the marginal PMF of X , is given by

$$P(X = k) = \int_0^1 P(X = k \mid p) f(p) dp = \int_0^1 \binom{n}{k} p^k (1-p)^{n-k} f(p) dp$$

For $a = b = 1$ (which gives a $\text{Unif}(0, 1)$ prior on p), we showed in the Bayes' billiards story that $P(X = k) = 1/(n+1)$, i.e., X is Discrete Uniform on $\{0, 1, \dots, n\}$. For a and b any positive integers, we can again use Bayes' billiards to find $P(X = k)$; interestingly, X then turns out to have a Negative Hypergeometric distribution (the story and PMF of which we saw in Example 4.4.7).

But it does not seem easy to find $P(X = k)$ in general for real a, b . Relatedly, we still have not evaluated $\beta(a, b)$ in general. Are we stuck?

Actually, the calculation is much easier than it appears at first. The conditional PDF $f(p \mid X = k)$ is a function of p , which means everything that doesn't depend on p is just a constant. We can drop all these constants and find the PDF up to a multiplicative constant (and then the normalizing constant is whatever it needs to be to make the PDF integrate to 1). This gives

$$f(p \mid X = k) \propto p^{a+k-1} (1-p)^{b+n-k-1}$$

which is the $\text{Beta}(a+k, b+n-k)$ PDF, up to a multiplicative constant. Therefore, the posterior distribution of p is

$$p \mid X = k \sim \text{Beta}(a + k, b + n - k)$$

The posterior distribution of p after observing $X = k$ is still a Beta distribution! This is a special relationship between the Beta and Binomial distributions called conjugacy: if we have a Beta prior distribution on p and data that are conditionally Binomial given p , then when going from prior to posterior, we don't leave the family of Beta distributions. We say that the Beta is the conjugate prior of the Binomial.

Furthermore, notice the very simple formula for updating the distribution of p . We just add the number of observed successes, k , to the first parameter of the Beta distribution, and the number of observed failures, $n - k$, to the second parameter. So a and b have a concrete interpretation in this context, at least when a and b are positive integers: think of $a - 1$ as the number of prior successes and $b - 1$ as the number of prior failures in earlier experiments (these prior experiments could be real or imagined). Adding on k successes and $n - k$ failures from the current

experiment, we have $a + k - 1$ successes and $b + n - k - 1$ failures, and the $\text{Beta}(a, b)$ prior gets updated to a $\text{Beta}(a + k, b + n - k)$ posterior.

As in Section 2.6, we can sequentially update our beliefs as we get more and more evidence: we start with a prior distribution and update it to get a posterior distribution, which becomes the new prior distribution, which we update to get a new posterior distribution, etc. The beauty here is that all of this can be done within the Beta family of distributions, with easy updates to the parameters based on tallying the observed successes and failures.

For concreteness, Figure 8.6 shows the case where the prior is $\text{Beta}(1, 1)$ (which is equivalent to $\text{Unif}(0, 1)$, as noted earlier), and we observe $n = 5$ coin tosses, all of which happen to land Heads. Then the posterior is $\text{Beta}(6, 1)$, which is plotted on the right half of Figure 8.6. Notice how the posterior distribution incorporates the evidence from the coin tosses: larger values of p have higher density, consistent with the fact that we observed all Heads.




FIGURE 8.6

Beta is conjugate prior of Binomial. Left: prior is $\text{Unif}(0, 1)$. Right: after observing 5 Heads in 5 tosses, posterior is $\text{Beta}(6, 1)$.

This model is a continuous analog of Example 2.3.7, our very first example of Bayes' rule. In that example, we also had a coin whose probability of Heads p was unknown, but our prior information led us to believe that p could only take on one of two possible values, $1/2$ or $3/4$. For this reason, our prior distribution on p —though we didn't call it that at the time!—was discrete. In particular, our prior PMF was

$$P(p = 1/2) = 1/2$$

$$P(p = 3/4) = 1/2$$

After observing three Heads in a row, we updated this PMF to obtain the posterior PMF, which assigned a probability of 0.23 to $p = 1/2$ and a probability of 0.77

to $p = 3/4$. The same logic applies to the example in this chapter, except that we now give p a continuous prior distribution, which is appropriate if we believe that p could possibly take on any value between 0 and 1.

To conclude the section, here are two examples that illustrate how Beta-Binomial conjugacy and Bayes' billiards can make some seemingly complicated problems much easier.

Example 8.3.4 (Bayes' serum). A new treatment, Bayes' serum, has just been developed for the disease conditionitis. A clinical trial is about to be conducted, to study how effective the treatment is. Bayes' serum will be applied to n patients who have conditionitis. Given p , the patients' outcomes are independent, with each patient having probability p of being cured by the treatment. But p is unknown. To quantify our uncertainty about p , we model p as a random variable, with prior distribution $p \sim \text{Unif}(0, 1)$.

(a) Find the probability that exactly k out of the n patients will be cured by the treatment (unconditionally, not given p).

(b) Now suppose that the treatment is extremely effective in the clinical trial: all n patients are cured! Given this information, find the probability that p exceeds $1/2$. Your answer should be fully simplified, and expressed only in terms of n .

Solution:

(a) By LOTP followed by Bayes' billiards or pattern-matching to a Beta PDF, the probability is

$$\int_0^1 \binom{n}{k} p^k (1-p)^{n-k} dp = \frac{1}{n+1}$$

(b) By Beta-Binomial conjugacy, the posterior distribution of p given the data is $\text{Beta}(1+n, 1)$, which has PDF $(n+1)p^n$ for $0 < p < 1$. Thus,

$$P\left(p > \frac{1}{2}\right) = \int_{1/2}^1 (n+1)p^n dp = p^{n+1} \Big|_{1/2}^1 = 1 - \frac{1}{2^{n+1}}$$

Example 8.3.5. A basketball player will shoot $N \sim \text{Pois}(\lambda)$ free throws in a game tomorrow. Let X_j be the indicator of her making their j th free throw, and $X = X_1 + \cdots + X_N$ be the total number of free throws she makes in the game (so $X = 0$ if $N = 0$). To model our uncertainty about how good a free throw shooter she is, let $p \sim \text{Beta}(a, b)$. Given p , the player has probability p of making a free throw and probability $q = 1 - p$ of missing it. Assume that $X_1, X_2, \dots$ are conditionally independent given p , and that N is independent of $p, X_1, X_2, \dots$.

(a) Find the conditional distribution of X given N, p .

(b) Find the conditional distribution of X given p .

(c) Find the conditional distribution of X given N , for the case $a = b = 1$.

(d) Find the conditional distribution of p given X, N .

Solution:

(a) By the story of the Binomial, $X \mid (N, p) \sim \text{Bin}(N, p)$.

(b) By the chicken-egg story, $X \mid p \sim \text{Pois}(\lambda p)$.

(c) For $k = 0, 1, \dots, n$, by LOTP and Bayes' billiards we have

$$P(X = k \mid N = n) = \int_0^1 \binom{n}{k} p^k (1-p)^{n-k} dp = \frac{1}{n+1}$$

So the conditional distribution of X given N is Discrete Uniform on $\{0, 1, \dots, N\}$.

(d) By Beta-Binomial conjugacy, $p \mid (X, N) \sim \text{Beta}(a + X, b + N - X)$.

4.8.4 8.4 Gamma

The Gamma distribution is a continuous distribution on the positive real line; it is a generalization of the Exponential distribution. While an Exponential r.v. represents the waiting time for the first success under conditions of memorylessness, we shall see that a Gamma r.v. represents the total waiting time for multiple successes.

Before writing down the PDF, we first introduce the gamma function, a very famous function in mathematics that extends the factorial function beyond the realm of nonnegative integers.

Definition 8.4.1 (Gamma function). The gamma function Γ is defined by

$$\Gamma(a) = \int_0^{\infty} x^a e^{-x} \frac{dx}{x}$$

for real numbers $a > 0$.

We could also cancel an x and write the integrand as $x^{a-1}e^{-x}$, but it turns out to be convenient having the $\frac{dx}{x}$ since it is common to make a transformation of the form $u = cx$, and then we have the handy fact that $\frac{du}{u} = \frac{dx}{x}$. Here are two important properties of the gamma function.

- $\Gamma(a+1) = a\Gamma(a)$ for all $a > 0$. This follows from integration by parts:

$$\Gamma(a+1) = \int_0^{\infty} x^a e^{-x} dx = -x^a e^{-x} \Big|_0^{\infty} + a \int_0^{\infty} x^{a-1} e^{-x} dx = 0 + a\Gamma(a)$$

- $\Gamma(n) = (n-1)!$ if n is a positive integer. This can be proved by induction, starting with $n = 1$ and using the recursive relation $\Gamma(a+1) = a\Gamma(a)$. Thus, if we evaluate the gamma function at positive integer values, we recover the factorial function (albeit shifted by 1).

Now let's suppose that on a whim, we decide to divide both sides of the above definition by $\Gamma(a)$. We have

$$1 = \int_0^{\infty} \frac{1}{\Gamma(a)} x^a e^{-x} \frac{dx}{x}$$

so the function under the integral is a valid PDF supported on $(0, \infty)$. This is the definition of the PDF of the Gamma distribution. Specifically, we say that X has the Gamma distribution with parameters a and 1, denoted $X \sim \text{Gamma}(a, 1)$, if its PDF is

$$f_X(x) = \frac{1}{\Gamma(a)} x^a e^{-x} \frac{1}{x}, \quad x > 0$$

From the $\text{Gamma}(a, 1)$ distribution, we obtain the general Gamma distribution by a scale transformation: if $X \sim \text{Gamma}(a, 1)$ and $\lambda > 0$, then the distribution of $Y = X/\lambda$ is called the $\text{Gamma}(a, \lambda)$ distribution. By the change of variables formula with $x = \lambda y$ and $dx/dy = \lambda$, the PDF of Y is

$$f_Y(y) = f_X(x) \left| \frac{dx}{dy} \right| = \frac{1}{\Gamma(a)} (\lambda y)^a e^{-\lambda y} \frac{1}{\lambda y} \lambda = \frac{1}{\Gamma(a)} (\lambda y)^a e^{-\lambda y} \frac{1}{y}, \quad y > 0$$

This is summarized in the following definition.

Definition 8.4.2 (Gamma distribution). An r.v. Y is said to have the Gamma distribution with parameters a and λ , where $a > 0$ and $\lambda > 0$, if its PDF is

$$f(y) = \frac{1}{\Gamma(a)} (\lambda y)^a e^{-\lambda y} \frac{1}{y}, \quad y > 0$$

We write $Y \sim \text{Gamma}(a, \lambda)$.

Taking $a = 1$, the $\text{Gamma}(1, \lambda)$ PDF is $f(y) = \lambda e^{-\lambda y}$ for $y > 0$, so the $\text{Gamma}(1, \lambda)$ and $\text{Expo}(\lambda)$ distributions are the same. The extra parameter a allows Gamma PDFs to have a greater variety of shapes. Figure 8.7 shows four Gamma PDFs. For small values of a , the PDF is skewed, but as a increases, the PDF starts to look more symmetrical and bell-shaped; we will learn the reason for this in Chapter 10. Increasing λ compresses the PDF toward smaller values, as we can see by comparing the $\text{Gamma}(3, 1)$ and $\text{Gamma}(3, 0.5)$ PDFs.

Let's find the mean, variance, and other moments of the Gamma distribution, starting with $X \sim \text{Gamma}(a, 1)$. We'll use properties of the gamma function as well as the technique of doing integrals by pattern recognition. For the mean, we write down the definition of $E(X)$,

$$E(X) = \int_0^{\infty} \frac{1}{\Gamma(a)} x^{a+1} e^{-x} \frac{dx}{x}$$

but instead of attempting a gruesome integration by parts, we recognize that after taking out $1/\Gamma(a)$, what's left is precisely the gamma function evaluated at $a + 1$. Therefore

$$E(X) = \frac{\Gamma(a+1)}{\Gamma(a)} = \frac{a\Gamma(a)}{\Gamma(a)} = a.$$




FIGURE 8.7

Gamma PDFs for various values of a and λ . Clockwise from top left: $\text{Gamma}(3, 1)$, $\text{Gamma}(3, 0.5)$, $\text{Gamma}(5, 0.5)$, $\text{Gamma}(10, 1)$.

Similarly, for the variance, LOTUS gives us an integral expression for the second moment,

$$E(X^2) = \int_0^{\infty} \frac{1}{\Gamma(a)} x^{a+2} e^{-x} \frac{dx}{x}$$

and we recognize the gamma function evaluated at $a + 2$. Therefore

$$E(X^2) = \frac{\Gamma(a+2)}{\Gamma(a)} = \frac{(a+1)a\Gamma(a)}{\Gamma(a)} = (a+1)a$$

and

$$\text{Var}(X) = (a+1)a - a^2 = a.$$

So for $X \sim \text{Gamma}(a, 1)$, $E(X) = \text{Var}(X) = a$.

The c th moment is no harder than the second moment; we just use LOTUS and recognize the definition of $\Gamma(a+c)$. This gives

$$E(X^c) = \int_0^\infty \frac{1}{\Gamma(a)} x^{a+c} e^{-x} \frac{dx}{x} = \frac{\Gamma(a+c)}{\Gamma(a)}$$

for all real c such that the integral converges, i.e., for $c > -a$.

We can now transform to $Y = X/\lambda \sim \text{Gamma}(a, \lambda)$ to get

$$\begin{aligned} E(Y) &= \frac{1}{\lambda} E(X) = \frac{a}{\lambda}, \\ \text{Var}(Y) &= \frac{1}{\lambda^2} \text{Var}(X) = \frac{a}{\lambda^2}, \\ E(Y^c) &= \frac{1}{\lambda^c} E(X^c) = \frac{1}{\lambda^c} \cdot \frac{\Gamma(a+c)}{\Gamma(a)}, \quad c > -a. \end{aligned}$$

Looking back at the Gamma PDF plots, they are consistent with our finding that the mean and variance are increasing in a and decreasing in λ .

So far, we've been learning about the Gamma distribution using the PDF, which allowed us to discern general patterns from PDF plots and to derive the mean and variance. But the PDF doesn't provide much insight about why we'd ever use the Gamma distribution, and it doesn't give us much of an interpretation for the parameters a and λ . For this, we need to connect the Gamma to other named distributions through stories. The rest of this section is devoted to stories for the Gamma distribution.

In the special case where a is an integer, we can represent a $\text{Gamma}(a, \lambda)$ r.v. as a sum (convolution) of i.i.d. $\text{Expo}(\lambda)$ r.v.s.

Theorem 8.4.3. Let $X_1, \dots, X_n$ be i.i.d. $\text{Expo}(\lambda)$. Then

$$X_1 + \dots + X_n \sim \text{Gamma}(n, \lambda).$$

Proof. The $\text{Expo}(\lambda)$ MGF is $\frac{\lambda}{\lambda-t}$ for $t < \lambda$, so the MGF of $X_1 + \dots + X_n$ is

$$M_n(t) = \left(\frac{\lambda}{\lambda-t} \right)^n$$

for $t < \lambda$. Let $Y \sim \text{Gamma}(n, \lambda)$; we'll show that the MGF of Y is the same as that of $X_1 + \dots + X_n$. By LOTUS,

$$E(e^{tY}) = \int_0^\infty e^{ty} \frac{1}{\Gamma(n)} (\lambda y)^n e^{-\lambda y} \frac{dy}{y}$$

Again, we'll use pattern recognition to do the integral. We just need to do algebraic manipulations until what's left inside the integral is a recognizable Gamma PDF:

$$\begin{aligned} E(e^{tY}) &= \int_0^\infty e^{ty} \frac{1}{\Gamma(n)} (\lambda y)^n e^{-\lambda y} \frac{dy}{y} \\ &= \frac{\lambda^n}{(\lambda-t)^n} \int_0^\infty \frac{1}{\Gamma(n)} e^{-(\lambda-t)y} ((\lambda-t)y)^n \frac{dy}{y} \end{aligned}$$

We pulled λ^n out of the integral, then multiplied by $(\lambda-t)^n$ on the inside while dividing by it on the outside. Now the expression inside the integral is the Gamma $(n, \lambda-t)$ PDF, assuming $t < \lambda$. Since PDFs integrate to 1, we have

$$E(e^{tY}) = \left(\frac{\lambda}{\lambda - t} \right)^n$$

for $t < \lambda$; if $t \geq \lambda$ the integral fails to converge.

We have shown that $X_1 + \cdots + X_n$ and $Y \sim \text{Gamma}(n, \lambda)$ have the same MGF. Since the MGF determines the distribution, $X_1 + \cdots + X_n \sim \text{Gamma}(n, \lambda)$.

Alternatively, we can compute the convolution integral inductively. It suffices to consider the case $\lambda = 1$, since after having done that case we can rescale the r.v.s to prove the result for general λ . Let $T_n = X_1 + \cdots + X_n$. We will show by induction that $T_n \sim \text{Gamma}(n, 1)$ for all $n \geq 1$. For $n = 1$, this is true since the $\text{Gamma}(1, 1)$ PDF is e^{-x} (for $x > 0$), which is the same as the $\text{Expo}(1)$ PDF. Now assume that $T_n \sim \text{Gamma}(n, 1)$, and show that $T_{n+1} \sim \text{Gamma}(n+1, 1)$. The PDF of $T_{n+1} = T_n + X_{n+1}$ is given by the convolution integral $\int_0^\infty f_{T_n}(x)f_{X_{n+1}}(t-x)dx = \frac{1}{\Gamma(n)} \int_0^t x^{n-1}e^{-x}e^{-(t-x)}dx = \frac{e^{-t}}{\Gamma(n)} \int_0^t x^{n-1}dx = \frac{t^n e^{-t}}{\Gamma(n+1)}$ for $t > 0$, which completes the induction.

Thus, if $Y \sim \text{Gamma}(a, \lambda)$ with a an integer, we can represent Y as a sum of i.i.d. $\text{Expo}(\lambda)$ r.v.s, $X_1 + \cdots + X_a$, and get the mean and variance right away:

$$E(Y) = E(X_1 + \cdots + X_a) = aE(X_1) = \frac{a}{\lambda}$$

$$\text{Var}(Y) = \text{Var}(X_1 + \cdots + X_a) = a \text{Var}(X_1) = \frac{a}{\lambda^2}$$

in agreement with the results we derived earlier for general a .

Theorem 8.4.3 also allows us to connect the Gamma distribution to the story of the Poisson process. We showed in Chapter 5 that in a Poisson process of rate λ , the interarrival times are i.i.d. $\text{Expo}(\lambda)$ r.v.s. But the total waiting time T_n for the n th arrival is the sum of the first n interarrival times; for instance, Figure 8.8 illustrates how T_3 is the sum of the 3 interarrival times X_1, X_2, X_3 . Therefore, by the theorem, $T_n \sim \text{Gamma}(n, \lambda)$. The interarrival times in a Poisson process are Exponential r.v.s, while the raw arrival times are Gamma r.v.s.




FIGURE 8.8 Poisson process. The interarrival times X_j are i.i.d. $\text{Expo}(\lambda)$, while the raw arrival times T_j are $\text{Gamma}(j, \lambda)$.

8.4.4. Unlike the X_j , the T_j are not independent, since they are constrained to be increasing; nor are they identically distributed.

At last, we have an interpretation for the parameters of the $\text{Gamma}(a, \lambda)$ distribution. In the Poisson process story, a is the number of successes we are waiting for, and λ is the rate at which successes arrive; $Y \sim \text{Gamma}(a, \lambda)$ is the total waiting time for the a th arrival in a Poisson process of rate λ .

A consequence of this story is that a convolution of Gammas with the same λ is still Gamma. Exercise 29 explores this fact from several perspectives.

When we introduced the Exponential distribution, we viewed it as the continuous analog of the Geometric distribution: the Geometric waits for the first success in discrete time, and the Exponential waits for the first success in continuous time. Likewise, we can now say that the Gamma distribution is the continuous analog of the Negative Binomial distribution: the Negative Binomial is a sum of Geometric waiting times, and the Gamma is a sum of Exponential waiting times. In Exercise 54 you will use MGFs to show that the Gamma distribution can be obtained as a continuous limit of the Negative Binomial distribution.

A final story about the Gamma is that it shares the same special relationship with the Poisson that the Beta shares with the Binomial: the Gamma is the conjugate prior of the Poisson. Earlier we saw that the Beta distribution can represent an unknown probability of success because its support is $(0, 1)$. The Gamma distribution, on the other hand, can represent an unknown rate in a Poisson process because its support is $(0, \infty)$.

To investigate, we'll return to Blotchville, where buses arrive in a Poisson process of rate λ . Previously it was assumed that $\lambda = 1/10$, so that the times between buses were i.i.d. Exponentials with mean 10 minutes, but now we'll assume Fred doesn't know the rate λ at which buses arrive and needs to figure it out. Fred will follow the Bayesian approach and treat the unknown rate as a random variable.

Story 8.4.5 (Gamma-Poisson conjugacy). In Blotchville, buses arrive at a certain bus stop according to a Poisson process with rate λ buses per hour, where λ is unknown. Based on his very memorable adventures in Blotchville, Fred quantifies his uncertainty about λ using the prior $\lambda \sim \text{Gamma}(r_0, b_0)$, where r_0 and b_0 are known, positive constants with r_0 an integer.

To better understand the bus system, Fred wants to learn more about λ . He is a very patient person, and decides that he will sit at the bus stop for t hours and count how many buses arrive in this time interval. Let Y be the number of buses in this time interval, and suppose Fred observes that $Y = y$.

- Find Fred's hybrid joint distribution for Y and λ .
- Find Fred's marginal distribution for Y .
- Find Fred's posterior distribution for λ , i.e., his conditional distribution of λ given the data y .
- Find Fred's posterior mean $E(\lambda \mid Y = y)$ and posterior variance $\text{Var}(\lambda \mid Y = y)$.

Solution: Notice the similarities between the structure of this problem and that of the hybrid joint distribution from Example 7.1.26. We know that $\lambda \sim \text{Gamma}(r_0, b_0)$ marginally, and by definition of Poisson process, conditional on knowing the true rate λ , the number of buses in an interval of length t is distributed $\text{Pois}(\lambda t)$. In other words, what we're given is


$$\begin{aligned}\lambda &\sim \text{Gamma}(r_0, b_0) \\ Y \mid \lambda &\sim \text{Pois}(\lambda t).\end{aligned}$$

Then we are asked to flip it around: find the marginal distribution of Y and the conditional distribution of λ given $Y = y$, which is the posterior distribution. This is characteristic of Bayesian inference: we have a prior distribution for the unknown parameters (in this case, a Gamma distribution for λ) and a model for the data conditional on the unknown parameters (in this case, a Poisson distribution for Y given λ), and we use Bayes' rule to get the distribution of the unknowns conditional on the observed data. So let's get started.

- Let f_0 be the prior PDF of λ . The hybrid joint distribution of Y and λ is

$$f(y, \lambda) = P(Y = y \mid \lambda) f_0(\lambda) = \frac{e^{-\lambda t} (\lambda t)^y}{y!} \frac{(b_0 \lambda)^{r_0} e^{-b_0 \lambda}}{\lambda \Gamma(r_0)},$$

for $y = 0, 1, 2, \dots$ and $\lambda > 0$. The hybrid joint distribution is plotted in Figure 8.9;






FIGURE 8.9

Hybrid joint distribution of Y and λ . (a) Conditioning on a particular value of λ , the relative heights form a Poisson PMF. (b) In the other direction, conditioning on $Y = y$ gives the posterior distribution of λ .

there is a conditional PMF of Y for every value of λ and a conditional PDF of λ for every value of Y .

(b) To get the marginal PMF of Y , we integrate out λ from the hybrid joint distribution; this is also a form of LOTP. This gives

$$\begin{aligned} P(Y = y) &= \int_0^\infty P(Y = y \mid \lambda) f_0(\lambda) d\lambda \\ &= \int_0^\infty \frac{e^{-\lambda t} (\lambda t)^y}{y!} \frac{(b_0 \lambda)^{r_0} e^{-b_0 \lambda}}{\Gamma(r_0)} \frac{d\lambda}{\lambda} \end{aligned}$$

Let's do the integral by pattern recognition, focusing in on the terms involving λ . We spot λ^{r_0+y} and $e^{-(b_0+t)\lambda}$ lurking in the integrand, which suggests patternmatching to a $\text{Gamma}(r_0 + y, b_0 + t)$ PDF. Pull out all the terms that don't depend on λ , then multiply by whatever it takes to get the desired PDF inside the integral, remembering to multiply by the reciprocal on the outside:

$$\begin{aligned} P(Y = y) &= \frac{t^y b_0^{r_0}}{y! \Gamma(r_0)} \int_0^\infty e^{-(b_0+t)\lambda} \lambda^{r_0+y} \frac{d\lambda}{\lambda} \\ &= \frac{\Gamma(r_0 + y)}{y! \Gamma(r_0)} \frac{t^y b_0^{r_0}}{(b_0 + t)^{r_0+y}} \int_0^\infty \frac{1}{\Gamma(r_0 + y)} e^{-(b_0+t)\lambda} ((b_0 + t) \lambda)^{r_0+y} \frac{d\lambda}{\lambda} \\ &= \frac{(r_0 + y - 1)!}{(r_0 - 1)! y!} \left(\frac{t}{b_0 + t} \right)^y \left(\frac{b_0}{b_0 + t} \right)^{r_0} \end{aligned}$$

In the last step, we used the property $\Gamma(n) = (n - 1)!$, which is applicable because r_0 is an integer. This is the $\text{NBin}(r_0, b_0 / (b_0 + t))$ PMF, so the marginal distribution of Y is Negative Binomial with parameters r_0 and $b_0 / (b_0 + t)$.

(c) By Bayes' rule, the posterior PDF of λ is given by

$$f_1(\lambda \mid y) = \frac{P(Y = y \mid \lambda) f_0(\lambda)}{P(Y = y)}$$

We found $P(Y = y)$ in the previous part, but since it does not depend on λ , we can just treat it as part of the normalizing constant. Absorbing this and other multiplicative factors that don't depend on λ into the normalizing constant,

$$f_1(\lambda \mid y) \propto e^{-\lambda t} \lambda^y \lambda^{r_0} e^{-b_0 \lambda} \frac{1}{\lambda} = e^{-(b_0+t)\lambda} \lambda^{r_0+y} \frac{1}{\lambda},$$

which shows that the posterior distribution of λ is $\text{Gamma}(r_0 + y, b_0 + t)$.

When going from prior to posterior, the distribution of λ stays in the Gamma family, so the Gamma is indeed the conjugate prior for the Poisson.

Now that we have the posterior PDF of λ , we have a more elegant approach to solving (b). Rearranging Bayes' rule, the marginal PMF of Y is

$$P(Y = y) = \frac{P(Y = y | \lambda)f_0(\lambda)}{f_1(\lambda | y)}$$

where we know the numerator from the statement of the problem and the denominator from the calculation we just did. Plugging in these ingredients and simplifying again yields

$$Y \sim \text{NBin}(r_0, b_0 / (b_0 + t)).$$

(d) Since conditional PDFs are PDFs, it is perfectly fine to calculate the expectation and variance of λ with respect to the posterior distribution. The mean and variance of the $\text{Gamma}(r_0 + y, b_0 + t)$ distribution give us

$$E(\lambda | Y = y) = \frac{r_0 + y}{b_0 + t} \quad \text{and} \quad \text{Var}(\lambda | Y = y) = \frac{r_0 + y}{(b_0 + t)^2}.$$

This example gives another interpretation for the parameters in a Gamma when it is being used as a conjugate prior. Fred's $\text{Gamma}(r_0, b_0)$ prior got updated to a $\text{Gamma}(r_0 + y, b_0 + t)$ posterior after observing y arrivals in t hours. We can imagine that in the past, Fred observed r_0 buses arrive in b_0 hours; then after the new data, he has observed $r_0 + y$ buses in $b_0 + t$ hours. So we can interpret r_0 as the number of prior arrivals and b_0 as the total time required for those prior arrivals.

4.8.5 8.5 Beta-gamma Connections

In this section, we will unite the Beta and Gamma distributions with a common story. As an added bonus, the story will give us an expression for the normalizing constant of the $\text{Beta}(a, b)$ PDF in terms of gamma functions, and it will allow us to easily find the expectation of the $\text{Beta}(a, b)$ distribution.

Story 8.5.1 (Bank-post office). While running errands, you need to go to the bank, then to the post office. Let $X \sim \text{Gamma}(a, \lambda)$ be your waiting time in line at the bank, and let $Y \sim \text{Gamma}(b, \lambda)$ be your waiting time in line at the post office (with the same λ for both). Assume X and Y are independent. What is the joint distribution of $T = X + Y$ (your total wait at the bank and post office) and $W = \frac{X}{X+Y}$ (the fraction of your waiting time spent at the bank)?

Solution:

We'll do a change of variables in two dimensions to get the joint PDF of T and W . Let $t = x + y, w = \frac{x}{x+y}$. Then $x = tw, y = t(1 - w)$, and

$$\frac{\partial(x, y)}{\partial(t, w)} = \begin{pmatrix} w & t \\ 1 - w & -t \end{pmatrix},$$

which has an absolute determinant of t . Therefore

$$\begin{aligned}
f_{T,W}(t, w) &= f_{X,Y}(x, y) \cdot \left\| \frac{\partial(x, y)}{\partial(t, w)} \right\| \\
&= f_X(x) f_Y(y) \cdot t \\
&= \frac{1}{\Gamma(a)} (\lambda x)^a e^{-\lambda x} \frac{1}{x} \cdot \frac{1}{\Gamma(b)} (\lambda y)^b e^{-\lambda y} \frac{1}{y} \cdot t \\
&= \frac{1}{\Gamma(a)} (\lambda t w)^a e^{-\lambda t w} \frac{1}{t w} \cdot \frac{1}{\Gamma(b)} (\lambda t (1-w))^b e^{-\lambda t (1-w)} \frac{1}{t(1-w)} \cdot t
\end{aligned}$$

Let's group all the terms involving w together, and all the terms involving t together:

$$\begin{aligned}
f_{T,W}(t, w) &= \frac{1}{\Gamma(a)\Gamma(b)} w^{a-1} (1-w)^{b-1} (\lambda t)^{a+b} e^{-\lambda t} \frac{1}{t} \\
&= \left(\frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} w^{a-1} (1-w)^{b-1} \right) \left(\frac{1}{\Gamma(a+b)} (\lambda t)^{a+b} e^{-\lambda t} \frac{1}{t} \right)
\end{aligned}$$

for $0 < w < 1$ and $t > 0$. The form of the joint PDF, together with Proposition 7.1.21, tells us several things:

1. Since the joint PDF factors into a function of t times a function of w , we have that T and W are independent: the total waiting time is independent of the fraction of time spent at the bank.
2. We recognize the marginal PDF of T and deduce that $T \sim \text{Gamma}(a+b, \lambda)$.
3. The PDF of W is

$$f_W(w) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} w^{a-1} (1-w)^{b-1}, \quad 0 < w < 1$$

by Proposition 7.1.21 or just by integrating out T from the joint PDF of T and W . This PDF is proportional to the $\text{Beta}(a, b)$ PDF, so it is the $\text{Beta}(a, b)$ PDF ! Note that as a byproduct of the calculation we have just done, we have found the normalizing constant of the Beta distribution:

$$\frac{1}{\beta(a, b)} = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)}$$

is the constant that goes in front of the $\text{Beta}(a, b)$ PDF.

To summarize, the bank-post office story tells us that when we add independent Gamma r.v.s X and Y with the same rate λ , the total $X+Y$ has a Gamma distribution, the fraction $X/(X+Y)$ has a Beta distribution, and the total is independent of the fraction.

We can use this result to find the mean of $W \sim \text{Beta}(a, b)$ without the slightest trace of calculus.

Example 8.5.2 (Beta expectation). With notation as above, note that since T and W are independent, they are uncorrelated: $E(TW) = E(T)E(W)$. Writing this in terms of X and Y , we have

$$\begin{aligned}
E\left((X+Y) \cdot \frac{X}{X+Y}\right) &= E(X+Y)E\left(\frac{X}{X+Y}\right), \\
E(X) &= E(X+Y)E\left(\frac{X}{X+Y}\right), \\
\frac{E(X)}{E(X+Y)} &= E\left(\frac{X}{X+Y}\right).
\end{aligned}$$

Ordinarily, the last equality would be a horrendous blunder: faced with an expectation like $E(X/(X+Y))$, we are not generally permitted to move the E into the numerator and denominator as we please. In this case, however, the bank-post office story justifies the move, so finding the expectation of W happily reduces to finding the expectations of X and $X+Y$:

$$E(W) = E\left(\frac{X}{X+Y}\right) = \frac{E(X)}{E(X+Y)} = \frac{a/\lambda}{a/\lambda + b/\lambda} = \frac{a}{a+b}.$$

Another approach is to proceed from the definition of expectation:

$$E(W) = \int_0^1 \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} w^a (1-w)^{b-1} dw$$

By pattern recognition, the integrand is a $\text{Beta}(a+1, b)$ PDF, up to a normalizing constant. After obtaining an exact match for the PDF, we apply properties of the gamma function:

$$\begin{aligned} E(W) &= \frac{\Gamma(a+b)}{\Gamma(a)} \frac{\Gamma(a+1)}{\Gamma(a+b+1)} \int_0^1 \frac{\Gamma(a+b+1)}{\Gamma(a+1)\Gamma(b)} w^a (1-w)^{b-1} dw \\ &= \frac{\Gamma(a+b)}{\Gamma(a)} \frac{a\Gamma(a)}{(a+b)\Gamma(a+b)} \\ &= \frac{a}{a+b} \end{aligned}$$

In Exercise 30, you will use this approach to find the variance and the other moments of the Beta distribution.

Now that we know the Beta normalizing constant, we can also quickly obtain the PMF of the Beta-Binomial distribution.

Example 8.5.3 (Beta-Binomial PMF). Let $X | p \sim \text{Bin}(n, p)$, with $p \sim \text{Beta}(a, b)$. As mentioned in Story 8.3.3, X has a Beta-Binomial distribution. Find the marginal distribution of X .

Solution: Let $f(p)$ be the $\text{Beta}(a, b)$ PDF. Then

$$\begin{aligned} P(X = k) &= \int_0^1 P(X = k | p) f(p) dp \\ &= \frac{1}{\beta(a, b)} \int_0^1 \binom{n}{k} p^k (1-p)^{n-k} p^{a-1} (1-p)^{b-1} dp \\ &= \frac{\binom{n}{k}}{\beta(a, b)} \int_0^1 p^{a+k-1} (1-p)^{b+n-k-1} dp \\ &= \binom{n}{k} \frac{\beta(a+k, b+n-k)}{\beta(a, b)}, \end{aligned}$$

for $k = 0, 1, \dots, n$.

4.8.6 8.6 Order Statistics

The final transformation we will consider in this chapter is the transformation that takes n random variables $X_1, \dots, X_n$ and sorts them in order, producing the transformed r.v.s $\min(X_1, \dots, X_n), \dots, \max(X_1, \dots, X_n)$. The transformed r.v.s are called the order statistics,² and they are often useful when we are concerned with the distribution of extreme values, as we alluded to earlier.

Furthermore, like the sample mean $\bar{X}_n$, the order statistics serve as useful summaries of an experiment, since we can use them to determine the cutoffs for the worst 5% of observations, the worst 25%, the best 25%, and so forth (such cutoffs are called the quantiles of the sample).

Definition 8.6.1 (Order statistics). For r.v.s $X_1, X_2, \dots, X_n$, the order statistics are the random variables $X_{(1)}, X_{(2)}, \dots, X_{(n)}$, where

$$\begin{aligned} X_{(1)} &= \min(X_1, \dots, X_n), \\ X_{(2)} &\text{ is the second-smallest of } X_1, \dots, X_n, \\ &\vdots \\ X_{(n-1)} &\text{ is the second-largest of } X_1, \dots, X_n, \\ X_{(n)} &= \max(X_1, \dots, X_n). \end{aligned}$$

Note that $X_{(1)} \leq X_{(2)} \leq \dots \leq X_{(n)}$ by definition. We call $X_{(j)}$ the j th order statistic. If n is odd, $X_{((n+1)/2)}$ is called the sample median of $X_1, \dots, X_n$.

8.6.2. The order statistics $X_{(1)}, \dots, X_{(n)}$ are r.v.s, and each $X_{(j)}$ is a function of $X_1, \dots, X_n$. Even if the original r.v.s are independent, the order statistics are dependent: if we know that $X_{(1)} = 100$, then $X_{(n)}$ is forced to be at least 100.

We will focus our attention on the case where $X_1, \dots, X_n$ are i.i.d. continuous r.v.s. The reason is that with discrete r.v.s, there is a positive probability of tied values; with continuous r.v.s, the probability of a tie is exactly 0, which makes matters much easier. Thus, for the rest of this section, assume $X_1, \dots, X_n$ are i.i.d. and continuous, with CDF F and PDF f . We will derive the marginal CDF and PDF of each individual order statistic $X_{(j)}$, as well as the joint PDF of $(X_{(1)}, \dots, X_{(n)})$.

A complication we run into right away is that the transformation to order statistics is not invertible: starting with $\min(X, Y) = 3$ and $\max(X, Y) = 5$, we can't tell whether the original values of X and Y were 3 and 5, respectively, or 5 and 3. Therefore the change of variables formula from $\mathbb{R}^n$ to $\mathbb{R}^n$ does not apply. Instead we will take a direct approach, using pictures to guide us when necessary.


Let's start with the CDF of $X_{(n)} = \max(X_1, \dots, X_n)$. Since $X_{(n)}$ is less than x if and only if all of the X_j are less than x , the CDF of $X_{(n)}$ is

$$\begin{aligned} F_{X_{(n)}}(x) &= P(\max(X_1, \dots, X_n) \leq x) \\ &= P(X_1 \leq x, \dots, X_n \leq x) \\ &= P(X_1 \leq x) \dots P(X_n \leq x) \\ &= (F(x))^n, \end{aligned}$$

where F is the CDF of the individual X_i . Similarly, $X_{(1)} = \min(X_1, \dots, X_n)$ exceeds x if and only if all of the X_j exceed x , so the CDF of $X_{(1)}$ is

$$\begin{aligned} F_{X_{(1)}}(x) &= 1 - P(\min(X_1, \dots, X_n) > x) \\ &= 1 - P(X_1 > x, \dots, X_n > x) \\ &= 1 - (1 - F(x))^n. \end{aligned}$$

The same logic lets us find the CDF of $X_{(j)}$. For the event $X_{(j)} \leq x$ to occur, we need at least j of the X_i to fall to the left of x . This is illustrated in Figure 8.10.



⁰² This term sometimes causes confusion. In statistics (the field of study), any function of the data is called a statistic. If $X_1, \dots, X_n$ are the data, then $\min(X_1, \dots, X_n)$ is a statistic, and so is $\max(X_1, \dots, X_n)$. They are called order statistics because we get them by sorting the data in order (from smallest to largest).

FIGURE 8.10 The event $X_{(j)} \leq x$ is equivalent to the event "at least j X_i 's fall to the left of x ".

Since it appears that the number of X_i to the left of x will be important to us, let's define a new random variable, N , to keep track of just that: define N to be the number of X_i that land to the left of x . Each X_i lands to the left of x with probability $F(x)$, independently. If we define success as landing to the left of x , we have n independent Bernoulli trials with probability $F(x)$ of success, so $N \sim \text{Bin}(n, F(x))$. Then, by the Binomial PMF,

$$\begin{aligned} P(X_{(j)} \leq x) &= P(\text{at least } j \text{ of the } X_i \text{ are to the left of } x) \\ &= P(N \geq j) \\ &= \sum_{k=j}^n \binom{n}{k} F(x)^k (1 - F(x))^{n-k}. \end{aligned}$$

We thus have the following result for the CDF of $X_{(j)}$.

Theorem 8.6.3 (CDF of order statistic). Let $X_1, \dots, X_n$ be i.i.d. continuous r.v.s with CDF F . Then the CDF of the j th order statistic $X_{(j)}$ is

$$P(X_{(j)} \leq x) = \sum_{k=j}^n \binom{n}{k} F(x)^k (1 - F(x))^{n-k}.$$

To get the PDF of $X_{(j)}$, we can differentiate the CDF with respect to x , but the resulting expression is ugly (though it can be simplified). Instead we will take a more direct approach. Consider $f_{X_{(j)}}(x)dx$, the probability that the j th order statistic falls into an infinitesimal interval of length dx around x . The only way this can happen is illustrated in Figure 8.11. We need one of the X_i to fall into the infinitesimal interval around x , and we need exactly $j-1$ of the X_i to fall to the left of x , leaving the remaining $n-j$ to fall to the right of x .




FIGURE 8.11 In order for $X_{(j)}$ to fall within a small interval of x , we require that one of the X_i fall within the small interval and that exactly $j-1$ fall to the left of x .

What is the probability of this extremely specific event? Let's break up the experiment into stages.

- First, we choose which one of the X_i will fall into the infinitesimal interval around x . There are n such choices, each of which occurs with probability $f(x)dx$, where f is the PDF of the X_i .
- Next, we choose exactly $j-1$ out of the remaining $n-1$ to fall to the left of x . There are $\binom{n-1}{j-1}$ such choices, each with probability $F(x)^{j-1}(1 - F(x))^{n-j}$ by the $\text{Bin}(n, F(x))$ PMF.

We multiply the probabilities of the two stages to get

$$f_{X_{(j)}}(x)dx = n f(x)dx \binom{n-1}{j-1} F(x)^{j-1} (1 - F(x))^{n-j}$$

Dropping the dx 's from both sides gives us the PDF we desire.

Theorem 8.6.4 (PDF of order statistic). Let $X_1, \dots, X_n$ be i.i.d. continuous r.v.s with CDF F and PDF f . Then the marginal PDF of the j th order statistic $X_{(j)}$ is

$$f_{X_{(j)}}(x) = n \binom{n-1}{j-1} f(x) F(x)^{j-1} (1 - F(x))^{n-j}.$$

In general, the order statistics of $X_1, \dots, X_n$ will not follow a named distribution, but the order statistics of the standard Uniform distribution are an exception.

Example 8.6.5 (Order statistics of Uniforms). Let $U_1, \dots, U_n$ be i.i.d. $\text{Unif}(0, 1)$. Then for $0 \leq x \leq 1$, $f(x) = 1$ and $F(x) = x$, so the PDF of $U_{(j)}$ is

$$f_{U_{(j)}}(x) = n \binom{n-1}{j-1} x^{j-1} (1-x)^{n-j}.$$

This is the $\text{Beta}(j, n-j+1)$ PDF! So $U_{(j)} \sim \text{Beta}(j, n-j+1)$, and $E(U_{(j)}) = \frac{j}{n+1}$. The simple case $n = 2$ is consistent with Example 7.2.2, where we used 2D LOTUS to show that for i.i.d. $U_1, U_2 \sim \text{Unif}(0, 1)$,

$$E(\max(U_1, U_2)) = 2/3, E(\min(U_1, U_2)) = 1/3.$$

Now that we know $\max(U_1, U_2)$ and $\min(U_1, U_2)$ follow Beta distributions, the expectation of the Beta distribution confirms our earlier findings.

4.8.7 8.7 Recap

In this chapter we discussed three broad classes of transformations:

- invertible transformations $\mathbf{Y} = g(\mathbf{X})$ of continuous random vectors, which can be handled with the change of variables formula (under some technical assumptions, most notably that the partial derivatives $\partial x_i / \partial y_j$ exist and are continuous);
- convolutions, for which we can determine the distribution using (in decreasing order of preference) stories, MGFs, or convolution sums/integrals;
- the transformation of i.i.d. continuous r.v.s to their order statistics.

Figure 8.12 illustrates connections between the original random vector (X, Y) and the transformed random vector $(Z, W) = g(X, Y)$, where g is an invertible transformation satisfying certain technical assumptions. The change of variables formula uses Jacobians to take us back and forth between the joint PDF of (X, Y) and the joint PDF of (Z, W) .

Let A be a region in the support of (X, Y) , and $B = \{g(x, y) : (x, y) \in A\}$ be the corresponding region in the support of (Z, W) . Then $(X, Y) \in A$ is the same event as $(Z, W) \in B$, so

$$P((X, Y) \in A) = P((Z, W) \in B).$$

To find this probability, we can either integrate the joint PDF of (X, Y) over A or integrate the joint PDF of (Z, W) over B .




FIGURE 8.12

Let $(Z, W) = g(X, Y)$, where g is an invertible transformation satisfying certain technical assumptions. The change of variables formula lets us go back and forth between the joint PDFs of (X, Y) and (Z, W) .

In this chapter, as in many others, we made extensive use of Bayes' rule and LOTP, especially in continuous or hybrid forms. And we often used the strategy of integration by pattern recognition. Since any valid PDF must integrate to 1, and by now we know lots of valid PDFs, we can often use probability to help us do calculus in addition to using calculus to help us do probability!

The two new distributions we introduced are the Beta and Gamma, which are laden with stories and connections to other distributions. The Beta is a generalization of the $\text{Unif}(0, 1)$ distribution, and it has the following stories.


- Order statistics of the Uniform: The j th order statistic of n i.i.d. $\text{Unif}(0, 1)$ r.v.s is distributed $\text{Beta}(j, n - j + 1)$.
- Unknown probability, conjugate prior of the Binomial: If $p \sim \text{Beta}(a, b)$ and $X | p \sim \text{Bin}(n, p)$, then $p | X = k \sim \text{Beta}(a + k, b + n - k)$. The posterior distribution of p stays within the Beta family of distributions after updating based on Binomial data, a property known as conjugacy. The parameters a and b can be interpreted as the prior number of successes and failures, respectively.

The Gamma is a generalization of the Exponential distribution, and it has the following stories.

- Poisson process: In a Poisson process of rate λ , the total waiting time for n arrivals is distributed $\text{Gamma}(n, \lambda)$. Thus the Gamma is the continuous analog of the Negative Binomial distribution.
- Unknown rate, conjugate prior of the Poisson: If $\lambda \sim \text{Gamma}(r_0, b_0)$ and $Y | \lambda \sim \text{Pois}(\lambda t)$, then $\lambda | Y = y \sim \text{Gamma}(r_0 + y, b_0 + t)$. The posterior distribution of λ stays within the Gamma family of distributions after updating based on Poisson data. The parameters r_0 and b_0 can be interpreted as the prior number of observed successes and the total waiting time for those successes, respectively.

The Beta and Gamma distributions are related by the bank-post office story, which says that if $X \sim \text{Gamma}(a, \lambda)$, $Y \sim \text{Gamma}(b, \lambda)$ are independent, then $X + Y \sim \text{Gamma}(a + b, \lambda)$, $\frac{X}{X+Y} \sim \text{Beta}(a, b)$, with $X + Y$ and $\frac{X}{X+Y}$ independent.

The diagram of connections, which we last saw in Chapter 5, is hereby updated to include the Beta and Gamma distributions. Distributions listed in parentheses are special cases of the ones not in parentheses.



4.8.8 8.8r

Beta and Gamma Distributions

The Beta and Gamma distributions are programmed into R.

- `dbeta`, `pbeta`, `rbeta`: To evaluate the $\text{Beta}(a, b)$ PDF or CDF at x , we use `dbeta(x, a, b)` and `pbeta(x, a, b)`. To generate n realizations from the $\text{Beta}(a, b)$ distribution, we use `rbeta(n, a, b)`.
- `dgamma`, `pgamma`, `rgamma`: To evaluate the $\text{Gamma}(a, \lambda)$ PDF or CDF at x , we use `dgamma(x, a, lambda)` or `pgamma(x, a, lambda)`. To generate n realizations from the $\text{Gamma}(a, \lambda)$ distribution, we use `rgamma(n, a, lambda)`.

For example, we can check that the $\text{Gamma}(3, 2)$ distribution has mean $3/2$ and variance $3/4$. To do this, we generate a large number of $\text{Gamma}(3, 2)$ random variables using `rgamma`, then compute their mean and var:

```
y <- rgamma(10^5, 3, 2)
mean(y)
var(y)
```

Did you get values that were close to 1.5 and 0.75, respectively?

Convolution of Uniforms

Using R, we can quickly verify that for $X, Y \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(0, 1)$, the distribution of $T = X + Y$ is triangular in shape:

```
x <- runif(10^5)
y <- runif(10^5)
hist(x+y)
```

The histogram looks like an ascending and then descending staircase, a discrete approximation to a triangle.

Bayes' Billiards

In the Bayes' billiards story, we have n white balls and 1 gray ball, throw them onto the unit interval completely at random, and count the number of white balls to the left of the gray ball. Letting p be the position of the gray ball and X be the number of white balls to the left of the gray ball, we have

$$p \sim \text{Unif}(0, 1)$$

$$X \mid p \sim \text{Bin}(n, p).$$

By performing this experiment a large number of times, we can verify the results we derived in this chapter about the marginal PMF of X and the posterior PDF of p given $X = x$. We'll let the number of simulations be called `nsim`, to avoid a name conflict with the number of white balls, `n`, which we set equal to 10 :

```
nsim <- 10^5
n <- 10
```

We simulate 10^5 values of p , then simulate 10^5 values from the conditional distribution of X given p :

```
p <- runif(nsim)
x <- rbinom(nsim, n, p)
```

Notice that we feed the entire vector `p` into `rbinom`. This means that the first element of `x` is generated using the first element of `p`, the second element of `x` is generated using the second element of `p`, and so forth. Thus, conditional on a particular element of `p`, the corresponding element of `x` is Binomial, but the elements of `p` are themselves Uniform, exactly as the model specifies.

According to the Bayes' billiards argument, the marginal distribution of X should be Discrete Uniform on the integers 0 through n . Is this in fact the case? We can make a histogram of `x` to check! Because the distribution of X is discrete, we tell R to make the histogram breaks at $-0.5, 0.5, 1.5, \dots$ so that each bar is centered at an integer value:

```
hist(x, breaks=seq(-0.5, n+0.5, 1))
```

Indeed, all the histogram bars are approximately equal in height, consistent with a Discrete Uniform distribution.

Now for the posterior distribution of p given $X = x$. Conditioning is very simple in R. To consider only the simulated values of p where the value of X was 3, we use square brackets, like this: `p[x == 3]`. In particular, we can create a histogram of these values to see what the posterior distribution of p given $X = 3$ looks like; try `hist ( p[x == 3] )`.

According to the Beta-Binomial conjugacy result, the true posterior distribution is $p \mid X = 3 \sim \text{Beta}(4, 8)$. We can plot the histogram of `p[x == 3]` next to a histogram of simulated values from the $\text{Beta}(4, 8)$ distribution to confirm that they look similar:

```
par(mfrow=c(1,2))
hist(p[x==3])
hist(rbeta(10^4, 4, 8))
```

The first line tells R we want two side-by-side plots, and the second and third lines create the histograms.

Simulating Order Statistics

Simulating order statistics in R is easy: we simply simulate i.i.d. r.v.s and use sort to sort them in increasing order. For example,

```
sort(rnorm(10))
```

produces one realization of $X_{(1)}, \dots, X_{(10)}$, where $X_1, \dots, X_{10}$ are i.i.d. $\mathcal{N}(0, 1)$.

If we want to plot a histogram of realizations of, say, $X_{(9)}$, we'll need to use replicate:

```
order_stats <- replicate(10^4, sort(rnorm(10)))
```

This creates a matrix, `order_stats`, with 10 rows. The i th row of the matrix contains 10^4 realizations of $X_{(i)}$. Now we can create a histogram of $X_{(9)}$, simply by selecting row 9 of the matrix:

```
x9 <- order_stats[9,]  
hist(x9)
```

We can also compute summaries like `mean(x9)` and `var(x9)`.

4.8.9 8.9 Exercises

Exercises marked with (S) have detailed solutions at <http://stat110.net>.

Change of Variables

- Find the PDF of e^{-X} for $X \sim \text{Expo}(1)$.
- Find the PDF of X^7 for $X \sim \text{Expo}(\lambda)$.
- Find the PDF of Z^3 for $Z \sim \mathcal{N}(0, 1)$.
- (S) Find the PDF of Z^4 for $Z \sim \mathcal{N}(0, 1)$.
- Find the PDF of $|Z|$ for $Z \sim \mathcal{N}(0, 1)$.
- (S) Let $U \sim \text{Unif}(0, 1)$. Find the PDFs of U^2 and $\sqrt{U}$.
- Let $U \sim \text{Unif}(0, \frac{\pi}{2})$. Find the PDF of $\sin(U)$.
- (a) Find the distribution of X^2 for $X \sim \text{DUnif}(0, 1, \dots, n)$.
(b) Find the distribution of X^2 for $X \sim \text{DUnif}(-n, -n + 1, \dots, 0, 1, \dots, n)$.
- Let $X \sim \text{Bern}(p)$ and let a and b be constants with $a < b$. Find a simple transformation of X that yields an r.v. that equals a with probability $1 - p$ and equals b with probability p .
- Let $X \sim \text{Pois}(\lambda)$ and Y be the indicator of X being odd. Find the PMF of Y .

Hint: Find $P(Y = 0) - P(Y = 1)$ by writing $P(Y = 0)$ and $P(Y = 1)$ as series and then using the fact that $(-1)^k$ is 1 if k is even and -1 if k is odd.

- Let T be a continuous r.v. and $V = 1/T$. Show that their CDFs are related as follows:

$$F_V(v) = \begin{cases} F_T(0) + 1 - F_T\left(\frac{1}{v}\right) & \text{for } v > 0 \\ F_T(0) & \text{for } v = 0 \\ F_T(0) - F_T\left(\frac{1}{v}\right) & \text{for } v < 0 \end{cases}$$

12. Let T be the ratio X/Y of two i.i.d. $\mathcal{N}(0, 1)$ r.v.s. X, Y . This is the Cauchy distribution and, as shown in Example 7.1.25, it has PDF

$$f_T(t) = \frac{1}{\pi(1+t^2)}$$

(a) Use the result of the previous problem to find the CDF of $1/T$. Then use calculus to find the PDF of $1/T$. (Note that the one-dimensional change of variables formula does not apply directly, since the function $g(t) = 1/t$, even though it has $g'(t) < 0$ for all $t \neq 0$, is undefined at $t = 0$ and is not a strictly decreasing function on its domain.)

(b) Show that $1/T$ has the same distribution as T without using calculus, in 140 characters or fewer.

13. Let X and Y be i.i.d. $\text{Expo}(\lambda)$, and $T = \log(X/Y)$. Find the CDF and PDF of T .

14. Let X and Y have joint PDF $f_{X,Y}(x, y)$, and transform $(X, Y) \mapsto (T, W)$ linearly by letting

$$T = aX + bY \text{ and } W = cX + dY$$

where a, b, c, d are constants such that $ad - bc \neq 0$.

(a) Find the joint PDF $f_{T,W}(t, w)$ (in terms of $f_{X,Y}$, though your answer should be written as a function of t and w).

(b) For the case where $T = X + Y, W = X - Y$, show that

$$f_{T,W}(t, w) = \frac{1}{2} f_{X,Y} \left(\frac{t+w}{2}, \frac{t-w}{2} \right)$$

15. (S) Let X, Y be continuous r.v.s with a spherically symmetric joint distribution, which means that the joint PDF is of the form $f(x, y) = g(x^2 + y^2)$ for some function g . Let (R, θ) be the polar coordinates of (X, Y) , so $R^2 = X^2 + Y^2$ is the squared distance from the origin and θ is the angle (in $[0, 2\pi)$), with $X = R \cos \theta, Y = R \sin \theta$.

(a) Explain intuitively why R and θ are independent. Then prove this by finding the joint PDF of (R, θ) .

(b) What is the joint PDF of (R, θ) when (X, Y) is Uniform in the unit disk $\{(x, y) : x^2 + y^2 \leq 1\}$?

(c) What is the joint PDF of (R, θ) when X and Y are i.i.d. $\mathcal{N}(0, 1)$?

16. Let X and Y be i.i.d. $\mathcal{N}(0, 1)$ r.v.s, $T = X + Y$, and $W = X - Y$. We know from Example 7.5.8 that T and W are independent $\mathcal{N}(0, 2)$ r.v.s (note that (T, W) is Multivariate Normal with $\text{Cov}(T, W) = 0$). Give another proof of this fact, using the change of variables theorem.

17. Let X and Y be i.i.d. $\mathcal{N}(0, 1)$ r.v.s, and (R, θ) be the polar coordinates for the point (X, Y) , so $X = R \cos \theta$ and $Y = R \sin \theta$ with $R \geq 0$ and $\theta \in [0, 2\pi)$. Find the joint PDF of R^2 and θ . Also find the marginal distributions of R^2 and θ , giving their names (and parameters) if they are distributions we have studied before.

18. Let X and Y be independent positive r.v.s, with PDFs f_X and f_Y , respectively. Let T be the ratio X/Y .

(a) Find the joint PDF of T and X , using a Jacobian.

(b) Find the marginal PDF of T , as a single integral.

19. Let X and Y be i.i.d. $\text{Expo}(\lambda)$, and transform them to $T = X + Y, W = X/Y$.

(a) Find the joint PDF of T and W . Are they independent?

(b) Find the marginal PDFs of T and W .

Convolutions

20. Let $U \sim \text{Unif}(0, 1)$ and $X \sim \text{Expo}(1)$, independently. Find the PDF of $U + X$.
21. Let X and Y be i.i.d. $\text{Expo}(1)$. Use a convolution integral to show that the PDF of $L = X - Y$ is $f(t) = \frac{1}{2}e^{-|t|}$ for all real t ; this is known as the Laplace distribution.
22. Use a convolution integral to show that if $X \sim \mathcal{N}(\mu_1, \sigma^2)$ and $Y \sim \mathcal{N}(\mu_2, \sigma^2)$ are independent, then $T = X + Y \sim \mathcal{N}(\mu_1 + \mu_2, 2\sigma^2)$ (to simplify the calculation, we are assuming that the variances are equal). You can use a standardization (location-scale) idea to reduce to the standard Normal case before setting up the integral.

Hint: Complete the square.

23. (S) Let X and Y be independent positive r.v.s, with PDFs f_X and f_Y , respectively, and consider the product $T = XY$. When asked to find the PDF of T , Jacobno argues that: "It's like a convolution, with a product instead of a sum. To have $T = t$ we need $X = x$ and $Y = t/x$ for some x ; that has probability $f_X(x)f_Y(t/x)$, so summing up these possibilities we get that the PDF of T is $\int_0^\infty f_X(x)f_Y(t/x)dx$."

Evaluate Jacobno's argument, while getting the PDF of T (as an integral) in 2 ways:

- (a) using the continuous version of the law of total probability to get the CDF, and then taking the derivative (you can assume that swapping the derivative and integral is valid);
- (b) by taking the log of both sides of $T = XY$ and doing a convolution (and then converting back to get the PDF of T).

24. Let X and Y be i.i.d. Discrete Uniform r.v.s on $\{0, 1, \dots, n\}$, where n is a positive integer. Find the PMF of $T = X + Y$.

Hint: In finding $P(T = k)$, it helps to consider the cases $0 \leq k \leq n$ and $n + 1 \leq k \leq 2n$ separately. Be careful about the range of summation in the convolution sum.

25. Let X and Y be i.i.d. $\text{Unif}(0, 1)$, and let $W = X - Y$.
 - (a) Find the mean and variance of W , without yet deriving the PDF.
 - (b) Show that the distribution of W is symmetric about 0, without yet deriving the PDF.
 - (c) Find the PDF of W .
 - (d) Use the PDF of W to verify your results from (a) and (b).
 - (e) How does the distribution of W relate to the distribution of $X + Y$, the Triangle distribution derived in Example 8.2.5? Give a precise description, e.g., using the concepts of location and scale.

26. Let X and Y be i.i.d. $\text{Unif}(0, 1)$, and $T = X + Y$. We derived the distribution of T (a Triangle distribution) in Example 8.2.5, using a convolution integral. Since (X, Y) is Uniform in the unit square $\{(x, y) : 0 < x < 1, 0 < y < 1\}$, we can also interpret $P((X, Y) \in A)$ as the area of A , for any region A within the unit square. Use this idea to find the CDF of T , by interpreting the CDF (evaluated at some point) as an area.

27. Let X, Y, Z be i.i.d. $\text{Unif}(0, 1)$, and $W = X + Y + Z$. Find the PDF of W .

Hint: We already know the PDF of $X + Y$. Be careful about limits of integration in the convolution integral; there are 3 cases that should be considered separately.

Beta and Gamma

28. (S) Let $B \sim \text{Beta}(a, b)$. Find the distribution of $1 - B$ in two ways: (a) using a change of variables and (b) using a story proof. Also explain why the result makes sense in terms of Beta being the conjugate prior for the Binomial.
29. (S) Let $X \sim \text{Gamma}(a, \lambda)$ and $Y \sim \text{Gamma}(b, \lambda)$ be independent, with a and b integers. Show that $X + Y \sim \text{Gamma}(a + b, \lambda)$ in three ways: (a) with a convolution integral; (b) with MGFs; (c) with a story proof.

30. Let $B \sim \text{Beta}(a, b)$. Use integration by pattern recognition to find $E(B^k)$ for positive integers k . In particular, show that

$$\text{Var}(B) = \frac{ab}{(a+b)^2(a+b+1)}$$

31. (S) Fred waits $X \sim \text{Gamma}(a, \lambda)$ minutes for the bus to work, and then waits $Y \sim \text{Gamma}(b, \lambda)$ for the bus going home, with X and Y independent. Is the ratio X/Y independent of the total wait time $X + Y$?
32. (S) The F -test is a very widely used statistical test based on the $F(m, n)$ distribution, which is the distribution of $\frac{X/m}{Y/n}$ with $X \sim \text{Gamma}(\frac{m}{2}, \frac{1}{2})$, $Y \sim \text{Gamma}(\frac{n}{2}, \frac{1}{2})$. Find the distribution of $mV/(n + mV)$ for $V \sim F(m, n)$.
33. (S) Customers arrive at the Leftorium store according to a Poisson process with rate λ customers per hour. The true value of λ is unknown, so we treat it as a random variable. Suppose that our prior beliefs about λ can be expressed as $\lambda \sim \text{Expo}(3)$. Let X be the number of customers who arrive at the Leftorium between 1 pm and 3 pm tomorrow. Given that $X = 2$ is observed, find the posterior PDF of λ .
34. (S) Let X and Y be independent, positive r.v.s. with finite expected values.
 (a) Give an example where $E\left(\frac{X}{X+Y}\right) \neq \frac{E(X)}{E(X+Y)}$, computing both sides exactly.

Hint: Start by thinking about the simplest examples you can think of!

- (b) If X and Y are i.i.d., then is it necessarily true that $E\left(\frac{X}{X+Y}\right) = \frac{E(X)}{E(X+Y)}$?
- (c) Now let $X \sim \text{Gamma}(a, \lambda)$ and $Y \sim \text{Gamma}(b, \lambda)$. Show without using calculus that

$$E\left(\frac{X^c}{(X+Y)^c}\right) = \frac{E(X^c)}{E((X+Y)^c)}$$

for every real $c > 0$.

35. Let $T = X^{1/\gamma}$, with $X \sim \text{Expo}(\lambda)$ and $\lambda, \gamma > 0$. So T has a Weibull distribution, as discussed in Example 6.5.5. Using LOTUS and the definition of the gamma function, we showed in this chapter that for $Y \sim \text{Gamma}(a, \lambda)$,

$$E(Y^c) = \frac{1}{\lambda^c} \cdot \frac{\Gamma(a+c)}{\Gamma(a)}$$

for all real $c > -a$. Use this result to show that

$$E(T) = \frac{\Gamma(1+1/\gamma)}{\lambda^{1/\gamma}} \text{ and } \text{Var}(T) = \frac{\Gamma(1+2/\gamma) - (\Gamma(1+2/\gamma))^2}{\lambda^{2/\gamma}}.$$

36. Alice walks into a post office with 2 clerks. Both clerks are in the midst of serving customers, but Alice is next in line. The clerk on the left takes an $\text{Expo}(\lambda_1)$ time to serve a customer, and the clerk on the right takes an $\text{Expo}(\lambda_2)$ time to serve a customer. Let T_1 be the time until the clerk on the left is done serving their current customer, and define T_2 likewise for the clerk on the right.
 (a) If $\lambda_1 = \lambda_2$, is T_1/T_2 independent of $T_1 + T_2$?

Hint: Note that $T_1/T_2 = (T_1/(T_1 + T_2)) / (T_2/(T_1 + T_2))$.

- (b) Find $P(T_1 < T_2)$ (do not assume $\lambda_1 = \lambda_2$ here or in the next part, but do check that your answers make sense in that special case).

(c) Find the expected total amount of time that Alice spends in the post office (assuming that she leaves immediately after she is done being served).

37. Let $X \sim \text{Pois}(\lambda t)$ and $Y \sim \text{Gamma}(j, \lambda)$, where j is a positive integer. Show using a story about a Poisson process that

$$P(X \geq j) = P(Y \leq t)$$

38. Visitors arrive at a certain scenic park according to a Poisson process with rate λ visitors per hour. Fred has just arrived (independent of anyone else), and will stay for an $\text{Expo}(\lambda_2)$ number of hours. Find the distribution of the number of other visitors who arrive at the park while Fred is there.

39. (a) Let $p \sim \text{Beta}(a, b)$, where a and b are positive real numbers. Find $E(p^2(1-p)^2)$, fully simplified (Γ should not appear in your final answer).

Two teams, A and B , have an upcoming match. They will play five games and the winner will be declared to be the team that wins the majority of games. Given p , the outcomes of games are independent, with probability p of team A winning and $1-p$ of team B winning. But you don't know p , so you decide to model it as an r.v., with $p \sim \text{Unif}(0, 1)$ a priori (before you have observed any data).

To learn more about p , you look through the historical records of previous games between these two teams, and find that the previous outcomes were, in chronological order, $AAABBAABAB$. (Assume that the true value of p has not been changing over time and will be the same for the match, though your beliefs about p may change over time.)

(b) Does your posterior distribution for p , given the historical record of games between A and B , depend on the specific order of outcomes or only on the fact that A won exactly 6 of the 10 games on record? Explain.

(c) Find the posterior distribution for p , given the historical data.

The posterior distribution for p from (c) becomes your new prior distribution, and the match is about to begin!

(d) Conditional on p , is the indicator of A winning the first game of the match positively correlated with, uncorrelated with, or negatively correlated with the indicator of A winning the second game? What about if we only condition on the historical data?

(e) Given the historical data, what is the expected value for the probability that the match is not yet decided when going into the fifth game (viewing this probability as an r.v. rather than a number, to reflect our uncertainty about it)?

40. An engineer is studying the reliability of a product by performing a sequence of n trials. Reliability is defined as the probability of success. In each trial, the product succeeds with probability p and fails with probability $1-p$. The trials are conditionally independent given p . Here p is unknown (else the study would be unnecessary!). The engineer takes a Bayesian approach, with $p \sim \text{Unif}(0, 1)$ as prior.

Let r be a desired reliability level and c be the corresponding confidence level, in the sense that, given the data, the probability is c that the true reliability p is at least r . For example, if $r = 0.9, c = 0.95$, we can be 95% sure, given the data, that the product is at least 90% reliable. Suppose that it is observed that the product succeeds all n times. Find a simple equation for c as a function of r .

Order Statistics

41. (S) Let $X \sim \text{Bin}(n, p)$ and $B \sim \text{Beta}(j, n-j+1)$, where n is a positive integer and j is a positive integer with $j \leq n$. Show using a story about order statistics that

$$P(X \geq j) = P(B \leq p).$$

This shows that the CDF of the continuous r.v. B is closely related to the CDF of the discrete r.v. X , and is another connection between the Beta and Binomial.

42. Show that for i.i.d. continuous r.v.s X, Y, Z ,

$$P(X < \min(Y, Z)) + P(Y < \min(X, Z)) + P(Z < \min(X, Y)) = 1.$$

43. Show that

$$\int_0^x \frac{n!}{(j-1)!(n-j)!} t^{j-1} (1-t)^{n-j} dt = \sum_{k=j}^n \binom{n}{k} x^k (1-x)^{n-k}$$

without using calculus, for all $x \in [0, 1]$ and j, n positive integers with $j \leq n$.

44. Let $X_1, \dots, X_n$ be i.i.d. continuous r.v.s with PDF f and a strictly increasing CDF F . Suppose that we know that the j th order statistic of n i.i.d. $\text{Unif}(0, 1)$ r.v.s is a $\text{Beta}(j, n-j+1)$, but we have forgotten the formula and derivation for the distribution of the j th order statistic of $X_1, \dots, X_n$. Show how we can recover the PDF of $X_{(j)}$ quickly using a change of variables.

45. (S) Let X and Y be independent $\text{Expo}(\lambda)$ r.v.s and $M = \max(X, Y)$. Show that M has the same distribution as $X + \frac{1}{2}Y$, in two ways: (a) using calculus and (b) by remembering the memoryless property and other properties of the Exponential.

46. (S) (a) If X and Y are i.i.d. continuous r.v.s with CDF $F(x)$ and PDF $f(x)$, then $M = \max(X, Y)$ has PDF $2F(x)f(x)$. Now let X and Y be discrete and i.i.d., with CDF $F(x)$ and PMF $f(x)$. Explain in words why the PMF of M is not $2F(x)f(x)$.

(b) Let X and Y be i.i.d. $\text{Bern}(1/2)$ r.v.s, $M = \max(X, Y)$ and $L = \min(X, Y)$. Find the joint PMF of M and L , i.e., $P(M = a, L = b)$, and the marginal PMFs of M and L .

47. Let $X_1, X_2, \dots$ be i.i.d. r.v.s with CDF F , and let $M_n = \max(X_1, X_2, \dots, X_n)$. Find the joint distribution of M_n and M_{n+1} , for each $n \geq 1$.

48. (S) Let $X_1, X_2, \dots, X_n$ be i.i.d. r.v.s with CDF F and PDF f . Find the joint PDF of the order statistics $X_{(i)}$ and $X_{(j)}$ for $1 \leq i < j \leq n$, by drawing and thinking about a picture.

49. (S) Two women are pregnant, both with the same due date. On a timeline, define time 0 to be the instant when the due date begins. Suppose that the time when the woman gives birth has a Normal distribution, centered at 0 and with standard deviation 8 days. The two birth times are i.i.d. Let T be the time of the first of the two births (in days).

(a) Show that

$$E(T) = \frac{-8}{\sqrt{\pi}}$$

Hint: For any two random variables X and Y , we have $\max(X, Y) + \min(X, Y) = X + Y$ and $\max(X, Y) - \min(X, Y) = |X - Y|$. Example 7.2.3 derives the expected distance between two i.i.d. $\mathcal{N}(0, 1)$ r.v.s.

(b) Find $\text{Var}(T)$, in terms of integrals. You can leave your answers unsimplified for this part, but it can be shown that the answer works out to

$$\text{Var}(T) = 64 \left(1 - \frac{1}{\pi} \right)$$

50. We are about to observe random variables $Y_1, Y_2, \dots, Y_n$, i.i.d. from a continuous distribution. We will need to predict an independent future observation Y_{new} , which will also

have the same distribution. The distribution is unknown, so we will construct our prediction using $Y_1, Y_2, \dots, Y_n$ rather than the distribution of Y_{new} . In forming a prediction, we do not want to report only a single number; rather, we want to give a predictive interval with "high confidence" of containing Y_{new} . One approach to this is via order statistics.

(a) For fixed j and k with $1 \leq j < k \leq n$, find $P(Y_{\text{new}} \in [Y_{(j)}, Y_{(k)}])$.

Hint: By symmetry, all orderings of $Y_1, \dots, Y_n, Y_{\text{new}}$ are equally likely.

(b) Let $n = 99$. Construct a predictive interval, as a function of $Y_1, \dots, Y_n$, such that the probability of the interval containing Y_{new} is 0.95.

51. Let $X_1, \dots, X_n$ be i.i.d. continuous r.v.s with n odd. Show that the median of the distribution of the sample median of the X_i 's is the median of the distribution of the X_i 's.

Hint: Start by reading the problem carefully; it is crucial to distinguish between the median of a distribution (as defined in Chapter 6) and the sample median of a collection of r.v.s (as defined in this chapter). Of course they are closely related: the sample median of i.i.d. r.v.s is a very natural way to estimate the true median of the distribution that the r.v.s are drawn from. Two approaches to evaluating a sum that might come up are (i) use the first story proof example and first story proof exercise from Chapter 1, or (ii) use the fact that, by the story of the Binomial, $Y \sim \text{Bin}(n, 1/2)$ implies $n - Y \sim \text{Bin}(n, 1/2)$.

Mixed Practice

52. Let $U_1, U_2, \dots, U_n$ be i.i.d. $\text{Unif}(0, 1)$, and let $X_j = -\log(U_j)$ for all j .

(a) Find the distribution of X_j . What is its name?

(b) Find the distribution of the product $U_1 U_2 \dots U_n$.

Hint: First take the log.

53. (S) A DNA sequence can be represented as a sequence of letters, where the alphabet has 4 letters: A, C, T, G. Suppose such a sequence is generated randomly, where the letters are independent and the probabilities of A, C, T, G are p_1, p_2, p_3, p_4 , respectively.

(a) In a DNA sequence of length 115, what is the expected number of occurrences of the expression "CATCAT" (in terms of the p_j)? (Note that, for example, the expression "CATCATCAT" counts as 2 occurrences.)

(b) What is the probability that the first A appears earlier than the first C appears, as letters are generated one by one (in terms of the p_j)?

(c) For this part, assume that the p_j are unknown. Suppose we treat p_2 as a $\text{Unif}(0, 1)$ r.v. before observing any data, and that then the first 3 letters observed are "CAT". Given this information, what is the probability that the next letter is C?

54. (S) Consider independent Bernoulli trials with probability p of success for each. Let X be the number of failures incurred before getting a total of r successes.

(a) Determine what happens to the distribution of $\frac{p}{1-p}X$ as $p \rightarrow 0$, using MGFs; what is the PDF of the limiting distribution, and its name and parameters if it is one we have studied?

Hint: Start by finding the $\text{Geom}(p)$ MGF. Then find the MGF of $\frac{p}{1-p}X$, and use the fact that if the MGFs of Y_n converges to the MGF of Y , then the CDF of Y_n converges to the CDF of Y .

(b) Explain intuitively why the result of (a) makes sense.

4.9 9 Conditional Expectation

Given that you've read the earlier chapters, you already know what conditional expectation is: expectation, but using conditional probabilities. This is an essential concept, for reasons analogous to why we need conditional probability:

- Conditional expectation is a powerful tool for calculating expectations. Using strategies such as conditioning on what we wish we knew and first-step analysis, we can often decompose complicated expectation problems into simpler pieces.
- Conditional expectation is a relevant quantity in its own right, allowing us to predict or estimate unknowns based on whatever evidence is currently available. For example, in statistics we often want to predict a response variable (such as test scores or earnings) based on explanatory variables (such as number of practice problems solved or enrollment in a job training program).

There are two different but closely linked notions of conditional expectation:

- Conditional expectation $E(Y | A)$ given an event: let Y be an r.v., and A be an event. If we learn that A occurred, our updated expectation for Y is denoted by $E(Y | A)$ and is computed analogously to $E(Y)$, except using conditional probabilities given A .
- Conditional expectation $E(Y | X)$ given a random variable: a more subtle question is how to define $E(Y | X)$, where X and Y are both r.v.s. Intuitively, $E(Y | X)$ is the r.v. that best predicts Y using only the information available from X .

In this chapter, we explore the definitions, properties, intuitions, and applications of both forms of conditional expectation.

4.9.1 9.1 Conditional Expectation Given an Event

Recall that the expectation $E(Y)$ of a discrete r.v. Y is a weighted average of its possible values, where the weights are the PMF values $P(Y = y)$. After learning that an event A occurred, we want to use weights that have been updated to reflect this new information. The definition of $E(Y | A)$ simply replaces the probability $P(Y = y)$ with the conditional probability $P(Y = y | A)$.

Similarly, if Y is continuous, $E(Y)$ is still a weighted average of the possible values of Y , with an integral in place of a sum and the PDF value $f(y)$ in place of a PMF value. If we learn that A occurred, we update the expectation for Y by replacing $f(y)$ with the conditional PDF $f(y | A)$.

Definition 9.1.1 (Conditional expectation given an event). Let A be an event with positive probability. If Y is a discrete r.v., then the conditional expectation of Y given A is

$$E(Y | A) = \sum_y y P(Y = y | A),$$

where the sum is over the support of Y . If Y is a continuous r.v. with PDF f , then

$$E(Y | A) = \int_{-\infty}^{\infty} y f(y | A) dy$$

where the conditional PDF $f(y | A)$ is defined as the derivative of the conditional CDF $F(y | A) = P(Y \leq y | A)$, and can also be computed by a hybrid version of Bayes' rule:

$$f(y | A) = \frac{P(A | Y = y) f(y)}{P(A)}.$$

Intuition 9.1.2. To gain intuition for $E(Y | A)$, let's consider approximating it via simulation (or via the frequentist perspective, based on repeating the same experiment many times).

Imagine generating a large number n of replications of the experiment for which Y is a numerical summary. We then have Y -values $y_1, \dots, y_n$, and we can approximate

$$E(Y) \approx \frac{1}{n} \sum_{j=1}^n y_j$$

To approximate $E(Y | A)$, we restrict to the replications where A occurred, and average only those Y -values. This can be written as

$$E(Y | A) \approx \frac{\sum_{j=1}^n y_j I_j}{\sum_{j=1}^n I_j},$$

where I_j is the indicator of A occurring in the j th replication. This is undefined if A never occurred in the simulation, which makes sense since then there is no simulation data about what the " A occurred" scenario is like. We would like to have n large enough so that there are many occurrences of A (if A is a rare event, more sophisticated techniques for approximating $E(Y | A)$ may be needed).

The principle is simple though: $E(Y | A)$ is approximately the average of Y in a large number of simulation runs in which A occurred.

9.1.3. Confusing conditional expectation and unconditional expectation is a dangerous mistake. More generally, not keeping careful track of what you should be conditioning on and what you are conditioning on is a recipe for disaster.

For a life-or-death example of the previous biohazard, consider life expectancy.

Example 9.1.4 (Life expectancy). Fred is 30 years old, and he hears that the average life expectancy in his country is 80 years. Should he conclude that, on average, he has 50 years of life left? No, there is a crucial piece of information that he must condition on: the fact that he has lived to age 30 already. Letting T be Fred's lifespan, we have the cheerful news that

$$E(T) < E(T | T \geq 30)$$

The left-hand side is Fred's life expectancy at birth (it implicitly conditions on the fact that he is born), and the right-hand side is Fred's life expectancy given that he reaches age 30.

A harder question is how to decide on an appropriate estimate to use for $E(T)$. Is it just 80, the overall average for his country? In almost every country, women have a longer average life expectancy than men, so it makes sense to condition on Fred being a man. But should we also condition on what city he was born in? Should we condition on racial and financial information about his parents, or the time of day when he was born? Intuitively, we would like estimates that are both accurate and relevant for Fred, but there is a tradeoff since if we condition on more characteristics of Fred, then there are fewer people who match those characteristics to use as data for estimating the life expectancy.

Now consider some specific numbers for the United States. A Social Security Administration study estimated that between 1900 and 2000, the average life expectancy at birth in the U.S. for men increased from 46 to 74, and for women increased from 49 to 79. Tremendous gains! But much of the gain is due to decreases in child mortality. For a 30-year-old person in 1900, the average number of years remaining was 35 for a man and 36 for a woman; in 2000, the corresponding numbers were 46 for a man and 50 for a woman.

There are some subtle statistical issues in obtaining these estimates. For example, how were estimates for life expectancy for someone born in 2000 obtained without waiting at least until the year 2100? Estimating survival distributions is a very important topic in biostatistics and actuarial science.

The law of total probability allows us to get unconditional probabilities by slicing up the sample space and computing conditional probabilities in each slice. The same idea works for computing unconditional expectations.

Theorem 9.1.5 (Law of total expectation). Let $A_1, \dots, A_n$ be a partition of a sample space, with $P(A_i) > 0$ for all i , and let Y be a random variable on this sample space. Then

$$E(Y) = \sum_{i=1}^n E(Y | A_i) P(A_i)$$


In fact, since all probabilities are expectations by the fundamental bridge, the law of total probability is a special case of the law of total expectation. To see this, let $Y = I_B$ for an event B ; then the above theorem says

$$P(B) = E(I_B) = \sum_{i=1}^n E(I_B | A_i) P(A_i) = \sum_{i=1}^n P(B | A_i) P(A_i),$$


which is exactly LOTP. The law of total expectation is, in turn, a special case of a major result called Adam's law (Theorem 9.3.7), so we will not prove it yet.

There are many interesting examples of using wishful thinking to break up an unconditional expectation into conditional expectations. We begin with two cautionary tales about the importance of conditioning carefully and not destroying information without justification.

Example 9.1.6 (Two-envelope paradox). A stranger presents you with two identical-looking, sealed envelopes, each of which contains a check for some positive amount of money. You are informed that one of the envelopes contains exactly twice as much money as the other. You can choose either envelope. Which do you prefer: the one on the left or the one on the right? (Assume that the expected amount of money in each envelope is finite - certainly a good assumption in the real world!)



X



Y

FIGURE 9.1 Two envelopes, where one contains twice as much money as the other. Either $Y = 2X$ or $Y = X/2$, with equal probabilities. Which would you prefer?

Solution:

Let X and Y be the amounts in the left and right envelopes, respectively. By symmetry, there is no reason to prefer one envelope over the other (we are assuming there is no prior information that the stranger is left-handed and left-handed people prefer putting more money on the left). Concluding by symmetry that $E(X) = E(Y)$, it seems that you should not care which envelope you get.

But as you daydream about what's inside the envelopes, another argument occurs to you: suppose that the left envelope has \$100. Then the right envelope either has \$50 or \$200. The average of \$50 and \$200 is \$125, so it seems then that the right envelope is better. But there was nothing special about \$100 here; for any value x for the left envelope, the average of $2x$ and $x/2$ is greater than x , suggesting that the right envelope is better. This is bizarre though, since not only does it contradict the symmetry argument, but also the same reasoning could be applied starting with the right envelope, leading to switching back and forth forever!

Let us try to formalize this argument to see what's going on. We have $Y = 2X$ or $Y = X/2$, with equal probabilities. By Theorem 9.1.5,

$$E(Y) = E(Y \mid Y = 2X) \cdot \frac{1}{2} + E(Y \mid Y = X/2) \cdot \frac{1}{2}$$

One might then think that this is

$$E(2X) \cdot \frac{1}{2} + E(X/2) \cdot \frac{1}{2} = \frac{5}{4}E(X)$$

suggesting a 25% gain from switching from the left to the right envelope. But there is a blunder in that calculation: $E(Y \mid Y = 2X) = E(2X \mid Y = 2X)$, but there is no justification for dropping the $Y = 2X$ condition after plugging in $2X$ for Y .

To put it another way, let I be the indicator of the event $Y = 2X$, so that $E(Y \mid Y = 2X) = E(2X \mid I = 1)$. If we know that X is independent of I , then we can drop the condition $I = 1$. But in fact we have just proven that X and I can't be independent: if they were, we'd have a paradox! Surprisingly, observing X gives information about whether X is the bigger value or the smaller value. If we learn that X is very large, we might guess that X is larger than Y , but what is considered very large? Is 10^{12} very large, even though it is tiny compared with 10^{100} ? The two-envelope paradox says that no matter what the distribution of X is, there are reasonable ways to define "very large" relative to that distribution.

In Exercise 8 you will look at a related problem, in which the amounts of money in the two envelopes are i.i.d. random variables. You'll show that if you are allowed to look inside one of the envelopes and then decide whether to switch, there is a strategy that allows you to get the better envelope more than 50% of the time!

The next example vividly illustrates the importance of conditioning on all the information. The phenomenon revealed here arises in many real-life decisions about what to buy and what investments to make.

Example 9.1.7 (Mystery prize). You are approached by another stranger, who gives you an opportunity to bid on a mystery box containing a mystery prize! The value of the prize is completely unknown, except that it is worth at least nothing, and at most a million dollars. So the true value V of the prize is considered to be Uniform on $[0, 1]$ (measured in millions of dollars).

You can choose to bid any amount b (in millions of dollars). You have the chance to get the prize for considerably less than it is worth, but you could also lose money if you bid too much.

Specifically, if $b < 2V/3$, then the bid is rejected and nothing is gained or lost. If $b \geq 2V/3$, then the bid is accepted and your net payoff is $V - b$ (since you pay b to get a prize worth V). What is your optimal bid b , to maximize the expected payoff?

Solution:

Your bid $b \geq 0$ must be a predetermined constant (not based on V , since V is




FIGURE 9.2 When bidding on an unknown asset, beware the winner's curse, and condition on the relevant information.

unknown!). To find the expected payoff W , condition on whether the bid is accepted. The payoff is $V - b$ if the bid is accepted and 0 if the bid is rejected. So

$$\begin{aligned} E(W) &= E(W \mid b \geq 2V/3)P(b \geq 2V/3) + E(W \mid b < 2V/3)P(b < 2V/3) \\ &= E(V - b \mid b \geq 2V/3)P(b \geq 2V/3) + 0 \\ &= (E(V \mid V \leq 3b/2) - b)P(V \leq 3b/2). \end{aligned}$$

For $b \geq 2/3$, the event $V \leq 3b/2$ has probability 1, so the right-hand side is $1/2 - b$, which is negative. Now assume $b < 2/3$. Then $V \leq 3b/2$ has probability $3b/2$. Given that $V \leq 3b/2$, the conditional distribution of V is Uniform on $[0, 3b/2]$. Therefore,

$$E(W) = (E(V \mid V \leq 3b/2) - b)P(V \leq 3b/2) = (3b/4 - b)(3b/2) = -3b^2/8.$$

The above expression is negative except at $b = 0$, so the optimal bid is 0 : you shouldn't play this game!

Alternatively, condition on which of the following events occurs: $A = \{V < b/2\}$, $B = \{b/2 \leq V \leq 3b/2\}$, $C = \{V > 3b/2\}$. We have

$$\begin{aligned} E(W \mid A) &= E(V - b \mid A) < E(b/2 - b \mid A) = -b/2 \leq 0, \\ E(W \mid B) &= E\left(\frac{b/2 + 3b/2}{2} - b \mid B\right) = 0, \\ E(W \mid C) &= 0, \end{aligned}$$

so we should just set $b = 0$ and walk away.

The moral of this story is to condition on all the information. It is crucial in the above calculation to use $E(V \mid V \leq 3b/2)$ rather than $E(V) = 1/2$; knowing that the bid was accepted gives information about how much the mystery prize is worth, so we shouldn't destroy that information. This problem is related to the so-called winner's curse, which says that the winner

in an auction with incomplete information tends to profit less than they expect (unless they understand probability!). This is because in many settings, the expected value of the item that they bid on given that they won the bid is less than the unconditional expected value they originally had in mind. For $b \geq 2/3$, conditioning on $V \leq 3b/2$ does nothing since we know in advance that $V \leq 1$, but such a bid is ludicrously high. For any $b < 2/3$, finding out that your bid is accepted lowers your expectation:

$$E(V \mid V \leq 3b/2) < E(V)$$

The remaining examples use first-step analysis to calculate unconditional expectations. First, as promised in Chapter 4, we derive the expectation of the Geometric distribution using first-step analysis.

Example 9.1.8 (Geometric expectation redux). Let $X \sim \text{Geom}(p)$. Interpret X as the number of Tails before the first Heads in a sequence of coin flips with probability p of Heads. To get $E(X)$, we condition on the outcome of the first toss: if it lands Heads, then X is 0 and we're done; if it lands Tails, then we've wasted one toss and are back to where we started, by memorylessness. Therefore,

$$\begin{aligned} E(X) &= E(X \mid \text{first toss } H) \cdot p + E(X \mid \text{first toss } T) \cdot q \\ &= 0 \cdot p + (1 + E(X)) \cdot q \end{aligned}$$

which gives $E(X) = q/p$.

The next example derives expected waiting times for some more complicated patterns, using two steps of conditioning.

Example 9.1.9 (Time until HH vs. HT). You toss a fair coin repeatedly. What is the expected number of tosses until the pattern HT appears for the first time? What about the expected number of tosses until HH appears for the first time?

Solution:

Let W_{HT} be the number of tosses until HT appears. As we can see from Figure 9.3, W_{HT} is the waiting time for the first Heads, which we'll call W_1 , plus the additional waiting time for the first Tails after the first Heads, which we'll call W_2 . By the story of the First Success distribution, W_1 and W_2 are i.i.d. $\text{FS}(1/2)$, so $E(W_1) = E(W_2) = 2$ and $E(W_{HT}) = 4$.

Finding the expected waiting time for HH , $E(W_{HH})$, is more complicated. We can't apply the same logic as for $E(W_{HT})$: as shown in Figure 9.4, if the first Heads is immediately followed by Tails, our progress is destroyed and we must start from scratch. But this is progress for us in solving the problem, since the fact that the




FIGURE 9.3 Waiting time for HT is the waiting time for the first Heads, W_1 , plus the additional waiting time for the next Tails, W_2 . Durable partial progress is possible!

system can get reset suggests the strategy of first-step analysis. Let's condition on the outcome of the first toss:

$$E(W_{HH}) = E(W_{HH} \mid \text{first toss } H) \frac{1}{2} + E(W_{HH} \mid \text{first toss } T) \frac{1}{2}.$$




FIGURE 9.4 When waiting for HH , partial progress can easily be destroyed. For the second term, $E(W_{HH} \mid \text{first toss } T) = 1 + E(W_{HH})$ by memorylessness. For the first term, we compute $E(W_{HH} \mid 1 \text{st toss } H)$ by further conditioning on the outcome of the second toss. If the second toss is Heads, we have obtained HH in two tosses. If the second toss is Tails, we've wasted two tosses and have to start all over! This gives

$$E(W_{HH} \mid \text{first toss } H) = 2 \cdot \frac{1}{2} + (2 + E(W_{HH})) \cdot \frac{1}{2}.$$

Therefore,

$$E(W_{HH}) = \left(2 \cdot \frac{1}{2} + (2 + E(W_{HH})) \cdot \frac{1}{2} \right) \frac{1}{2} + (1 + E(W_{HH})) \frac{1}{2}.$$

Solving for $E(W_{HH})$, we get $E(W_{HH}) = 6$.

It might seem surprising at first that the expected waiting time for HH is greater than the expected waiting time for HT . How do we reconcile this with the fact that in two tosses of the coin, HH and HT both have a $1/4$ chance of appearing? Why aren't the average waiting times the same by symmetry?

As we solved this problem, we in fact noticed an important asymmetry. When waiting for HT , once we get the first Heads, we've achieved partial progress that cannot be destroyed: if the Heads is followed by another Heads, we're in the same position as before, and if the Heads is followed by a Tails, we're done. By contrast, when waiting for HH , even after getting the first Heads, we could be sent back to square one if the Heads is followed by a Tails. This suggests the average waiting time for HH should be longer. Symmetry implies that the average waiting time for HH is the same as that for TT , and that for HT is the same as that for TH , but it does not imply that the average waiting times for HH and HT are the same.

More intuition into what's going on can be obtained by considering a long string of coin flips, as in Figure 9.5. We notice right away that appearances of HH can overlap, while appearances of HT must be disjoint. For example, $HHHHHH$ has 5 occurrences of HH , but $HTHTHT$ has only 3 occurrences of HT . Since there are the same average number of HH s and HT s, but HH s sometimes clump together, the average waiting time for HH must be larger than that of HT to compensate.

(add a proper figure)

HHTHHTTHHHHTHTHTTHTT HHTHHT HHHHTHTHTTHTT

FIGURE 9.5 Clumping. (a) Appearances of HH can overlap. (b) Appearances of HT must be disjoint.

Related problems occur in information theory when compressing a message, and in genetics when looking for recurring patterns (called motifs) in DNA sequences. $\square$

Our final example in this section uses wishful thinking for both probabilities and expectations to study a question about a random walk.

Example 9.1.10 (Random walk on the integers). An immortal drunk man wanders around randomly on the integers. He starts at the origin, and at each step he moves 1 unit to the right or 1 unit to the left, with equal probabilities, independently of all his previous steps. Let b be a googolplex (this is 10^g , where $g = 10^{100}$ is a googol).

(a) Find a simple expression for the probability that the immortal drunk visits b before returning to the origin for the first time.

(b) Find the expected number of times that the immortal drunk visits b before returning to the origin for the first time.

Solution:

(a) Let B be the event that the drunk man visits b before returning to the origin for the first time and let L be the event that his first move is to the left. Then

$P(B | L) = 0$ since any path from -1 to b must pass through 0 . For $P(B | L^c)$, we are exactly in the setting of the gambler's ruin problem, where player A starts with \$1, player B starts with $$(b - 1)$, and the rounds are fair. Applying that result, we have$

$$P(B) = P(B | L)P(L) + P(B | L^c)P(L^c) = \frac{1}{b} \cdot \frac{1}{2} = \frac{1}{2b}$$

(b) Let N be the number of visits to b before returning to the origin for the first time, and let $p = 1/(2b)$ be the probability found in (a). Then

$$E(N) = E(N | N = 0)P(N = 0) + E(N | N \geq 1)P(N \geq 1) = pE(N | N \geq 1).$$

The conditional distribution of N given $N \geq 1$ is $\text{FS}(p)$: given that the man reaches b , by symmetry there is probability p of returning to the origin before visiting b again (call this "success") and probability $1 - p$ of returning to b again before returning to the origin (call this "failure"). Note that the trials are independent since the situation is the same each time he is at b , independent of the past history. Thus $E(N | N \geq 1) = 1/p$, and

$$E(N) = pE(N | N \geq 1) = p \cdot \frac{1}{p} = 1$$

Surprisingly, the result doesn't depend on the value of b , and our proof didn't require knowing the value of p .

4.9.2 9.2 Conditional Expectation Given an r.v.

In this section we introduce conditional expectation given a random variable. That is, we want to understand what it means to write $E(Y | X)$ for an r.v. X . We will see that $E(Y | X)$ is a random variable that is, in a certain sense, our best prediction of Y , assuming we get to know X .

The key to understanding $E(Y | X)$ is first to understand $E(Y | X = x)$. Since $X = x$ is an event, $E(Y | X = x)$ is just the conditional expectation of Y given this event, and it can be computed using the conditional distribution of Y given $X = x$.

If Y is discrete, we use the conditional PMF $P(Y = y \mid X = x)$ in place of the unconditional PMF $P(Y = y)$:

$$E(Y \mid X = x) = \sum_y y P(Y = y \mid X = x)$$

Analogously, if Y is continuous, we use the conditional PDF $f_{Y|X}(y \mid x)$ in place of the unconditional PDF:

$$E(Y \mid X = x) = \int_{-\infty}^{\infty} y f_{Y|X}(y \mid x) dy$$

Notice that because we sum or integrate over y , $E(Y \mid X = x)$ is a function of x only. We can give this function a name, like g : let $g(x) = E(Y \mid X = x)$. We define $E(Y \mid X)$ as the random variable obtained by finding the form of the function $g(x)$, then plugging in X for x .

Definition 9.2.1 (Conditional expectation given an r.v.). Let $g(x) = E(Y \mid X = x)$. Then the conditional expectation of Y given X , denoted $E(Y \mid X)$, is defined to be the random variable $g(X)$. In other words, if after doing the experiment X crystallizes into x , then $E(Y \mid X)$ crystallizes into $g(x)$.

9.2.2. The notation in this definition sometimes causes confusion. It does not say " $g(x) = E(Y \mid X = x)$, so $g(X) = E(Y \mid X = X)$, which equals $E(Y)$ because $X = X$ is always true". Rather, we should first compute the function $g(x)$, then plug in X for x . For example, if $g(x) = x^2$, then $g(X) = X^2$. A similar biohazard is 5.3.2, about the meaning of $F(X)$ in the universality of the Uniform.

9.2.3. By definition, $E(Y \mid X)$ is a function of X , so it is a random variable. (This does not mean there are no examples where $E(Y \mid X)$ is a constant. A constant is a degenerate r.v., and a constant function of X . For example, if X and Y are independent then $E(Y \mid X) = E(Y)$, which is a constant.) Thus it makes sense to compute quantities like $E(E(Y \mid X))$ and $\text{Var}(E(Y \mid X))$, the mean and variance of the r.v. $E(Y \mid X)$. It is easy to be ensnared by category errors when working with conditional expectation, so it is important to keep in mind that conditional expectations of the form $E(Y \mid A)$ are numbers, while those of the form $E(Y \mid X)$ are random variables.

Here are some quick examples of how to calculate conditional expectation. In both examples, we don't need to do a sum or integral to get $E(Y \mid X = x)$ because a more direct approach is available.

Example 9.2.4. A stick of length 1 is broken at a point X chosen uniformly at random. Given that $X = x$, we then choose another breakpoint Y uniformly on the interval $[0, x]$. Find $E(Y \mid X)$, and its mean and variance.

Solution:

From the description of the experiment, $X \sim \text{Unif}(0, 1)$ and $Y \mid X = x \sim \text{Unif}(0, x)$. Then $E(Y \mid X = x) = x/2$, so by plugging in X for x , we have

$$E(Y \mid X) = X/2$$

The expected value of $E(Y \mid X)$ is

$$E(E(Y \mid X)) = E(X/2) = 1/4$$

(We will show in the next section that a general property of conditional expectation is that $E(E(Y \mid X)) = E(Y)$, so it also follows that $E(Y) = 1/4$.) The variance of $E(Y \mid X)$ is

$$\text{Var}(E(Y | X)) = \text{Var}(X/2) = 1/48$$

Example 9.2.5. For $X, Y \stackrel{\text{i.i.d.}}{\sim} \text{Expo}(\lambda)$, find $E(\max(X, Y) | \min(X, Y))$.

Solution:

Let $M = \max(X, Y)$ and $L = \min(X, Y)$. By the memoryless property, $M - L$ is independent of L , and $M - L \sim \text{Expo}(\lambda)$ (see Example 7.3.6). Therefore

$$E(M | L = l) = E(L | L = l) + E(M - L | L = l) = l + E(M - L) = l + \frac{1}{\lambda},$$

and $E(M | L) = L + \frac{1}{\lambda}$.

4.9.3 9.3 Properties of Conditional Expectation

Conditional expectation has some very useful properties.

- Dropping what's independent: If X and Y are independent, then $E(Y | X) = E(Y)$.
- Taking out what's known: For any function h , $E(h(X)Y | X) = h(X)E(Y | X)$.
- Linearity: $E(Y_1 + Y_2 | X) = E(Y_1 | X) + E(Y_2 | X)$, and $E(cY | X) = cE(Y | X)$ for c a constant (the latter is a special case of taking out what's known).
- Adam's law: $E(E(Y | X)) = E(Y)$.
- Projection interpretation: The r.v. $Y - E(Y | X)$, which is called the residual from using X to predict Y , is uncorrelated with $h(X)$ for any function h .

Let's discuss each property individually.

Theorem 9.3.1 (Dropping what's independent). If X and Y are independent, then $E(Y | X) = E(Y)$.

This is true because independence implies $E(Y | X = x) = E(Y)$ for all x , hence $E(Y | X) = E(Y)$. Intuitively, if X provides no information about Y , then our best guess for Y , even if we get to know X , is still the unconditional mean $E(Y)$. However, the converse is false: a counterexample is given in Example 9.3.3 below.

Theorem 9.3.2 (Taking out what's known). For any function h ,

$$E(h(X)Y | X) = h(X)E(Y | X).$$

Intuitively, when we take expectations given X , we are treating X as if it has crystallized into a known constant. Then any function of X , say $h(X)$, also acts like a known constant while we are conditioning on X . Taking out what's known is the conditional version of the unconditional fact that $E(cY) = cE(Y)$. The difference is that $E(cY) = cE(Y)$ asserts that two numbers are equal, while taking out what's known asserts that two random variables are equal.

Example 9.3.3. Let $Z \sim \mathcal{N}(0, 1)$ and $Y = Z^2$. Find $E(Y | Z)$ and $E(Z | Y)$.

Solution: Since Y is a function of Z , $E(Y | Z) = E(Z^2 | Z) = Z^2$ by taking out what's known. To get $E(Z | Y)$, notice that conditional on $Y = y$, Z equals $\sqrt{y}$ or $-\sqrt{y}$ with equal probabilities by the symmetry of the standard Normal, so $E(Z | Y = y) = 0$ and $E(Z | Y) = 0$.

In this case, although Y provides a lot of information about Z , narrowing down the possible values of Z to just two values, Y only tells us about the magnitude of Z and not its sign.

For this reason, $E(Z | Y) = E(Z)$ despite the dependence between Z and Y . This example illustrates that the converse of Theorem 9.3.1 is false.

Theorem 9.3.4 (Linearity). $E(Y_1 + Y_2 | X) = E(Y_1 | X) + E(Y_2 | X)$.

This result is the conditional version of the unconditional fact that $E(Y_1 + Y_2) = E(Y_1) + E(Y_2)$, and is true since conditional probabilities are probabilities.

9.3.5. It is incorrect to write " $E(Y | X_1 + X_2) = E(Y | X_1) + E(Y | X_2)$ "; linearity applies on the left side of the conditioning bar, not on the right side!

Example 9.3.6. Let $X_1, \dots, X_n$ be i.i.d., and $S_n = X_1 + \dots + X_n$. Find $E(X_1 | S_n)$.

Solution:

By symmetry,

$$E(X_1 | S_n) = E(X_2 | S_n) = \dots = E(X_n | S_n),$$

and by linearity,

$$E(X_1 | S_n) + \dots + E(X_n | S_n) = E(S_n | S_n) = S_n.$$

Therefore,

$$E(X_1 | S_n) = S_n/n = \bar{X}_n,$$

the sample mean of the X_j 's. This is an intuitive result: if we have 2 i.i.d. r.v.s X_1, X_2 and learn that $X_1 + X_2 = 10$, it makes sense to guess that X_1 is 5 (accounting for half of the total). Similarly, if we have n i.i.d. r.v.s and get to know their sum, our best guess for any one of them is the sample mean.

The next theorem connects conditional expectation to unconditional expectation. It goes by many names, including the law of total expectation, the law of iterated expectation (which has a terrible acronym for something glowing with truth), and the tower property. We call it Adam's law because it is used so frequently that it deserves a pithy name, and since it is often used in conjunction with another law we'll encounter soon, which has a complementary name.

Theorem 9.3.7 (Adam's law). For any r.v.s X and Y ,

$$E(E(Y | X)) = E(Y)$$

Proof. We present the proof in the case where X and Y are both discrete (the proofs for other cases are analogous). Let $E(Y | X) = g(X)$. We proceed by applying

LOTUS, expanding the definition of $g(x)$ to get a double sum, and then swapping the order of summation:

$$\begin{aligned} E(g(X)) &= \sum_x g(x)P(X = x) \\ &= \sum_x \left(\sum_y yP(Y = y | X = x) \right) P(X = x) \\ &= \sum_x \sum_y yP(X = x)P(Y = y | X = x) \\ &= \sum_y y \sum_x P(X = x, Y = y) \\ &= \sum_y yP(Y = y) = E(Y). \end{aligned}$$

Adam's law is a more compact, more general version of the law of total expectation (Theorem 9.1.5). For X discrete, the statements

$$E(Y) = \sum_x E(Y | X = x)P(X = x)$$

and

$$E(Y) = E(E(Y | X))$$

mean the same thing, since if we let $E(Y | X = x) = g(x)$, then

$$E(E(Y | X)) = E(g(X)) = \sum_x g(x)P(X = x) = \sum_x E(Y | X = x)P(X = x).$$

Armed with Adam's law, we have a powerful strategy for finding an expectation $E(Y)$, by conditioning on an r.v. X that we wish we knew. First obtain $E(Y | X)$ by treating X as known, and then take the expectation of $E(Y | X)$. We will see various examples of this later in the chapter.

Just as there are forms of Bayes' rule and LOTP with extra conditioning, as discussed in Chapter 2, there is a version of Adam's law with extra conditioning.

Theorem 9.3.8 (Adam's law with extra conditioning). For any r.v.s X, Y, Z ,

$$E(E(Y | X, Z) | Z) = E(Y | Z)$$

The above equation is Adam's law, except with extra conditioning on Z inserted everywhere. It is true because conditional probabilities are probabilities. So we are free to use Adam's law to help us find both unconditional expectations and conditional expectations.

Using Adam's law, we can also prove the last item on our list of properties of conditional expectation.

Theorem 9.3.9 (Projection interpretation). For any function h , the random variable $Y - E(Y | X)$ is uncorrelated with $h(X)$. Equivalently,

$$E((Y - E(Y | X))h(X)) = 0.$$

(This is equivalent since $E(Y - E(Y | X)) = 0$, by linearity and Adam's law.)

Proof. We have

$$\begin{aligned} E((Y - E(Y | X))h(X)) &= E(h(X)Y) - E(h(X)E(Y | X)) \\ &= E(h(X)Y) - E(E(h(X)Y | X)) \end{aligned}$$

by Theorem 9.3.2 (here we're "putting back what's known" in the inner expectation). By Adam's law, the second term is equal to $E(h(X)Y)$. $\square$




FIGURE 9.6 The conditional expectation $E(Y | X)$ is the projection of Y onto the space of all functions of X , shown here as a plane. The residual $Y - E(Y | X)$ is orthogonal to the plane: it's perpendicular to (uncorrelated with) any function of X .

From a geometric perspective, we can visualize Theorem 9.3.9 as in Figure 9.6. In a certain sense (described below), $E(Y | X)$ is the function of X that is closest to Y ; we say that $E(Y | X)$ is the projection of Y into the space of all functions of X . The "line" from Y to $E(Y | X)$ in the figure is orthogonal (perpendicular) to the "plane", since any other route from Y to $E(Y | X)$ would be longer. This orthogonality turns out to be the geometric interpretation of Theorem 9.3.9.

The details of this perspective are given in the next section, which is starred since it requires knowledge of linear algebra. But even without delving into the linear algebra, the projection picture gives some useful intuition. As mentioned earlier, we can think of $E(Y | X)$ as a prediction for Y based on X . This is an extremely common problem in statistics: predict or estimate the future observations or unknown parameters based on data. The projection interpretation of conditional expectation implies that $E(Y | X)$ is the best predictor of Y based on X , in the sense that it is the function of X with the lowest mean squared error (expected squared difference between Y and the prediction of Y).

Example 9.3.10 (Linear regression). An extremely widely used method for data analysis in statistics is linear regression. In its most basic form, the linear regression model uses a single explanatory variable X to predict a response variable Y , and it assumes that the conditional expectation of Y is linear in X :

$$E(Y | X) = a + bX.$$

(a) Show that an equivalent way to express this is to write

$$Y = a + bX + \epsilon,$$

where ϵ is an r.v. (called the error) with $E(\epsilon | X) = 0$.

(b) Solve for the constants a and b in terms of $E(X)$, $E(Y)$, $\text{Cov}(X, Y)$, and $\text{Var}(X)$.

Solution:

(a) Let $Y = a + bX + \epsilon$, with $E(\epsilon | X) = 0$. Then by linearity,

$$E(Y | X) = E(a | X) + E(bX | X) + E(\epsilon | X) = a + bX.$$

Conversely, suppose that $E(Y | X) = a + bX$, and define

$$\epsilon = Y - (a + bX).$$

Then $Y = a + bX + \epsilon$, with

$$E(\epsilon | X) = E(Y | X) - E(a + bX | X) = E(Y | X) - (a + bX) = 0.$$

(b) First, by Adam's law, taking the expectation of both sides gives

$$E(Y) = a + bE(X).$$

Note that ϵ has mean 0 and X and ϵ are uncorrelated, since

$$E(\epsilon) = E(E(\epsilon | X)) = E(0) = 0$$

and

$$E(\epsilon X) = E(E(\epsilon X | X)) = E(X E(\epsilon | X)) = E(0) = 0.$$

Taking the covariance with X of both sides in $Y = a + bX + \epsilon$, we have

$$\text{Cov}(X, Y) = \text{Cov}(X, a) + b \text{Cov}(X, X) + \text{Cov}(X, \epsilon) = b \text{Var}(X)$$

Thus,

$$b = \frac{\text{Cov}(X, Y)}{\text{Var}(X)},$$

$$a = E(Y) - bE(X) = E(Y) - \frac{\text{Cov}(X, Y)}{\text{Var}(X)} \cdot E(X).$$

4.9.4 9.4 *geometric Interpretation of Conditional Expectation

This section explains in more detail the geometric perspective shown in Figure 9.6, using some concepts from linear algebra. Consider the vector space consisting of all random variables on a certain probability space, such that the random variables all have finite variance. Each vector or point in the space is a random variable (here we are using "vector" in the linear algebra sense, not in the sense of a random vector from Chapter 7). Define the inner product of two r.v.s U and V to be

$$\langle U, V \rangle = E(UV)$$

(For this definition to satisfy the axioms for an inner product, we need the convention that two r.v.s are considered the same if they are equal with probability 1.)

The squared length of an r.v. X is

$$\|X\|^2 = \langle X, X \rangle = EX^2$$

and the squared distance between two r.v.s U and V is

$$\|U - V\|^2 = E(U - V)^2.$$

The interpretations become especially nice if $E(U) = E(V) = 0$, since then:

- $\|U\|^2 = \text{Var}(U)$, and $\|U\| = \text{SD}(U)$.
- $\langle U, V \rangle = \text{Cov}(U, V)$, and the cosine of the "angle" between U and V is $\text{Corr}(U, V)$.
- U and V are orthogonal (i.e., $\langle U, V \rangle = 0$) if and only if they are uncorrelated.

To interpret $E(Y | X)$ geometrically, consider the space of all random variables (with finite variance) that can be expressed as functions of X . This is a subspace of the vector space. In Figure 9.6, the subspace of random variables of the form $h(X)$ is represented by a plane. To get $E(Y | X)$, we project Y onto the plane. Then the residual $Y - E(Y | X)$ is orthogonal to $h(X)$ for all functions h , and $E(Y | X)$ is the function of X that best predicts Y , where "best" here means that the mean squared error $E(Y - g(X))^2$ is minimized by choosing $g(X) = E(Y | X)$.

The projection interpretation is a helpful way to think about many of the properties of conditional expectation. For example, if $Y = h(X)$ is a function of X , then Y itself is already in the plane, so it is its own projection; this explains why

$$E(h(X) | X) = h(X).$$

We can think of unconditional expectation as a projection too: $E(Y)$ can be thought of as $E(Y | 0)$, the projection of Y onto the space of all constants (and indeed, $E(Y)$ is the constant c that minimizes $E(Y - c)^2$, as we proved in Theorem 6.1.4).

We can now also give a geometric interpretation for Adam's law: $E(Y)$ says to project Y in one step onto the space of all constants; $E(E(Y | X))$ says to do it in two steps, by first projecting onto the plane and then projecting $E(Y | X)$ onto the space of all constants, which is a line within that plane. Adam's law says that the one-step and two-step methods yield the same result.

In the next section we will introduce Eve's law, which serves the same purpose for variance as Adam's law does for expectation. As a preview and to further explore the geometric interpretation of conditional expectation, let's look at $\text{Var}(Y)$ from the perspective of this section. Assume that $E(Y) = 0$ (if $E(Y) \neq 0$, we can center Y by subtracting $E(Y)$; doing so has no effect on the variance of Y).

We can decompose Y into two orthogonal terms, the residual $Y - E(Y | X)$ and the conditional expectation $E(Y | X)$:

$$Y = (Y - E(Y | X)) + E(Y | X).$$

The two terms are orthogonal since $Y - E(Y | X)$ is uncorrelated with any function of X , and $E(Y | X)$ is a function of X . So by the Pythagorean theorem,

$$\|Y\|^2 = \|Y - E(Y | X)\|^2 + \|E(Y | X)\|^2.$$

That is,

$$\text{Var}(Y) = \text{Var}(Y - E(Y | X)) + \text{Var}(E(Y | X)).$$

As we will see in the next section, this identity is a form of Eve's law. So it turns out that Eve's law, which may look cryptic at first glance, can be interpreted as just being the Pythagorean theorem for a "triangle" whose sides are the vectors $Y - E(Y | X)$, $E(Y | X)$, and Y .

4.9.5 9.5 Conditional Variance

Once we've defined conditional expectation given an r.v., we have a natural way to define conditional variance given a random variable: replace all instances of $E(\cdot)$ in the definition of unconditional variance with $E(\cdot | X)$.

Definition 9.5.1 (Conditional variance). The conditional variance of Y given X is

$$\text{Var}(Y | X) = E((Y - E(Y | X))^2 | X).$$

This is equivalent to

$$\text{Var}(Y | X) = E(Y^2 | X) - (E(Y | X))^2$$

9.5.2. Like $E(Y | X)$, $\text{Var}(Y | X)$ is a random variable, and it is a function of X .

Since conditional variance is defined in terms of conditional expectations, we can use results about conditional expectation to help us calculate conditional variance. Here's an example.

Example 9.5.3. Let $Z \sim \mathcal{N}(0, 1)$ and $Y = Z^2$. Find $\text{Var}(Y | Z)$ and $\text{Var}(Z | Y)$.

Solution:

Without any calculations we can see that $\text{Var}(Y | Z) = 0$: conditional on Z , Y is a known constant, and the variance of a constant is 0 . By the same reasoning, $\text{Var}(h(Z) | Z) = 0$ for any function h .

To get $\text{Var}(Z | Y)$, apply the definition:

$$\text{Var}(Z | Z^2) = E(Z^2 | Z^2) - (E(Z | Z^2))^2.$$

The first term equals Z^2 . The second term equals 0 by symmetry, as we found in Example 9.3.3. Thus $\text{Var}(Z | Z^2) = Z^2$, which we can write as $\text{Var}(Z | Y) = Y$.

In the next example, we will practice working with conditional expectation and conditional variance in the context of the Bivariate Normal.

Example 9.5.4 (Conditional expectation and conditional variance in a BVN). Let (Z, W) be Bivariate Normal, with $\text{Corr}(Z, W) = \rho$ and Z, W marginally $\mathcal{N}(0, 1)$. Find $E(W | Z)$ and $\text{Var}(W | Z)$.

Solution: We can assume that (Z, W) has been constructed as in Example 7.5.10, since $E(W | Z)$ and $\text{Var}(W | Z)$ depend only on the joint distribution of (Z, W) , not on the specific method that was used to create (Z, W) . So let

$$\begin{aligned} Z &= X \\ W &= \rho X + \sqrt{1 - \rho^2} Y, \end{aligned}$$

with X, Y i.i.d. $\mathcal{N}(0, 1)$. We can then solve the problem very neatly, without having to resort to messy integrals based on the Bivariate Normal joint PDF. The conditional expectation is

$$E(W | Z) = E(W | X) = \rho X + \sqrt{1 - \rho^2} E(Y | X) = \rho X + \sqrt{1 - \rho^2} E(Y) = \rho Z$$

since X and Y are independent. And the conditional variance is

$$\text{Var}(W | Z) = \text{Var}(W | X) = \text{Var}(\sqrt{1 - \rho^2} Y | X) = (1 - \rho^2) \text{Var}(Y) = 1 - \rho^2,$$

since ρX acts as a constant if we are given X , and Y is independent of X . Interestingly, the same argument with the roles of Z and W reversed shows that

$$E(Z | W) = \rho W, \text{ and } \text{Var}(Z | W) = 1 - \rho^2.$$

One might have guessed that if we should multiply by ρ to go from an observed value of Z to a predicted value of W , then we should divide by ρ to go from an observed value of W to a predicted value of Z . But the above results say to multiply by the same ρ , regardless of whether using Z to predict W or vice versa! This is closely related to the fact that correlation is symmetric ($\text{Corr}(Z, W) = \rho = \text{Corr}(W, Z)$) and to an important concept in statistics known as regression toward the mean.

We learned in the previous section that Adam's law relates conditional expectation to unconditional expectation. A companion result for Adam's law is Eve's law, which relates conditional variance to unconditional variance.

Theorem 9.5.5 (Eve's law). For any r.v.s X and Y ,

$$\text{Var}(Y) = E(\text{Var}(Y | X)) + \text{Var}(E(Y | X)).$$

The ordering of E 's and Var 's on the right-hand side spells EVVE, whence the name Eve's law. Eve's law is also known as the law of total variance or the variance decomposition formula.

Proof. Let $g(X) = E(Y | X)$. By Adam's law, $E(g(X)) = E(Y)$. Then

$$\begin{aligned} E(\text{Var}(Y | X)) &= E(E(Y^2 | X) - g(X)^2) = E(Y^2) - E(g(X)^2), \\ \text{Var}(E(Y | X)) &= E(g(X)^2) - (Eg(X))^2 = E(g(X)^2) - (EY)^2. \end{aligned}$$

Adding these equations, we have Eve's law. $\square$

To visualize Eve's law, imagine a population where each person has a value of X and a value of Y . We can divide this population into subpopulations, one for each possible value of X . For example, if X represents age and Y represents height, we can group people based on age. Then there are two sources contributing to the variation in people's heights in the overall population. First, within each age group, people have different heights. The average amount of variation in height within each age group is the within-group variation, $E(\text{Var}(Y | X))$. Second, across age groups, the average heights are different. The variance of average heights across age groups is the between-group variation, $\text{Var}(E(Y | X))$. Eve's law says that to get the total variance of Y , we simply add these two sources of variation.

Figure 9.7 illustrates Eve's law in the simple case where we have three age groups. The average amount of scatter within each of the groups $X = 1$, $X = 2$, and $X = 3$ is the within-group variation, $E(\text{Var}(Y | X))$. The variance of the group means $E(Y | X = 1)$, $E(Y | X = 2)$, and $E(Y | X = 3)$ is the between-group variation, $\text{Var}(E(Y | X))$.

Another way to think about Eve's law is in terms of prediction. If we wanted to predict someone's height based on their age alone, the ideal scenario would be if everyone within an age group had exactly the same height, while different age groups had different heights. Then, given someone's age, we would be able to predict their height perfectly. In other words, the ideal scenario for prediction is no within-group




FIGURE 9.7

Eve's law says that total variance is the sum of within-group and between-group variation. variation in height, since the within-group variation cannot be explained by age differences. For this reason, within-group variation is also called unexplained variation, and between-group variation is also called explained variation. Eve's law says that the overall variance of Y is the sum of unexplained and explained variation.

We can also write Eve's law in the form

$$\text{Var}(Y) = \text{Var}(Y - E(Y | X)) + \text{Var}(E(Y | X)),$$

since, letting W be the residual $Y - E(Y | X)$,

$$\text{Var}(Y - E(Y | X)) = E(W^2) = E(E(W^2 | X)) = E(\text{Var}(Y | X)).$$

Again this says that we can decompose variance into within-group variation plus between-group variation.

9.5.6. Let Y be an r.v. and A be an event. It is wrong to say " $\text{Var}(Y) = \text{Var}(Y | A)P(A) +$

$\text{Var}(Y | A^c) P(A^c)$, even though this looks analogous to the law of total expectation. (For a simple counterexample, let $Y \sim \text{Bern}(1/2)$ and A be the event $Y = 0$. Then $\text{Var}(Y | A)$ and $\text{Var}(Y | A^c)$ are both 0, but $\text{Var}(Y) = 1/4$.)

Instead, we should use Eve's law if we want to condition on whether or not A occurred: letting I be the indicator of A ,

$$\text{Var}(Y) = E(\text{Var}(Y | I)) + \text{Var}(E(Y | I)).$$

To see how this expression relates to the "wrong expression", let $p = P(A)$, $q = P(A^c)$, $a = E(Y | A)$, $b = E(Y | A^c)$, $v = \text{Var}(Y | A)$, $w = \text{Var}(Y | A^c)$. Then $E(Y | I)$ is a with probability p and b with probability q , and $\text{Var}(Y | I)$ is v with probability p and w with probability q . So

$$E(\text{Var}(Y | I)) = vp + wq = \text{Var}(Y | A)P(A) + \text{Var}(Y | A^c)P(A^c),$$

which is exactly the "wrong expression", and $\text{Var}(Y)$ consists of this plus the term

$$\text{Var}(E(Y | I)) = a^2p + b^2q - (ap + bq)^2.$$

It is crucial to account for both within-group and between-group variation.

4.9.6 9.6 Adam and Eve Examples

We conclude this chapter with several examples showing how Adam's law and Eve's law allow us to find the mean and variance of complicated r.v.s, especially in situations that involve multiple levels of randomness.

In our first example, the r.v. of interest is a random sum: the sum of a random number of random variables. There are thus two levels of randomness: first, each term in the sum is a random variable; second, the number of terms in the sum is also a random variable.

Example 9.6.1 (Random sum). A store receives N customers in a day, where N is an r.v. with finite mean and variance. Let X_j be the amount spent by the j th customer at the store. Assume that each X_j has mean μ and variance σ^2 , and that N and all the X_j are independent of one another. Find the mean and variance of the random sum $X = \sum_{j=1}^N X_j$, which is the store's total revenue in a day, in terms of μ , σ^2 , $E(N)$, and $\text{Var}(N)$.

Solution:

Since X is a sum, our first impulse might be to claim " $E(X) = N\mu$ by linearity". Alas, this would be a category error, since $E(X)$ is a number and $N\mu$ is a random variable. The key is that X is not merely a sum, but a random sum; the number of terms we are adding up is itself random, whereas linearity applies to sums with a fixed number of terms.

Yet this category error actually suggests the correct strategy: if only we were allowed to treat N as a constant, then linearity would apply. So let's condition on N . By linearity of conditional expectation,

$$E(X | N) = E\left(\sum_{j=1}^N X_j | N\right) = \sum_{j=1}^N E(X_j | N) = \sum_{j=1}^N E(X_j) = N\mu.$$

We used the independence of the X_j and N to assert $E(X_j | N) = E(X_j)$ for all j . Note that the statement " $E(X | N) = N\mu$ " is not a category error because both sides of the equality are r.v.s that are functions of N . Finally, by Adam's law,

$$E(X) = E(E(X | N)) = E(N\mu) = \mu E(N).$$

This is a pleasing result: the average total revenue is the average amount spent per customer, multiplied by the average number of customers.

For $\text{Var}(X)$, we again condition on N to get $\text{Var}(X | N)$:

$$\text{Var}(X | N) = \text{Var}\left(\sum_{j=1}^N X_j | N\right) = \sum_{j=1}^N \text{Var}(X_j | N) = \sum_{j=1}^N \text{Var}(X_j) = N\sigma^2.$$

Eve's law then tells us how to obtain the unconditional variance of X :

$$\begin{aligned} \text{Var}(X) &= E(\text{Var}(X | N)) + \text{Var}(E(X | N)) \\ &= E(N\sigma^2) + \text{Var}(N\mu) \\ &= \sigma^2 E(N) + \mu^2 \text{Var}(N). \end{aligned}$$

In the next example, two levels of randomness arise because our experiment takes place in two stages. We sample a city from a group of cities, then sample citizens within the city. This is an example of a multilevel model.

Example 9.6.2 (Random sample from a random city). To study the prevalence of a disease in several cities of interest within a certain county, we pick a city at random, then pick a random sample of n people from that city. This is a form of a widely used survey technique known as cluster sampling.

Let Q be the proportion of diseased people in the chosen city, and let X be the number of diseased people in the sample. As illustrated in Figure 9.8 (where white dots represent healthy individuals and black dots represent diseased individuals), different cities may have very different prevalences. Since each city has its own disease prevalence, Q is a random variable. Suppose that $Q \sim \text{Unif}(0, 1)$. Also assume that conditional on Q , each individual in the sample independently has probability Q of having the disease; this is true if we sample with replacement from the chosen city, and is approximately true if we sample without replacement but the population size is large. Find $E(X)$ and $\text{Var}(X)$.




FIGURE 9.8 A certain oval-shaped country has 4 cities. Each city has healthy people (represented as white dots) and diseased people (represented as black dots). A random city is chosen, and then a random sample of n people is chosen from within that city. There are two components to the variability in the number of diseased people in the sample: variation due to different cities having different disease prevalence, and variation due to the randomness of the sample within the chosen city.

Solution:

With our assumptions, $X | Q \sim \text{Bin}(n, Q)$; this notation says that conditional on knowing the disease prevalence in the chosen city, we can treat Q as a constant, and each sampled individual is an independent Bernoulli trial with probability Q of success. Using the mean and variance of the Binomial distribution, $E(X | Q) = nQ$ and $\text{Var}(X | Q) = nQ(1 - Q)$. Furthermore, using the moments of the standard Uniform distribution, $E(Q) = 1/2$, $E(Q^2) = 1/3$, and $\text{Var}(Q) = 1/12$. Now we can apply Adam's law and Eve's law to get the unconditional mean and variance of X :

$$\begin{aligned} E(X) &= E(E(X | Q)) = E(nQ) = \frac{n}{2} \\ \text{Var}(X) &= E(\text{Var}(X | Q)) + \text{Var}(E(X | Q)) \\ &= E(nQ(1 - Q)) + \text{Var}(nQ) \\ &= nE(Q) - nE(Q^2) + n^2 \text{Var}(Q) \\ &= \frac{n}{2} - \frac{n}{3} + \frac{n^2}{12} \end{aligned}$$

Note that the structure of this problem is identical to that in the story of Bayes' billiards. Therefore, we actually know the distribution of X , not just its mean and variance: X is Discrete Uniform on $\{0, 1, 2, \dots, n\}$. But the Adam-and-Eve approach can be applied when Q has a more complicated distribution, or with more levels in the multilevel model, whether or not it is feasible to work out the distribution of X . For example, we could have people within cities within counties within states within countries.

Last but not least, we revisit Story 8.4.5, the Gamma-Poisson problem from the previous chapter.

Example 9.6.3 (Gamma-Poisson revisited). Recall that Fred decided to find out about the rate of Blotchville's Poisson process of buses by waiting at the bus stop for t hours and counting the number of buses Y . He then used the data to update his prior distribution $\lambda \sim \text{Gamma}(r_0, b_0)$. Thus, Fred was using the two-level model

$$\begin{aligned} \lambda &\sim \text{Gamma}(r_0, b_0) \\ Y | \lambda &\sim \text{Pois}(\lambda t). \end{aligned}$$

We found that under Fred's model, the marginal distribution of Y is Negative Binomial with parameters $r = r_0$ and $p = b_0 / (b_0 + t)$. In particular,

$$\begin{aligned} E(Y) &= \frac{rq}{p} = \frac{r_0 t}{b_0} \\ \text{Var}(Y) &= \frac{rq}{p^2} = \frac{r_0 t (b_0 + t)}{b_0^2}. \end{aligned}$$

Let's independently verify this with Adam's law and Eve's law. Using results about the Poisson distribution, the conditional mean and variance of Y given λ are $E(Y | \lambda) = \text{Var}(Y | \lambda) = \lambda t$. Using results about the Gamma distribution, the marginal mean and variance of λ are $E(\lambda) = r_0 / b_0$ and $\text{Var}(\lambda) = r_0 / b_0^2$. For Adam and Eve, this is all that is required:

$$\begin{aligned}
E(Y) &= E(E(Y \mid \lambda)) = E(\lambda t) = \frac{r_0 t}{b_0}, \\
\text{Var}(Y) &= E(\text{Var}(Y \mid \lambda)) + \text{Var}(E(Y \mid \lambda)) \\
&= E(\lambda t) + \text{Var}(\lambda t) \\
&= \frac{r_0 t}{b_0} + \frac{r_0 t^2}{b_0^2} = \frac{r_0 t (b_0 + t)}{b_0^2},
\end{aligned}$$

which is consistent with our earlier answers. The difference is that when using Adam and Eve, we don't need to know that Y is Negative Binomial! If we had been too lazy to derive the marginal distribution of Y , or if we weren't so lucky as to have a named distribution for Y , Adam and Eve would still deliver the mean and variance of Y (though not the PMF).

Lastly, let's compare the mean and variance of Y under the two-level model to the mean and variance we would get if Fred were absolutely sure of the true value of λ . In other words, suppose we replaced λ by its mean, $E(\lambda) = r_0/b_0$, making λ a constant instead of an r.v. Then the marginal distribution of the number of buses (which we'll call $\tilde{Y}$ under the new assumptions) would just be Poisson with parameter $r_0 t/b_0$. Then we would have

$$\begin{aligned}
E(\tilde{Y}) &= \frac{r_0 t}{b_0}, \\
\text{Var}(\tilde{Y}) &= \frac{r_0 t}{b_0}.
\end{aligned}$$

Notice that $E(\tilde{Y}) = E(Y)$, but $\text{Var}(\tilde{Y}) < \text{Var}(Y)$: the extra term $r_0 t^2/b_0^2$ from Eve's law is missing. Intuitively, when we fix λ at its mean, we are eliminating a level of uncertainty in the model, and this causes a reduction in the unconditional variance.

Figure 9.9 overlays the plots of two PMFs, that of $Y \sim \text{NBin}(r_0, b_0/(b_0 + t))$ in gray and that of $\tilde{Y} \sim \text{Pois}(r_0 t/b_0)$ in black. The values of the parameters are arbitrarily chosen to be $r_0 = 5, b_0 = 1, t = 2$. These two PMFs have the same center of mass, but the PMF of Y is noticeably more dispersed.

4.9.7 9.7 Recap

To calculate an unconditional expectation, we can divide up the sample space and use the law of total expectation

$$E(Y) = \sum_{i=1}^n E(Y \mid A_i) P(A_i),$$




FIGURE 9.9

PMF of $Y \sim \text{NBin}(r_0, b_0)$ in gray and $\tilde{Y} \sim \text{Pois}(r_0 t / b_0)$ in black, where $r_0 = 5, b_0 = 1, t = 2$.

but we must be careful not to destroy information in subsequent steps (such as by forgetting in the midst of a long calculation to condition on something that needs to be conditioned on). In problems with a recursive structure, we can also use first-step analysis for expectations.

The conditional expectation $E(Y | X)$ and conditional variance $\text{Var}(Y | X)$ are random variables that are functions of X ; they are obtained by treating X as if it were a known constant. If X and Y are independent, then $E(Y | X) = E(Y)$ and $\text{Var}(Y | X) = \text{Var}(Y)$. Conditional expectation has the properties

$$\begin{aligned} E(h(X)Y | X) &= h(X)E(Y | X) \\ E(Y_1 + Y_2 | X) &= E(Y_1 | X) + E(Y_2 | X), \end{aligned}$$

analogous to the properties $E(cY) = cE(Y)$ and $E(Y_1 + Y_2) = E(Y_1) + E(Y_2)$ for unconditional expectation. The conditional expectation $E(Y | X)$ is also the random variable that makes the residual $Y - E(Y | X)$ uncorrelated with any function of X , which means we can interpret it geometrically as a projection.

Finally, Adam's law and Eve's law,

$$\begin{aligned} E(Y) &= E(E(Y | X)) \\ \text{Var}(Y) &= E(\text{Var}(Y | X)) + \text{Var}(E(Y | X)), \end{aligned}$$

often help us calculate $E(Y)$ and $\text{Var}(Y)$ in problems that feature multiple forms or levels of randomness.

Figure 9.10 illustrates how the number $E(Y | X = x)$ connects with the r.v. $E(Y | X)$, whose expectation is $E(Y)$ by Adam's law. Additionally, it shows how the ingredients in Eve's law are formed and come together to give a useful decomposition of $\text{Var}(Y)$ in terms of quantities that condition on X .




FIGURE 9.10

We often observe an r.v. X and want to predict another r.v. Y based on the information about X . If we observe that $X = x$, then we can condition on that event and use $E(Y | X = x)$ as our prediction. The conditional expectation $E(Y | X)$ is the r.v. that takes the value $E(Y | X = x)$ when $X = x$. Adam's law lets us compute the unconditional expectation $E(Y)$ by starting with the conditional expectation $E(Y | X)$. Similarly, Eve's law lets us compute $\text{Var}(Y)$ in terms of quantities that condition on X .

4.9.8 9.8 R

Mystery Prize Simulation

We can use simulation to show that in Example 9.1.7, the example of bidding on a mystery prize with unknown value, any bid will lead to a negative payout on average. First choose a bid b (we chose 0.6); then simulate a large number of hypothetical mystery prizes and store them in v :

```
b <- 0.6
```

```
nsim <- 10^5
```

```
v <- runif(nsim)
```

The bid is accepted if $b > (2/3) * v$. To get the average profit conditional on an accepted bid, we use square brackets to keep only those values of v satisfying the condition:

```
mean(v[b > (2/3) * v]) - b
```

This value is negative regardless of b , as you can check by experimenting with different values of b .

Time Until hH vs. hT

To verify the results of Example 9.1.9, we can start by generating a long sequence of fair coin tosses. This is done with the sample command. We use paste with the collapse="" argument to turn these tosses into a single string of H 's and T 's:

```
paste(sample(c("H","T"),100,replace=TRUE),collapse="")
```

A sequence of length 100 is enough to virtually guarantee that both HH and HT will have appeared at least once.

To determine how many tosses are required on average to see HH and HT , we need to generate many sequences of coin tosses. For this, we use our familiar friend replicate:

```
r <- replicate(10^3,paste(sample(c("H","T"),100,replace=T), collapse=""))
```

Now r contains a thousand sequences of coin tosses, each of length 100. To find the first appearance of HH in each of these sequences, you can use the `str_locate` command from the `stringr` package. After you've installed and loaded the package,

```
t <- str_locate(r,"HH")
```

creates a two-column table t , whose columns contain the starting and ending positions of the first appearance of HH in each sequence of coin tosses. (Use `head(t)` to display the first few rows of the table and get an idea of what your results look like.) What we want are the ending positions, given by the second column. In particular, we want the average value of the second column, which is an approximation of the average waiting time for HH :

```
mean(t[,2])
```

Is your answer around 6? Trying again with "HT" instead of "HH", is your answer around 4?

Linear Regression

In Example 9.3.10, we derived formulas for the slope and intercept of a linear regression model, which can be used to predict a response variable using an explanatory variable. Let's try to apply these formulas to a simulated dataset:

```
x <- rnorm(100)
```

```
y <- 3 + 5*x + rnorm(100)
```

The vector x contains 100 realizations of the random variable $X \sim \mathcal{N}(0, 1)$, and the vector y contains 100 realizations of the random variable $Y = a + bX + \epsilon$ where $\epsilon \sim \mathcal{N}(0, 1)$. As we can see, the true values of a and b for this dataset are 3 and 5, respectively. We can visualize the data as a scatterplot with `plot(x, y)`.

Now let's see if we can get good estimates of the true a and b , using the formulas in Example 9.3.10:

```
b <- cov(x,y) / var(x)
```

```
a <- mean(y) - b*mean(x)
```

Here `cov(x,y)`, `var(x)`, and `mean(x)` provide the sample covariance, sample variance, and sample mean, estimating the quantities $\text{Cov}(X, Y)$, $\text{Var}(X)$, and $E(X)$, respectively. (We have discussed sample mean and sample variance in detail in earlier chapters. Sample covariance is defined analogously, and is a natural way to estimate the true covariance.)

You should find that b is close to 5 and a is close to 3. These estimated values define the line of best fit. The `abline` command lets us plot the line of best fit on top of our scatterplot:

```
plot(x,y)
```

```
abline(a=a,b=b)
```

The first argument to `abline` is the intercept of the line, and the second argument is the slope.

4.9.9 9.9 Exercises

Exercises marked with (S) have detailed solutions at <http://stat110.net>.

Conditional Expectation Given an Event

- Fred wants to travel from Blotchville to Blissville, and is deciding between 3 options (involving different routes or different forms of transportation). The j th option would take an average of μ_j hours, with a standard deviation of σ_j hours. Fred randomly chooses between the 3 options, with equal probabilities. Let T be how long it takes for him to get from Blotchville to Blissville.
 - Find $E(T)$. Is it simply $(\mu_1 + \mu_2 + \mu_3)/3$, the average of the expectations?
 - Find $\text{Var}(T)$. Is it simply $(\sigma_1^2 + \sigma_2^2 + \sigma_3^2)/3$, the average of the variances?
- While Fred is sleeping one night, X legitimate emails and Y spam emails are sent to him. Suppose that X and Y are independent, with $X \sim \text{Pois}(10)$ and $Y \sim \text{Pois}(40)$. When he wakes up, he observes that he has 30 new emails in his inbox. Given this information, what is the expected value of how many new legitimate emails he has?
- A group of 21 women and 14 men are enrolled in a medical study. Each of them has a certain disease with probability p , independently. It is then found (through extremely reliable testing) that exactly 5 of the people have the disease. Given this information, what is the expected number of women who have the disease?
- A researcher studying crime is interested in how often people have gotten arrested. Let $X \sim \text{Pois}(\lambda)$ be the number of times that a random person got arrested in the last 10 years. However, data from police records are being used for the researcher's study, and people who were never arrested in the last 10 years do not appear in the records. In other words, the police records have a selection bias: they only contain information on people who have been arrested in the last 10 years.
 So averaging the numbers of arrests for people in the police records does not directly estimate $E(X)$; it makes more sense to think of the police records as giving us information about the conditional distribution of how many times a person was arrested, given that the person was arrested at least once in the last 10 years. The conditional distribution of X , given that $X \geq 1$, is called a truncated Poisson distribution (see Exercise 14 from Chapter 3 for another example of this distribution).
 - Find $E(X \mid X \geq 1)$
 - Find $\text{Var}(X \mid X \geq 1)$.
- A fair 20-sided die is rolled repeatedly, until a gambler decides to stop. The gambler pays \$1 per roll, and receives the amount shown on the die when the gambler stops (e.g., if the die is rolled 7 times and the gambler decides to stop then, with an 18 as the value of the last roll, then the net payoff is $\$18 - \$7 = \$11$). Suppose the gambler uses the following strategy: keep rolling until a value of m or greater is obtained, and then stop (where m is a fixed integer between 1 and 20).
 - What is the expected net payoff?

Hint: The average of consecutive integers $a, a+1, \dots, a+n$ is the same as the average of the first and last of these. See the math appendix for more information about series.

- Use R or other software to find the optimal value of m .
- Let $X \sim \text{Expo}(\lambda)$. Find $E(X \mid X < 1)$ in two different ways:
 - by calculus, working with the conditional PDF of X given $X < 1$.
 - without calculus, by expanding $E(X)$ using the law of total expectation.
 - You are given an opportunity to bid on a mystery box containing a mystery prize! The value of the prize is completely unknown, except that it is worth at least nothing, and at most a million dollars. So the true value V of the prize is considered to be Uniform on $[0, 1]$ (measured in millions of dollars).

You can choose to bid any nonnegative amount b (in millions of dollars). If $b < \frac{1}{4}V$, then your bid is rejected and nothing is gained or lost. If $b \geq \frac{1}{4}V$, then your bid is accepted and your net payoff is $V - b$ (since you pay b to get a prize worth V).

Find your expected payoff as a function of b (be sure to specify it for all $b \geq 0$). Then find the optimal bid b , to maximize your expected payoff.

8. (S) You get to choose between two envelopes, each of which contains a check for some positive amount of money. Unlike in the two-envelope paradox, it is not given that one envelope contains twice as much money as the other envelope. Instead, assume that the two values were generated independently from some distribution on the positive real numbers, with no information given about what that distribution is.

After picking an envelope, you can open it and see how much money is inside (call this value x), and then you have the option of switching. As no information has been given about the distribution, it may seem impossible to have better than a 50% chance of picking the better envelope. Intuitively, we may want to switch if x is "small" and not switch if x is "large", but how do we define "small" and "large" in the grand scheme of all possible distributions? [The last sentence was a rhetorical question.]

Consider the following strategy for deciding whether to switch. Generate a threshold $T \sim \text{Expo}(1)$, and switch envelopes if and only if the observed value x is less than the value of T . Show that this strategy succeeds in picking the envelope with more money with probability strictly greater than $1/2$.

Hint: Let t be the value of T (generated by a random draw from the $\text{Expo}(1)$ distribution). First explain why the strategy works very well if t happens to be in between the two envelope values, and does no harm in any case (i.e., there is no case in which the strategy succeeds with probability strictly less than $1/2$).

9. There are two envelopes, each of which has a check for a $\text{Unif}(0, 1)$ amount of money, measured in thousands of dollars. The amounts in the two envelopes are independent. You get to choose an envelope and open it, and then you can either keep that amount or switch to the other envelope and get whatever amount is in that envelope.

Suppose that you use the following strategy: choose an envelope and open it. If you observe U , then stick with that envelope with probability U , and switch to the other envelope with probability $1 - U$.

(a) Find the probability that you get the larger of the two amounts.

(b) Find the expected value of what you will receive.

10. Suppose n people are bidding on a mystery prize that is up for auction. The bids are to be submitted in secret, and the individual who submits the highest bid wins the prize. The i th bidder receives a signal X_i , with $X_1, \dots, X_n$ i.i.d. The value of the prize, V , is defined to be the sum of the individual bidders' signals:

$$V = X_1 + \dots + X_n.$$

This is known in economics as the wallet game: we can imagine that the n people are bidding on the total amount of money in their wallets, and each person's signal is the amount of money in their own wallet. Of course, the wallet is a metaphor; the game can also be used to model company takeovers, where each of two companies bids to take over the other, and a company knows its own value but not the value of the other company. For this problem, assume the X_i are i.i.d. $\text{Unif}(0, 1)$.

(a) Before receiving her signal, what is bidder 1's unconditional expectation for V ?

(b) Conditional on receiving the signal $X_1 = x_1$, what is bidder 1's expectation for V ?

(c) Suppose each bidder submits a bid equal to their conditional expectation for V , i.e., bidder i bids $E(V | X_i = x_i)$. Conditional on receiving the signal $X_1 = x_1$ and winning the auction, what is bidder 1's expectation for V ? Explain intuitively why this quantity is always less than

the quantity calculated in (b).

11. (S) A coin with probability p of Heads is flipped repeatedly. For (a) and (b), suppose that p is a known constant, with $0 < p < 1$.

(a) What is the expected number of flips until the pattern HT is observed?

(b) What is the expected number of flips until the pattern HH is observed?

(c) Now suppose that p is unknown, and that we use a Beta(a, b) prior to reflect our uncertainty about p (where a and b are known constants and are greater than 2). In terms of a and b , find the corresponding answers to (a) and (b) in this setting.

12. A coin with probability p of Heads is flipped repeatedly, where $0 < p < 1$. The sequence of outcomes can be divided into runs (blocks of H 's or blocks of T 's), e.g., HHHTTT-THTTTTHH becomes

HHH	TTTT	H	TTT	HH

, which has 5 runs, with lengths 3, 4, 1, 3, 2, respectively. Assume that the coin is flipped at least until the start of the third run.

(a) Find the expected length of the first run.

(b) Find the expected length of the second run.

13. A fair 6-sided die is rolled once. Find the expected number of additional rolls needed to obtain a value at least as large as that of the first roll.

14. A fair 6-sided die is rolled repeatedly.

(a) Find the expected number of rolls needed to get a 1 followed right away by a 2.

Hint: Start by conditioning on whether or not the first roll is a 1.

(b) Find the expected number of rolls needed to get two consecutive 1's.

(c) Let a_n be the expected number of rolls needed to get the same value n times in a row (i.e., to obtain a streak of n consecutive j 's for some not-specified-in-advance value of j). Find a recursive formula for a_{n+1} in terms of a_n .

Hint: Divide the time until there are $n + 1$ consecutive appearances of the same value into two pieces: the time until there are n consecutive appearances, and the rest.

(d) Find a simple, explicit formula for a_n for all $n \geq 1$. What is a_7 (numerically)?

Conditional Expectation Given a Random Variable

15. (S) Let X_1, X_2 be i.i.d., and let $\bar{X} = \frac{1}{2}(X_1 + X_2)$ be the sample mean. In many statistics problems, it is useful or important to obtain a conditional expectation given $\bar{X}$. As an example of this, find $E(w_1 X_1 + w_2 X_2 \mid \bar{X})$, where w_1, w_2 are constants with $w_1 + w_2 = 1$.

16. Let $X_1, X_2, \dots$ be i.i.d. r.v.s with mean 0, and let $S_n = X_1 + \dots + X_n$. As shown in Example 9.3.6, the expected value of the first term given the sum of the first n terms is

$$E(X_1 \mid S_n) = \frac{S_n}{n}$$

Generalize this result by finding $E(S_k \mid S_n)$ for all positive integers k and n .

17. (S) Consider a group of n roommate pairs at a college (so there are $2n$ students). Each of these $2n$ students independently decides randomly whether to take a certain course, with probability p of success (where "success" is defined as taking the course).

Let N be the number of students among these $2n$ who take the course, and let X be the number of roommate pairs where both roommates in the pair take the course. Find $E(X)$ and $E(X \mid N)$.

18. (S) Show that $E((Y - E(Y \mid X))^2 \mid X) = E(Y^2 \mid X) - (E(Y \mid X))^2$, so these two expressions for $\text{Var}(Y \mid X)$ agree.

Hint for the variance: Adding a constant (or something acting as a constant) does not affect

variance.

19. Let X be the height of a randomly chosen adult man, and Y be his father's height, where X and Y have been standardized to have mean 0 and standard deviation 1. Suppose that (X, Y) is Bivariate Normal, with $X, Y \sim \mathcal{N}(0, 1)$ and $\text{Corr}(X, Y) = \rho$.

(a) Let $y = ax + b$ be the equation of the best line for predicting Y from X (in the sense of minimizing the mean squared error), e.g., if we were to observe $X = 1.3$ then we would predict that Y is $1.3a + b$. Now suppose that we want to use Y to predict X , rather than using X to predict Y . Give and explain an intuitive guess for what the slope is of the best line for predicting X from Y .

(b) Find a constant c (in terms of ρ) and an r.v. V such that $Y = cX + V$, with V independent of X .

Hint: Start by finding c such that $\text{Cov}(X, Y - cX) = 0$.

(c) Find a constant d (in terms of ρ) and an r.v. W such that $X = dY + W$, with W independent of Y .

(d) Find $E(Y | X)$ and $E(X | Y)$.

(e) Reconcile (a) and (d), if your intuitive guess in (a) differed from what the results of (d) implied. Give a clear and correct intuitive explanation of the relationship between the slope of the best line for predicting Y from X and the slope of the best line for predicting X from Y .

20. Let $\mathbf{X} \sim \text{Mult}_5(n, \mathbf{p})$.

(a) Find $E(X_1 | X_2)$ and $\text{Var}(X_1 | X_2)$.

(b) Find $E(X_1 | X_2 + X_3)$.

21. Let Y be a discrete r.v., A be an event with $0 < P(A) < 1$, and I_A be the indicator r.v. for A .

(a) Explain precisely how the r.v. $E(Y | I_A)$ relates to the numbers $E(Y | A)$ and $E(Y | A^c)$.

(b) Show that $E(Y | A) = E(YI_A) / P(A)$, directly from the definitions of expectation and conditional expectation.

Hint: Let $X = YI_A$, and then find an expression for the PMF of X .

(c) Use (b) to give a short proof of the fact that $E(Y) = E(Y | A)P(A) + E(Y | A^c)P(A^c)$.

22. Show that the following version of LOTP, which we encountered in Section 7.1, is also a consequence of Adam's law: for any event A and continuous r.v. X with PDF f_X ,

$$P(A) = \int_{-\infty}^{\infty} P(A | X = x) f_X(x) dx$$

Hint: Consider $E(I(A) | X = x)$.

23. (S) Let X and Y be random variables with finite variances, and let $W = Y - E(Y | X)$. This is a residual: the difference between the true value of Y and the predicted value of Y based on X .

(a) Compute $E(W)$ and $E(W | X)$.

(b) Compute $\text{Var}(W)$, for the case that $W | X \sim \mathcal{N}(0, X^2)$ with $X \sim \mathcal{N}(0, 1)$.

24. (S) One of two identical-looking coins is picked from a hat randomly, where one coin has probability p_1 of Heads and the other has probability p_2 of Heads. Let X be the number of Heads after flipping the chosen coin n times. Find the mean and variance of X .

25. Kelly makes a series of n bets, each of which she has probability p of winning, independently. Initially, she has x_0 dollars. Let X_j be the amount she has immediately after her j th bet is settled. Let f be a constant in $(0, 1)$, called the betting fraction. On each bet, Kelly wagers a fraction f of her wealth, and then she either wins or loses that amount. For example, if her current wealth is \$100 and $f = 0.25$, then she bets \$25 and either gains or loses that amount. (A famous choice when $p > 1/2$ is $f = 2p - 1$, which is known as the Kelly criterion.) Find $E(X_n)$ (in terms of n, p, f, x_0).

Hint: First find $E(X_{j+1} | X_j)$.

26. Let $N \sim \text{Pois}(\lambda_1)$ be the number of movies that will be released next year. Suppose that for each movie the number of tickets sold is $\text{Pois}(\lambda_2)$, independent of other movies and of N . Find the mean and variance of the number of movie tickets that will be sold next year.

27. A party is being held from 8:00 pm to midnight on a certain night, and $N \sim \text{Pois}(\lambda)$ people are going to show up. They will all arrive at uniformly random times while the party is going on, independently of each other and of N .

(a) Find the expected time at which the first person arrives, given that at least one person shows up. Give both an exact answer in terms of λ , measured in minutes after 8:00 pm, and an answer rounded to the nearest minute for $\lambda = 20$, expressed in time notation (e.g., 8:20 pm).

(b) Find the expected time at which the last person arrives, given that at least one person shows up. As in (a), give both an exact answer and an answer rounded to the nearest minute for $\lambda = 20$.

28. (S) We wish to estimate an unknown parameter θ , based on an r.v. X we will get to observe. As in the Bayesian perspective, assume that X and θ have a joint distribution. Let $\hat{\theta}$ be the estimator (which is a function of X). Then $\hat{\theta}$ is said to be unbiased if $E(\hat{\theta} | \theta) = \theta$, and $\hat{\theta}$ is said to be the Bayes procedure if $E(\theta | X) = \hat{\theta}$.

(a) Let $\hat{\theta}$ be unbiased. Find $E(\hat{\theta} - \theta)^2$ (the average squared difference between the estimator and the true value of θ), in terms of marginal moments of $\hat{\theta}$ and θ .

Hint: Condition on θ .

(b) Repeat (a), except in this part suppose that $\hat{\theta}$ is the Bayes procedure rather than assuming that it is unbiased.

Hint: Condition on X .

(c) Show that it is impossible for $\hat{\theta}$ to be both the Bayes procedure and unbiased, except in silly problems where we get to know θ perfectly by observing X .

Hint: If Y is a nonnegative r.v. with mean 0, then $P(Y = 0) = 1$.

29. Show that if $E(Y | X) = c$ is a constant, then X and Y are uncorrelated.

Hint: Use Adam's law to find $E(Y)$ and $E(XY)$.

30. Show by example that it is possible to have uncorrelated X and Y such that $E(Y | X)$ is not a constant.

Hint: Consider a standard Normal and its square.

31. (S) Emails arrive one at a time in an inbox. Let T_n be the time at which the n th email arrives (measured on a continuous scale from some starting point in time). Suppose that the waiting times between emails are i.i.d. $\text{Expo}(\lambda)$, i.e., $T_1, T_2 - T_1, T_3 - T_2, \dots$ are i.i.d. $\text{Expo}(\lambda)$.

Each email is non-spam with probability p , and spam with probability $q = 1 - p$ (independently of the other emails and of the waiting times). Let X be the time at which the first non-spam email arrives (so X is a continuous r.v., with $X = T_1$ if the 1st email is non-spam, $X = T_2$ if the 1st email is spam but the 2nd one isn't, etc.).

(a) Find the mean and variance of X .

(b) Find the MGF of X . What famous distribution does this imply that X has (be sure to state its parameter values)?

Hint for both parts: Let N be the number of emails until the first non-spam (including that one), and write X as a sum of N terms; then condition on N .

32. Customers arrive at a store according to a Poisson process of rate λ customers per hour. Each makes a purchase with probability p , independently. Given that a customer makes a purchase, the amount spent has mean μ (in dollars) and variance σ^2 .

(a) Find the mean and variance of how much a random customer spends (note that the customer may spend nothing).

(b) Find the mean and variance of the revenue the store obtains in an 8-hour time interval, using (a) and results from this chapter.

(c) Find the mean and variance of the revenue the store obtains in an 8-hour time interval, using the chicken-egg story and results from this chapter.

33. Fred's beloved computer will last an $\text{Expo}(\lambda)$ amount of time until it has a malfunction. When that happens, Fred will try to get it fixed. With probability p , he will be able to get it fixed. If he is able to get it fixed, the computer is good as new again and will last an additional, independent $\text{Expo}(\lambda)$ amount of time until the next malfunction (when again he is able to get it fixed with probability p , and so on). If after any malfunction Fred is unable to get it fixed, he will buy a new computer. Find the expected amount of time until Fred buys a new computer. (Assume that the time spent on computer diagnosis, repair, and shopping is negligible.)

34. A green die is rolled until it lands 1 for the first time. An orange die is rolled until it lands 6 for the first time. The dice are fair, six-sided dice. Let T_1 be the sum of the values of the rolls of the green die (including the 1 at the end) and T_6 be the sum of the values of the rolls of the orange die (including the 6 at the end). Two students are debating whether $E(T_1) = E(T_6)$ or $E(T_1) < E(T_6)$. They kindly gave permission to quote their arguments here.

Student A: We have $E(T_1) = E(T_6)$. By Adam's law, the expected sum of the rolls of a die is the expected number of rolls times the expected value of one roll, and each of these factors is the same for the two dice. In more detail, let N_1 be the number of rolls of the green die and N_6 be the number of rolls of the orange die. By Adam's law and linearity,

$$E(T_1) = E(E(T_1 | N_1)) = E(3.5N_1) = 3.5E(N_1)$$

and the same method applied to the orange die gives $3.5E(N_6)$, which equals $3.5E(N_1)$.

Student B: Actually, $E(T_1) < E(T_6)$. I agree that the expected number of rolls is the same for the two dice, but the key difference is that we know the last roll is a 1 for the green die and a 6 for the orange die. The expected totals are the same for the two dice excluding the last roll of each, and then including the last roll makes $E(T_1) < E(T_6)$.

(a) Discuss in words the extent to which Student A's argument is convincing and correct.

(b) Discuss in words the extent to which Student B's argument is convincing and correct.

(c) Give careful derivations of $E(T_1)$ and $E(T_6)$.

35. (S) Judit plays in a total of $N \sim \text{Geom}(s)$ chess tournaments in her career. Suppose that in each tournament she has probability p of winning the tournament, independently. Let T be the number of tournaments she wins in her career.

(a) Find the mean and variance of T .

(b) Find the MGF of T . What is the name of this distribution (with its parameters)?

36. In Story 8.4.5, we showed (among other things) that if $\lambda \sim \text{Gamma}(r_0, b_0)$ and $Y | \lambda \sim \text{Pois}(\lambda)$, then the marginal distribution of Y is $\text{NBin}(r_0, b_0/(b_0 + 1))$. Derive this result using Adam's law and MGFs.

Hint: Consider the conditional MGF of $Y | \lambda$.

37. Let $X_1, \dots, X_n$ be i.i.d. r.v.s with mean μ and variance σ^2 , and $n \geq 2$. A bootstrap sample of $X_1, \dots, X_n$ is a sample of n r.v.s $X_1^*, \dots, X_n^*$ formed from the X_j by sampling with replacement with equal probabilities. Let $\bar{X}^*$ denote the sample mean of the bootstrap sample:

$$\bar{X}^* = \frac{1}{n} (X_1^* + \dots + X_n^*)$$

(a) Calculate $E(X_j^*)$ and $\text{Var}(X_j^*)$ for each j .

(b) Calculate $E(\bar{X}^* | X_1, \dots, X_n)$ and $\text{Var}(\bar{X}^* | X_1, \dots, X_n)$.

Hint: Conditional on $X_1, \dots, X_n$, the X_j^* are independent, with a PMF that puts probability $1/n$ at each of the points $X_1, \dots, X_n$. As a check, your answers should be random variables that are functions of $X_1, \dots, X_n$.

(c) Calculate $E(\bar{X}^*)$ and $\text{Var}(\bar{X}^*)$.

(d) Explain intuitively why $\text{Var}(\bar{X}) < \text{Var}(\bar{X}^*)$.

38. An insurance company covers disasters in two neighboring regions, R_1 and R_2 . Let I_1 and I_2 be the indicator r.v.s for whether R_1 and R_2 are hit by the insured disaster, respectively. The indicators I_1 and I_2 may be dependent. Let $p_j = E(I_j)$ for $j = 1, 2$, and $p_{12} = E(I_1 I_2)$. The company reimburses a total cost of

$$C = I_1 \cdot T_1 + I_2 \cdot T_2$$

to these regions, where T_j has mean μ_j and variance σ_j^2 . Assume that T_1 and T_2 are independent of each other and that (T_1, T_2) is independent of (I_1, I_2) .

(a) Find $E(C)$.

(b) Find $\text{Var}(C)$.

39. (S) A certain stock has low volatility on some days and high volatility on other days. Suppose that the probability of a low volatility day is p and of a high volatility day is $q = 1 - p$, and that on low volatility days the percent change in the stock price is $\mathcal{N}(0, \sigma_1^2)$, while on high volatility days the percent change is $\mathcal{N}(0, \sigma_2^2)$, with $\sigma_1 < \sigma_2$.

Let X be the percent change of the stock on a certain day. The distribution is said to be a mixture of two Normal distributions, and a convenient way to represent X is as $X = I_1 X_1 + I_2 X_2$ where I_1 is the indicator r.v. of having a low volatility day, $I_2 = 1 - I_1$, $X_j \sim \mathcal{N}(0, \sigma_j^2)$, and I_1, X_1, X_2 are independent.

(a) Find $\text{Var}(X)$ in two ways: using Eve's law, and by using properties of covariance to calculate $\text{Cov}(I_1 X_1 + I_2 X_2, I_1 X_1 + I_2 X_2)$.

(b) Recall from Chapter 6 that the kurtosis of an r.v. Y with mean μ and standard deviation σ is defined by

$$\text{Kurt}(Y) = \frac{E(Y - \mu)^4}{\sigma^4} - 3$$

Find the kurtosis of X (in terms of $p, q, \sigma_1^2, \sigma_2^2$, fully simplified). The result will show that even though the kurtosis of any Normal distribution is 0, the kurtosis of X is positive and in fact can be very large depending on the parameter values.

40. Let X_1, X_2 , and Y be random variables, such that Y has finite variance. Let

$$A = E(Y | X_1) \text{ and } B = E(Y | X_1, X_2).$$

Show that

$$\text{Var}(A) \leq \text{Var}(B)$$

Also, check that this make sense in the extreme cases where Y is independent of X_1 and where $Y = h(X_2)$ for some function h .

Hint: Use Eve's law on B .

41. Show that for any r.v.s X and Y ,

$$E(Y | E(Y | X)) = E(Y | X)$$

This has a nice intuitive interpretation if we think of $E(Y | X)$ as the prediction we would make for Y based on X : given the prediction we would use for predicting Y from X , we no longer need to know X to predict Y -we can just use the prediction we have! For example, letting $E(Y | X) = g(X)$, if we observe $g(X) = 7$, then we may or may not know what X is (since g may not be one-to-one). But even without knowing X , we know that the prediction for Y based on X is 7.

Hint: Use Adam's law with extra conditioning.

42. A researcher wishes to know whether a new treatment for the disease conditionitis is more effective than the standard treatment. It is unfortunately not feasible to do a randomized

experiment, but the researcher does have the medical records of patients who received the new treatment and those who received the standard treatment. She is worried, though, that doctors tend to give the new treatment to younger, healthier patients. If this is the case, then naively comparing the outcomes of patients in the two groups would be like comparing apples and oranges.

Suppose each patient has background variables $\mathbf{X}$, which might be age, height and weight, and measurements relating to previous health status. Let Z be the indicator of receiving the new treatment. The researcher fears that Z is dependent on $\mathbf{X}$, i.e., that the distribution of $\mathbf{X}$ given $Z = 1$ is different from the distribution of $\mathbf{X}$ given $Z = 0$.

In order to compare apples to apples, the researcher wants to match every patient who received the new treatment to a patient with similar background variables who received the standard treatment. But $\mathbf{X}$ could be a high-dimensional random vector, which often makes it very difficult to find a match with a similar value of $\mathbf{X}$.

The propensity score reduces the possibly high-dimensional vector of background variables down to a single number (then it is much easier to match someone to a person with a similar propensity score than to match someone to a person with a similar value of $\mathbf{X}$). The propensity score of a person with background characteristics $\mathbf{X}$ is defined as

$$S = E(Z | \mathbf{X})$$

By the fundamental bridge, a person's propensity score is their probability of receiving the treatment, given their background characteristics. Show that conditional on S , the treatment indicator Z is independent of the background variables $\mathbf{X}$.

Hint: This problem relates to the previous one. Show that $P(Z = 1 | S, \mathbf{X}) = P(Z = 1 | S)$, which is equivalent to showing $E(Z | S, \mathbf{X}) = E(Z | S)$.

43. This exercise develops a useful identity for covariance, similar in spirit to Adam's law for expectation and Eve's law for variance. First define conditional covariance in a manner analogous to how we defined conditional variance:

$$\text{Cov}(X, Y | Z) = E((X - E(X | Z))(Y - E(Y | Z)) | Z)$$

(a) Show that

$$\text{Cov}(X, Y | Z) = E(XY | Z) - E(X | Z)E(Y | Z)$$

This should be true since it is the conditional version of the fact that

$$\text{Cov}(X, Y) = E(XY) - E(X)E(Y)$$

and conditional probabilities are probabilities, but for this problem you should prove it directly using properties of expectation and conditional expectation.

(b) ECCE, or the law of total covariance, says that

$$\text{Cov}(X, Y) = E(\text{Cov}(X, Y | Z)) + \text{Cov}(E(X | Z), E(Y | Z))$$

That is, the covariance of X and Y is the expected value of their conditional covariance plus the covariance of their conditional expectations, where all these conditional quantities are conditional on Z . Prove this identity.

Hint: We can assume without loss of generality that $E(X) = E(Y) = 0$, since adding a constant to an r.v. has no effect on its covariance with any r.v. Then expand out the covariances on the right-hand side of the identity and apply Adam's law.

Mixed Practice

44. A group of n friends often go out for dinner together. At their dinners, they play "credit card roulette" to decide who pays the bill. This means that at each dinner, one person is chosen uniformly at random to pay the entire bill (independently of what happens at the other dinners).
- Find the probability that in k dinners, no one will have to pay the bill more than once (do not simplify for the case $k \leq n$, but do simplify fully for the case $k > n$).
 - Find the expected number of dinners it takes in order for everyone to have paid at least once (you can leave your answer as a finite sum of simple-looking terms).
 - Alice and Bob are two of the friends. Find the covariance between how many times Alice pays and how many times Bob pays in k dinners.
45. As in the previous problem, a group of n friends play "credit card roulette" at their dinners. In this problem, let the number of dinners be a $\text{Pois}(\lambda)$ r.v.
- Alice is one of the friends. Find the correlation between how many dinners Alice pays for and how many free dinners Alice gets.
 - The costs of the dinners are i.i.d. $\text{Gamma}(a, b)$ r.v.s, independent of the number of dinners. Find the mean and variance of the total cost.
46. Joe will read $N \sim \text{Pois}(\lambda)$ books next year. Each book has a $\text{Pois}(\mu)$ number of pages, with book lengths independent of each other and independent of N .
- Find the expected number of book pages that Joe will read next year.
 - Find the variance of the number of book pages Joe will read next year.
 - For each of the N books, Joe likes it with probability p and dislikes it with probability $1 - p$, independently. Find the conditional distribution of how many of the N books Joe likes, given that he dislikes exactly d of the books.
47. Buses arrive at a certain bus stop according to a Poisson process of rate λ . Each bus has n seats and, at the instant when it arrives at the stop, has a $\text{Bin}(n, p)$ number of passengers. Assume that the numbers of passengers on different buses are independent of each other, and independent of the arrival times of the buses.
- Let N_t be the number of buses that arrive in the time interval $[0, t]$, and X_t be the total number of passengers on the buses that arrive in the time interval $[0, t]$.
- Find the mean and variance of N_t .
 - Find the mean and variance of X_t .
 - A bus is full if it has exactly n passengers when it arrives at the stop. Find the probability that exactly $a + b$ buses arrive in $[0, t]$, of which a are full and b are not full.
48. Paul and n other runners compete in a marathon. Their times are independent continuous r.v.s with CDF F .
- For $j = 1, 2, \dots, n$, let A_j be the event that anonymous runner j completes the race faster than Paul. Explain whether the events A_j are independent, and whether they are conditionally independent given Paul's time to finish the race.
 - For the rest of this problem, let N be the number of runners who finish faster than Paul. Find $E(N)$. (Your answer should depend only on n , since Paul's time is an r.v.)
 - Find the conditional distribution of N , given that Paul's time to finish the marathon is t .
 - Find $\text{Var}(N)$. (Your answer should depend only on n , since Paul's time is an r.v.)
- Hint: Let T be Paul's time, and use Eve's law to condition on T . Alternatively, use indicator r.v.s.
49. Emails arrive in an inbox according to a Poisson process of rate λ emails per hour.
- Find the name and parameters of the conditional distribution of the number of emails that

arrive in the first 2 hours of an 8 -hour time period, given that exactly n emails arrive in that time period.

(b) Each email is legitimate with probability p and spam with probability $q = 1 - p$, independently. Find the name and parameters of the conditional distribution of the number of legitimate emails that arrive in an 8 -hour time period, given that exactly s spams arrived in that time period.

(c) Reading an email takes a random amount of time, with mean μ hours and standard deviation σ hours. These reading times are i.i.d. and independent of the email arrival process. Find the (unconditional) mean and variance of the total time it takes to read all the emails that arrive in an 8 -hour time period.

50. An actuary wishes to estimate various quantities related to the number of insurance claims and the dollar amounts of those claims for someone named Fred. Suppose that Fred will make N claims next year, where $N \mid \lambda \sim \text{Pois}(\lambda)$. But λ is unknown, so the actuary, taking a Bayesian approach, gives λ a prior distribution based on past experience. Specifically, the prior is $\lambda \sim \text{Expo}(1)$. The dollar amount of a claim is Log-Normal with parameters μ and σ^2 (here μ and σ^2 are the mean and variance of the underlying Normal), with μ and σ^2 known. The dollar amounts of the claims are i.i.d. and independent of N .

(a) Find $E(N)$ and $\text{Var}(N)$ using properties of conditional expectation (your answers should not depend on λ , since λ is unknown and being treated as an r.v.!).

(b) Find the mean and variance of the total dollar amount of all the claims.

(c) Find the distribution of N . If it is a named distribution we have studied, give its name and parameters.

(d) Find the posterior distribution of λ , given that it is observed that Fred makes $N = n$ claims next year. If it is a named distribution we have studied, give its name and parameters.

51. (S) Let X_1, X_2, X_3 be independent with $X_i \sim \text{Expo}(\lambda_i)$ (so with possibly different rates). Recall from Chapter 7 that

$$P(X_1 < X_2) = \frac{\lambda_1}{\lambda_1 + \lambda_2}$$

(a) Find $E(X_1 + X_2 + X_3 \mid X_1 > 1, X_2 > 2, X_3 > 3)$ in terms of $\lambda_1, \lambda_2, \lambda_3$.

(b) Find $P(X_1 = \min(X_1, X_2, X_3))$, the probability that the first of the three Exponentials is the smallest.

Hint: Restate this in terms of X_1 and $\min(X_2, X_3)$.

(c) For the case $\lambda_1 = \lambda_2 = \lambda_3 = 1$, find the PDF of $\max(X_1, X_2, X_3)$. Is this one of the important distributions we have studied?

52. (S) A task is randomly assigned to one of two people (with probability $1/2$ for each person). If assigned to the first person, the task takes an $\text{Expo}(\lambda_1)$ length of time to complete (measured in hours), while if assigned to the second person it takes an $\text{Expo}(\lambda_2)$ length of time to complete (independent of how long the first person would have taken). Let T be the time taken to complete the task.

(a) Find the mean and variance of T .

(b) Suppose instead that the task is assigned to both people, and let X be the time taken to complete it (by whoever completes it first, with the two people working independently). It is observed that after 24 hours, the task has not yet been completed. Conditional on this information, what is the expected value of X ?

53. Suppose for this problem that "true IQ" is a meaningful concept rather than a reified social construct. Suppose that in the U.S. population, the distribution of true IQs is Normal with mean 100 and SD 15. A person is chosen at random from this population to take an IQ test. The test is a noisy measure of true ability: it's correct on average but has a Normal measurement error with SD 5.

Let μ be the person's true IQ, viewed as a random variable, and let Y be her score on the IQ test. Then we have

$$\begin{aligned} Y | \mu &\sim \mathcal{N}(\mu, 5^2) \\ \mu &\sim \mathcal{N}(100, 15^2) \end{aligned}$$

- (a) Find the unconditional mean and variance of Y .
 - (b) Find the marginal distribution of Y . One way is via the MGF.
 - (c) Find $\text{Cov}(\mu, Y)$.
54. (S) A certain genetic characteristic is of interest. It can be measured numerically. Let X_1 and X_2 be the values of the genetic characteristic for two twin boys. Given that they are identical twins, $X_1 = X_2$ and X_1 has mean 0 and variance σ^2 ; given that they are fraternal twins, X_1 and X_2 have mean 0, variance σ^2 , and correlation ρ . The probability that the twins are identical is $1/2$. Find $\text{Cov}(X_1, X_2)$ in terms of ρ, σ^2 .
55. (S) The Mass Cash lottery randomly chooses 5 of the numbers from $1, 2, \dots, 35$ each day (without repetitions within the choice of 5 numbers). Suppose that we want to know how long it will take until all numbers have been chosen. Let a_j be the average number of additional days needed if we are missing j numbers (so $a_0 = 0$ and a_{35} is the average number of days needed to collect all 35 numbers). Find a recursive formula for the a_j .
56. Two chess players, Vishy and Magnus, play a series of games. Given p , the game results are i.i.d. with probability p of Vishy winning, and probability $q = 1 - p$ of Magnus winning (assume that each game ends in a win for one of the two players). But p is unknown, so we will treat it as an r.v. To reflect our uncertainty about p , we use the prior $p \sim \text{Beta}(a, b)$, where a and b are known positive integers and $a \geq 2$.
- (a) Find the expected number of games needed in order for Vishy to win a game (including the win). Simplify fully; your final answer should not use factorials or Γ .
 - (b) Explain in terms of independence vs. conditional independence the direction of the inequality between the answer to (a) and $1 + E(G)$ for $G \sim \text{Geom}(\frac{a}{a+b})$.
 - (c) Find the conditional distribution of p given that Vishy wins exactly 7 out of the first 10 games.
57. Laplace's law of succession says that if $X_1, X_2, \dots, X_{n+1}$ are conditionally independent $\text{Bern}(p)$ r.v.s given p , but p is given a $\text{Unif}(0, 1)$ prior to reflect ignorance about its value, then

$$P(X_{n+1} = 1 | X_1 + \dots + X_n = k) = \frac{k+1}{n+2}$$

As an example, Laplace discussed the problem of predicting whether the sun will rise tomorrow, given that the sun did rise every time for all n days of recorded history; the above formula then gives $(n+1)/(n+2)$ as the probability of the sun rising tomorrow (of course, assuming independent trials with p unchanging over time may be a very unreasonable model for the sunrise problem).

- (a) Find the posterior distribution of p given $X_1 = x_1, X_2 = x_2, \dots, X_n = x_n$, and show that it only depends on the sum of the x_j (so we only need the one-dimensional quantity $x_1 + x_2 + \dots + x_n$ to obtain the posterior distribution, rather than needing all n data points).
 - (b) Prove Laplace's law of succession, using a form of the law of total probability to find $P(X_{n+1} = 1 | X_1 + \dots + X_n = k)$ by conditioning on p . (The next exercise, which is closely related, involves an equivalent Adam's law proof.)
58. Two basketball teams, A and B , play an n game match. Let X_j be the indicator of team A winning the j th game. Given p , the r.v.s $X_1, \dots, X_n$ are i.i.d. with $X_j | p \sim \text{Bern}(p)$. But p is unknown, so we will treat it as an r.v. Let the prior distribution be $p \sim \text{Unif}(0, 1)$, and let X be the number of wins for team A .
- (a) Find $E(X)$ and $\text{Var}(X)$.

(b) Use Adam's law to find the probability that team A will win game $j + 1$, given that they win exactly a of the first j games. (The previous exercise, which is closely related, involves an equivalent LOTP proof.)

Hint: Letting C be the event that team A wins exactly a of the first j games,

$$P(X_{j+1} = 1 \mid C) = E(X_{j+1} \mid C) = E(E(X_{j+1} \mid C, p) \mid C) = E(p \mid C)$$

(c) Find the PMF of X . (There are various ways to do this, including a very fast way to see it based on results from earlier chapters.)

(d) The Putnam exam from 2002 posed the following problem:

Shanille O'Keal shoots free throws on a basketball court. She hits the first and misses the second, and thereafter the [conditional] probability that she hits the next shot is equal to the proportion of shots she has hit so far. What is the probability she hits exactly 50 of her first 100 shots?

Solve this Putnam problem by applying the result of Part (c). Be sure to explain why it is valid to apply that result, despite the fact that the Putnam problem does not seem to be using the same model, e.g., it does not mention a prior distribution, let alone mention a $\text{Unif}(0, 1)$ prior.

59. Let $X \mid p \sim \text{Bin}(n, p)$, with $p \sim \text{Beta}(a, b)$. So X has a Beta-Binomial distribution, as mentioned in Story 8.3.3 and Example 8.5.3. Find $E(X)$ and $\text{Var}(X)$.

60. An election is being held. There are two candidates, A and B, and there are n voters. The probability of voting for Candidate A varies by city. There are m cities, labeled $1, 2, \dots, m$. The j th city has n_j voters, so $n_1 + n_2 + \dots + n_m = n$. Let X_j be the number of people in the j th city who vote for Candidate A, with $X_j \mid p_j \sim \text{Bin}(n_j, p_j)$. To reflect our uncertainty about the probability of voting in each city, we treat $p_1, \dots, p_m$ as r.v.s, with prior distribution asserting that they are i.i.d. $\text{Unif}(0, 1)$. Assume that $X_1, \dots, X_m$ are independent, both unconditionally and conditional on $p_1, \dots, p_m$. Let X be the total number of votes for Candidate A.

(a) Find the marginal distribution of X_1 and the posterior distribution of $p_1 \mid (X_1 = k_1)$.

(b) Find $E(X)$ and $\text{Var}(X)$ in terms of n and s , where $s = n_1^2 + n_2^2 + \dots + n_m^2$.

4.10 10 Inequalities and Limit Theorems

"What should I do if I can't calculate a probability or expectation exactly?" Almost everyone who uses probability has to deal with this sometimes. Don't panic. There are powerful strategies available: simulate it, bound it, or approximate it.

- Simulations using Monte Carlo: We have already seen many examples of simulations in this book; the R sections give numerous examples where a few lines of code and a few seconds on a computer suffice to get good approximate answers. "Monte Carlo" just means that the simulations use random numbers (the term originated from the Monte Carlo Casino in Monaco).

Monte Carlo simulation is an extremely powerful technique, and there are many problems where it is the only reasonable approach currently available. So why not always just do a simulation? Here are a few reasons:

1. The simulation may need to run for an extremely long time, even on a fast computer. A major extension known as Markov chain Monte Carlo, introduced in Chapter 12, greatly increases the range of problems for which Monte Carlo simulation is feasible. But even then, the simulation may need to run for a vast, unknown amount of time to get decent answers to the problem.

2. We may hope to get a good answer for all values of the parameters of the problem. For example, in the coupon collector problem (Example 4.3.12) we saw that with n toy types, it takes about $n \log n$ toys on average to get a complete set. This is a simple, memorable answer. It would be easy to simulate this for any specific n , e.g., we could run a coupon collector process for $n = 20$ and get an answer of around 60 . But this would not make it easy to see the general $n \log n$ result.
 3. A simulation result is easy to criticize: how do you know you ran it long enough? How do you know your result is close to the truth? How close is "close"? We may want provable guarantees instead.
- Bounds using inequalities: A bound on a probability gives a provable guarantee that the probability is in a certain range. In the first part of this chapter, we'll introduce several important inequalities in probability. These inequalities will often allow us to narrow down the range of possible values for the exact answer, that is, to determine an upper bound and/or lower bound. A bound may not provide a good approximation - if we have bounds of $[0.2, 0.6]$ for a probability we're trying to find, the exact answer could be anywhere within that range - but at least we know the exact answer is guaranteed to be inside the bounds.
 - Approximations using limit theorems: Later in the chapter, we will discuss the two most famous theorems in probability: the law of large numbers and the central limit theorem. Both tell us what happens to the sample mean as we obtain more and more data. Limit theorems let us make approximations which are likely to work well when we have a large number of data points. We conclude the chapter by using limit theorems to study a couple of important named distributions in statistics.

4.10.1 10.1 Inequalities

10.1.1 Cauchy-schwarz: a Marginal Bound on a Joint Expectation

The Cauchy-Schwarz inequality is one of the most famous inequalities in all of mathematics. In probability, it takes the following form.

Theorem 10.1.1 (Cauchy-Schwarz). For any r.v.s X and Y with finite variances,

$$|E(XY)| \leq \sqrt{E(X^2) E(Y^2)}$$

Proof. For any t ,

$$0 \leq E(Y - tX)^2 = E(Y^2) - 2tE(XY) + t^2E(X^2).$$

Where did t come from? The idea is to introduce t so that we have infinitely many inequalities, one for each value of t , and then we can use calculus to find the value of t that gives us the best inequality. Differentiating the right-hand side with respect to t and setting it equal to 0 , we get that $t = E(XY)/E(X^2)$ minimizes the righthand side, resulting in the tightest bound. Plugging in this value of t , we have the Cauchy-Schwarz inequality.

If X and Y are uncorrelated, then $E(XY) = E(X)E(Y)$, which depends only on the marginal expectations $E(X)$ and $E(Y)$. But in general, calculating $E(XY)$ exactly requires knowing the joint distribution of X and Y (and being able to work with it). The Cauchy-Schwarz inequality lets us bound $E(XY)$ in terms of the marginal second moments $E(X^2)$ and $E(Y^2)$.

If X and Y have mean 0 , then Cauchy-Schwarz already has a very familiar statistical interpretation: it says that their correlation is between -1 and 1 .

Example 10.1.2. Let $E(X) = E(Y) = 0$. Then

$$E(XY) = \text{Cov}(X, Y), E(X^2) = \text{Var}(X), E(Y^2) = \text{Var}(Y)$$

so Cauchy-Schwarz reduces to the statement $|\text{Corr}(X, Y)| \leq 1$. Of course, we already knew this from Theorem 7.3.5. Now let's see what happens if we drop the assumption that the means are 0. Applying Cauchy-Schwarz to the centered r.v.s $X - E(X)$ and $Y - E(Y)$, we again have that $|\text{Corr}(X, Y)| \leq 1$.

Cauchy-Schwarz can often be applied in creative ways. For example, if we write $X = X \cdot 1$, then Cauchy-Schwarz gives $|E(X \cdot 1)| \leq \sqrt{E(X^2)E(1^2)}$, which reduces to $E(X^2) \geq (EX)^2$. This gives a quick new proof that variances are nonnegative. As another example, we will obtain an upper bound on the probability of a nonnegative r.v. equaling 0.

Example 10.1.3 (Second moment method). Let X be a nonnegative r.v., and suppose that we want an upper bound on $P(X = 0)$. For example, X could be the number of questions that Fred gets wrong on an exam (then $P(X = 0)$ is the probability of Fred getting a perfect score), or X could be the number of pairs of people at a party with the same birthday (then $P(X = 0)$ is the probability of no birthday matches). Note that

$$X = XI(X > 0),$$

where $I(X > 0)$ is the indicator of $X > 0$. This is true since if $X = 0$ then both sides are 0, while if $X > 0$ then both sides are X . By Cauchy-Schwarz,

$$E(X) = E(XI(X > 0)) \leq \sqrt{E(X^2)E(I(X > 0))}.$$

Rearranging this and using the fundamental bridge, we have

$$P(X > 0) \geq \frac{(EX)^2}{E(X^2)}$$

or equivalently,

$$P(X = 0) \leq \frac{\text{Var}(X)}{E(X^2)}$$

Applying this bound is sometimes called the second moment method. For example, let's apply the bound in the case that

$$X = I_1 + \cdots + I_n,$$

where the I_j are uncorrelated indicator r.v.s. Let $p_j = E(I_j)$. Then

$$\text{Var}(X) = \sum_{j=1}^n \text{Var}(I_j) = \sum_{j=1}^n (p_j - p_j^2) = \sum_{j=1}^n p_j - \sum_{j=1}^n p_j^2 = \mu - c,$$

where $\mu = E(X)$, $c = \sum_{j=1}^n p_j^2$. Also, $E(X^2) = \text{Var}(X) + (EX)^2 = \mu^2 + \mu - c$. So

$$P(X = 0) \leq \frac{\text{Var}(X)}{E(X^2)} = \frac{\mu - c}{\mu^2 + \mu - c} \leq \frac{1}{\mu + 1},$$

where the last inequality is easily checked by cross-multiplying. In general, it is wrong to say "if X has a high mean, then it has a small chance of being 0", since it could be that X is usually 0 but has a small chance of being extremely large. But in our current setting, we have a simple, quantitative way to say that X having a high mean does mean that X is unlikely to be 0.

For example, suppose there are 14 people in a room. How likely is it that there are two people with the same birthday or birthdays one day apart? This is much harder to solve exactly than the birthday problem, so in Example 4.7.6 we used a Poisson approximation. But we may want a guarantee from a bound rather than worrying about whether the Poisson approximation is good enough. Let X be the number of "near birthday" pairs. Using indicator r.v.s, we have $E(X) = \binom{14}{2} \frac{3}{365}$. So

$$P(X = 0) \leq \frac{1}{E(X) + 1} < 0.573$$

The true answer for $P(X = 0)$ turns out to be 0.46 (to two decimal places), which is consistent with the bound.

Cauchy-Schwarz also allows us to deduce the existence of a joint MGF from the existence of marginal MGFs; this is another example of the benefit of being able to bound a joint distribution quantity by marginal distribution quantities.

Example 10.1.4 (Existence of joint MGF). Let X_1 and X_2 be jointly distributed r.v.s, not necessarily independent or identically distributed. Show that if X_1 and X_2 both have MGFs marginally, then the random vector (X_1, X_2) has a joint MGF.

Solution:

Recall from Chapter 7 that the joint MGF is defined as $M(s, t) = E(e^{sX_1 + tX_2})$, and exists if the expectation is finite in a box around the origin. The marginal MGFs are $E(e^{sX_1})$ and $E(e^{tX_2})$; each is required to be finite in an interval around the origin.

Suppose the MGFs of X_1 and X_2 are finite on $(-a, a)$. Fix s and t in $(-a/2, a/2)$. By Cauchy-Schwarz,

$$E(e^{sX_1 + tX_2}) \leq \sqrt{E(e^{2sX_1}) E(e^{2tX_2})}$$

The right-hand side is finite by assumption, so $E(e^{sX_1 + tX_2})$ is finite in the box $\{(s, t) : s, t \in (-a/2, a/2)\}$. Hence the joint MGF of (X_1, X_2) exists.

10.1.2 Jensen: an Inequality for Convexity

In 4.3.13, we discussed that for nonlinear functions g , $E(g(X))$ may be very different from $g(E(X))$. If g is either a convex function or a concave function, Jensen's inequality tells us exactly which of $E(g(X))$ and $g(E(X))$ is greater. See the math appendix for information about convex and concave functions. Often we can take the second derivative to test for convexity or concavity: assuming that g'' exists, g being convex is equivalent to $g''(x) \geq 0$ everywhere in the domain, and g being concave is equivalent to $g''(x) \leq 0$ everywhere in the domain.

Theorem 10.1.5 (Jensen). Let X be a random variable. If g is a convex function, then $E(g(X)) \geq g(E(X))$. If g is a concave function, then $E(g(X)) \leq g(E(X))$. In both cases, the only way that equality can hold is if there are constants a and b such that $g(X) = a + bX$ with probability 1.

Proof. If g is convex, then all lines that are tangent to g lie below g (see Figure 10.1). In particular, let $\mu = E(X)$, and consider the tangent line at the point $(\mu, g(\mu))$. (If g is differentiable at μ then the tangent line is unique; otherwise, choose any tangent line at μ .) Denoting this tangent line by $a + bx$, we have $g(x) \geq a + bx$ for all x by convexity, so $g(X) \geq a + bX$. Taking the expectation of both sides,

$$E(g(X)) \geq E(a + bX) = a + bE(X) = a + b\mu = g(\mu) = g(E(X)),$$

as desired. If g is concave, then $h = -g$ is convex, so we can apply what we just proved to h to see that the inequality for g is reversed from the convex case.

Lastly, assume that equality holds in the convex case. Let $Y = g(X) - a - bX$. Then Y is a nonnegative r.v. with $E(Y) = 0$, so $P(Y = 0) = 1$ (even a tiny nonzero chance of $Y > 0$ occurring would make $E(Y) > 0$). So equality holds if and only if $P(g(X) = a + bX) = 1$. For the concave case, we can use the same argument with $Y = a + bX - g(X)$.




FIGURE 10.1 Since g is convex, the tangent lines lie below the curve. In particular, the tangent line at the point $(\mu, g(\mu))$ lies below the curve.

Let's check Jensen's inequality in a couple of simple known cases.

- Since $g(x) = x^2$ is convex (its second derivative is 2), Jensen's inequality says $E(X^2) \geq (EX)^2$, which we already knew to be true since variances are nonnegative (or by Cauchy-Schwarz).
- In the St. Petersburg paradox from Chapter 4, we found that $E(2^N) > 2^{EN}$, where $N \sim \text{FS}(1/2)$. Jensen's inequality concurs since $g(x) = 2^x$ is convex (to find $g''(x)$, write $2^x = e^{x \log 2}$). Moreover, it tells us that the direction of the inequality would be the same no matter what distribution N has! The inequality will be strict unless N is constant (with probability 1).

If we ever forget the direction of Jensen's inequality, these simple cases make it easy to recover the correct direction. Here are a few more quick examples of Jensen's inequality:

- $E|X| \geq |EX|$,
- $E(1/X) \geq 1/(EX)$, for positive r.v.s X ,
- $E(\log(X)) \leq \log(EX)$, for positive r.v.s X .

As another example, we can use Jensen's inequality to see the direction of the bias if we estimate an unknown standard deviation using the sample standard deviation of a data set.

Example 10.1.6 (Bias of sample standard deviation). Let $X_1, \dots, X_n$ be i.i.d. random variables with variance σ^2 . Recall from Theorem 6.3.4 that the sample variance S_n^2 is unbiased for estimating σ^2 . That is, $E(S_n^2) = \sigma^2$. However, we are often more interested in estimating the standard deviation σ . A natural estimator for σ is the sample standard deviation, S_n .

Jensen's inequality shows us that S_n is biased for estimating σ . Moreover, it tells us which way the inequality goes:

$$E(S_n) = E(\sqrt{S_n^2}) \leq \sqrt{E(S_n^2)} = \sigma,$$

so the sample standard deviation tends to underestimate the true standard deviation. How biased it is depends on the distribution (so there is no universal way to fix the bias, in contrast to the fact that defining sample variance with $n - 1$ in the denominator makes it unbiased for all distributions). Fortunately, the bias is typically minor if the sample size is reasonably large.

One area in which Jensen's inequality is important is in information theory, the study of how to quantify information. The principles of information theory have become essential for communication and compression (e.g., for MP3s and cell phones). Here are a few examples of applications of Jensen's inequality.

Example 10.1.7 (Entropy). The surprise of learning that an event with probability p happened is defined as $\log_2(1/p)$, measured in a unit called bits. Low-probability events have high surprise, while an event with probability 1 has zero surprise. The log is there so that if we observe two independent events A and B , the total surprise is the same as the surprise from observing $A \cap B$. The log is base 2 so that if we learn that an event with probability $1/2$ happened, the surprise is 1, which corresponds to having received 1 bit of information.

Let X be a discrete r.v. whose distinct possible values are $a_1, a_2, \dots, a_n$, with probabilities $p_1, p_2, \dots, p_n$ respectively (so $p_1 + p_2 + \dots + p_n = 1$). The entropy of X is defined to be the average surprise of learning the value of X :

$$H(X) = \sum_{j=1}^n p_j \log_2(1/p_j)$$

Note that the entropy of X depends only on the probabilities p_j , not on the values a_j . So for example, $H(X^3) = H(X)$, since X^3 has distinct possible values $a_1^3, a_2^3, \dots, a_n^3$, with probabilities $p_1, p_2, \dots, p_n$ —the same list of p_j 's as for X .

Using Jensen's inequality, show that the maximum possible entropy for X is when its distribution is uniform over $a_1, a_2, \dots, a_n$, i.e., $p_j = 1/n$ for all j . This makes sense intuitively, since learning the value of X conveys the most information on average when X is equally likely to take any of its values, and the least possible information if X is a constant.

Solution:

Let $X \sim \text{DUnif}(a_1, \dots, a_n)$, so that

$$H(X) = \sum_{j=1}^n \frac{1}{n} \log_2(n) = \log_2(n)$$

Let Y be an r.v. that takes on values $1/p_1, \dots, 1/p_n$ with probabilities $p_1, \dots, p_n$, respectively (with the natural modification if the $1/p_j$ have some repeated values, e.g., if $1/p_1 = 1/p_2$ but none of the others are this value, then it gets $p_1 + p_2 = 2p_1$ as its probability). Then $H(Y) = E(\log_2(Y))$ by LOTUS, and $E(Y) = n$. So by Jensen's inequality,

$$H(Y) = E(\log_2(Y)) \leq \log_2(E(Y)) = \log_2(n) = H(X)$$

Since the entropy of an r.v. depends only on the probabilities p_j and not on the specific values that the r.v. takes on, the entropy of Y is unchanged if we alter the support from $1/p_1, \dots, 1/p_n$ to $a_1, \dots, a_n$. Therefore X , which is uniform on $a_1, \dots, a_n$, has entropy at least as large as that of any other r.v. with support $a_1, \dots, a_n$.

Example 10.1.8 (Kullback-Leibler divergence). Let $\mathbf{p} = (p_1, \dots, p_n)$ and $\mathbf{r} = (r_1, \dots, r_n)$ be probability vectors (so each is nonnegative and sums to 1). Think of each as a possible PMF for a random variable whose support consists of n distinct values. The Kullback-Leibler divergence between $\mathbf{p}$ and $\mathbf{r}$ is defined as

$$D(\mathbf{p}, \mathbf{r}) = \sum_{j=1}^n p_j \log_2 (1/r_j) - \sum_{j=1}^n p_j \log_2 (1/p_j).$$

This is the difference between the average surprise we will experience when the actual probabilities are $\mathbf{p}$ but we are instead working with $\mathbf{r}$ (for example, if $\mathbf{p}$ is unknown and $\mathbf{r}$ is our current guess for $\mathbf{p}$), and our average surprise when we work with $\mathbf{p}$. Show that the Kullback-Leibler divergence is nonnegative.

Solution:

Using properties of logs, we have

$$D(\mathbf{p}, \mathbf{r}) = - \sum_{j=1}^n p_j \log_2 \left(\frac{r_j}{p_j} \right).$$

Let Y be a random variable that takes on values r_j/p_j with probabilities p_j , so that $D(\mathbf{p}, \mathbf{r}) = -E(\log_2(Y))$ by LOTUS. Then by Jensen's inequality,

$$D(\mathbf{p}, \mathbf{r}) = -E(\log_2(Y)) \geq -\log_2(E(Y)) = -\log_2(1) = 0,$$

with equality if and only if $\mathbf{p} = \mathbf{r}$. This result tells us that we're more surprised on average when we work with the wrong probabilities than when we work with the correct probabilities.

Example 10.1.9 (Log probability scoring). Imagine that on a multiple-choice exam, instead of circling just one of the answer choices, you are asked to assign a probability of correctness to each choice. Your score on a particular question is the log of the probability that you assign to the correct answer. The maximum score for a particular question is 0, and the minimum score is $-\infty$, attained if you assign zero probability to the correct answer.¹

Suppose your personal probabilities of correctness for each of the n answer choices are $p_1, \dots, p_n$, where the p_j are positive and sum to 1. Show that your expected score on a question is maximized if you report your true probabilities p_j , not any other probabilities. In other words, under log probability scoring, you have no incentive to lie about your beliefs and pretend to be more or less confident than you really are (assuming that your goal is to maximize your expected score).

Solution:

This example is isomorphic to the previous one! Your expected score on a question is $\sum_{j=1}^n p_j \log p_j$ if you report your true probabilities $\mathbf{p}$, and $\sum_{j=1}^n p_j \log r_j$ if you report false probabilities $\mathbf{r}$. The difference between these two is precisely the Kullback-Leibler divergence between $\mathbf{p}$ and $\mathbf{r}$. This is always nonnegative, as we proved in the previous example. Therefore your expected score is maximized when you report your true probabilities.

10.1.3 Markov, Chebyshev, Chernoff: Bounds on Tail Probabilities

The inequalities in this section provide bounds on the probability of an r.v. taking on an "extreme" value in the right or left tail of a distribution.

⁰¹ Joe's philosophy professor as an undergraduate at Caltech, Alan Hájek, used precisely this system. He warned the class never to put a probability of zero, since a score of $-\infty$ would not only give a $-\infty$ on that exam, but also it would spill through and result in a $-\infty$ for the whole semester, since a weighted average with even a tiny positive weight on a $-\infty$ yields $-\infty$. Despite this warning, some students did put probability zero on the correct answers.

Theorem 10.1.10 (Markov). For any r.v. X and constant $a > 0$,

$$P(|X| \geq a) \leq \frac{E|X|}{a}$$

Proof. Let $Y = \frac{|X|}{a}$. We need to show that $P(Y \geq 1) \leq E(Y)$. Note that

$$I(Y \geq 1) \leq Y$$

since if $I(Y \geq 1) = 0$ then the inequality reduces to $Y \geq 0$, and if $I(Y \geq 1) = 1$ then $Y \geq 1$ (because the indicator says so). Taking the expectation of both sides, we have Markov's inequality.

For an intuitive interpretation, let X be the income of a randomly selected individual from a population. Taking $a = 2E(X)$, Markov's inequality says that $P(X \geq 2E(X)) \leq 1/2$, i.e., it is impossible for more than half the population to make at least twice the average income. This is clearly true, since if over half the population were earning at least twice the average income, the average income would be higher! Similarly, $P(X \geq 3E(X)) \leq 1/3$: you can't have more than $1/3$ of the population making at least three times the average income, since those people would already drive the average above what it is.

Markov's inequality is a very crude bound because it requires absolutely no assumptions about X . The right-hand side of the inequality could be greater than 1, or even infinite; this is not very helpful when trying to bound a number that we already know to be between 0 and 1. Surprisingly, the following two inequalities, which can be derived from Markov's inequality with almost no additional work, can often give us bounds that are much better than Markov's.

Theorem 10.1.11 (Chebyshev). Let X have mean μ and variance σ^2 . Then for any $a > 0$,

$$P(|X - \mu| \geq a) \leq \frac{\sigma^2}{a^2}$$

Proof. By Markov's inequality,

$$P(|X - \mu| \geq a) = P((X - \mu)^2 \geq a^2) \leq \frac{E(X - \mu)^2}{a^2} = \frac{\sigma^2}{a^2}.$$

Substituting $c\sigma$ for a , for $c > 0$, we have the following equivalent form of Chebyshev's inequality:

$$P(|X - \mu| \geq c\sigma) \leq \frac{1}{c^2}$$

This gives us an upper bound on the probability of an r.v. being more than c standard deviations away from its mean, e.g., there can't be more than a 25% chance of being 2 or more standard deviations from the mean.

The idea for proving Chebyshev from Markov was to square $|X - \mu|$ and then apply Markov. Similarly, it is often fruitful to perform other transformations before applying Markov. Chernoff's bound, which is widely used in engineering, uses this idea with an exponential function.

Theorem 10.1.12 (Chernoff). For any r.v. X and constants $a > 0$ and $t > 0$,

$$P(X \geq a) \leq \frac{E(e^{tX})}{e^{ta}}.$$

Proof. The transformation g with $g(x) = e^{tx}$ is invertible and strictly increasing. So by Markov's inequality, we have

$$P(X \geq a) = P(e^{tX} \geq e^{ta}) \leq \frac{E(e^{tX})}{e^{ta}}$$

At first it may not be clear what Chernoff's bound has to offer that Markov's inequality doesn't, but it has two very nice features:

1. The right-hand side can be optimized over t to give the tightest upper bound, as in the proof of Cauchy-Schwarz.
2. If the MGF of X exists, then the numerator in the bound is the MGF, and some of the useful properties of MGFs can come into play.

Let's now compare the three bounds just discussed by applying them to a simple example where the true probability is known.

Example 10.1.13 (Bounds on a Normal tail probability). Let $Z \sim \mathcal{N}(0, 1)$. By the 68-95-99.7% rule, we know that $P(|Z| > 3)$ is approximately 0.003 ; the exact value is $2 \cdot \Phi(-3)$. Let's see what upper bounds are obtained from Markov's, Chebyshev's, and Chernoff's inequalities.

- Markov: In Chapter 5, we found that $E|Z| = \sqrt{2/\pi}$. Then

$$P(|Z| > 3) \leq \frac{E|Z|}{3} = \frac{1}{3} \cdot \sqrt{\frac{2}{\pi}} \approx 0.27.$$

- Chebyshev:

$$P(|Z| > 3) \leq \frac{1}{9} \approx 0.11.$$

- Chernoff (after using symmetry of the Normal):

$$P(|Z| > 3) = 2P(Z > 3) \leq 2e^{-3t} E(e^{tZ}) = 2e^{-3t} \cdot e^{t^2/2}$$

using the MGF of the standard Normal distribution.

The right-hand side is minimized at $t = 3$, as found by setting the derivative equal to 0 , possibly after taking the log first (which is a good idea since it doesn't affect where the minimum occurs and it means we just have to minimize a quadratic polynomial). Plugging in $t = 3$, we have

$$P(|Z| > 3) \leq 2e^{-9/2} \approx 0.022.$$

All of these upper bounds are correct, but Chernoff's bound is the best by far. This example also illustrates the distinction between a bound and an approximation, as we explained in the introduction to this chapter. Markov's inequality tells us that the tail probability $P(|Z| > 3)$ is at most 0.27 , but it would be a blunder to say that $P(|Z| > 3)$ is approximately 0.27 -we'd be off by a factor of about 100 .

4.10.2 10.2 Law of Large Numbers

We turn next to two theorems, the law of large numbers and the central limit theorem, which describe the behavior of the sample mean of i.i.d. r.v.s as the sample size grows. Throughout this section and the next, assume we have i.i.d. $X_1, X_2, X_3, \dots$ with finite mean μ and finite variance σ^2 . For all positive integers n , let

$$\bar{X}_n = \frac{X_1 + \dots + X_n}{n}$$

be the sample mean of X_1 through X_n . The sample mean is itself an r.v., with mean μ and variance σ^2/n :

$$E(\bar{X}_n) = \frac{1}{n} E(X_1 + \cdots + X_n) = \frac{1}{n} (E(X_1) + \cdots + E(X_n)) = \mu,$$

$$\text{Var}(\bar{X}_n) = \frac{1}{n^2} \text{Var}(X_1 + \cdots + X_n) = \frac{1}{n^2} (\text{Var}(X_1) + \cdots + \text{Var}(X_n)) = \frac{\sigma^2}{n}.$$

The law of large numbers (LLN) says that as n grows, the sample mean $\bar{X}_n$ converges to the true mean μ (in a sense that is explained below). The LLN comes in two versions, "strong" (SLLN) and "weak" (WLLN), which use slightly different definitions of what it means for a sequence of r.v.s to converge to a number. We will state both versions, and prove the second using Chebyshev's inequality.

Theorem 10.2.1 (Strong law of large numbers). The sample mean $\bar{X}_n$ converges to the true mean μ pointwise, with probability 1. Recalling that r.v.s are functions from the sample space S to $\mathbb{R}$, this form of convergence says that $\bar{X}_n(s) \rightarrow \mu$ for each point $s \in S$, except that the convergence is allowed to fail on some set B_0 of exceptions, as long as $P(B_0) = 0$. In short, $P(\bar{X}_n \rightarrow \mu) = 1$.

Theorem 10.2.2 (Weak law of large numbers). For all $\epsilon > 0$, $P(|\bar{X}_n - \mu| > \epsilon) \rightarrow 0$ as $n \rightarrow \infty$. (This form of convergence is called convergence in probability.)

Proof. Fix $\epsilon > 0$. By Chebyshev's inequality,

$$P(|\bar{X}_n - \mu| > \epsilon) \leq \frac{\sigma^2}{n\epsilon^2}.$$

As $n \rightarrow \infty$, the right-hand side goes to 0, and so must the left-hand side. $\square$

The law of large numbers is essential for simulations, statistics, and science. Consider generating "data" from a large number of independent replications of an experiment, performed either by computer simulation or in the real world. Every time we use the average value in the replications of some quantity to approximate its theoretical average, we are implicitly appealing to the LLN.

Example 10.2.3 (Running proportion of Heads). Let $X_1, X_2, \dots$ be i.i.d. $\text{Bern}(1/2)$. Interpreting the X_j as indicators of Heads in a string of fair coin tosses, $\bar{X}_n$ is the proportion of Heads after n tosses. The SLLN says that with probability 1, when the sequence of r.v.s $\bar{X}_1, \bar{X}_2, \bar{X}_3, \dots$ crystallizes into a sequence of numbers, the sequence of numbers will converge to $1/2$. Mathematically, there are bizarre outcomes such as HHHHHH...and HHTHHTHHTHHT..., but collectively they have zero probability of occurring. The WLLN says that for any $\epsilon > 0$, the probability of $\bar{X}_n$ being more than ϵ away from $1/2$ can be made as small as we like by letting n grow.

As an illustration, we simulated six sequences of fair coin tosses and, for each sequence, computed $\bar{X}_n$ as a function of n . Of course, in real life we cannot simulate infinitely many coin tosses, so we stopped after 300 tosses. Figure 10.2 plots $\bar{X}_n$ as a function of n for each sequence.




FIGURE 10.2

Running proportion of Heads in 6 sequences of fair coin tosses. Dashed lines at 0.6 and 0.4 are plotted for reference. As the number of tosses increases, the proportion of Heads approaches $1/2$.

At the beginning, we can see that there is quite a bit of fluctuation in the running proportion of Heads. As the number of coin tosses increases, however, $\text{Var}(\bar{X}_n)$ gets smaller and smaller, and $\bar{X}_n$ approaches $1/2$. $\square$

10.2.4 (LLN does not contradict the fact that a coin is memoryless). In the above example, the law of large numbers states that the proportion of Heads converges to $1/2$, but this does not imply that after a long string of Heads, the coin is "due" for a Tails to balance things out. Rather, the convergence takes place through swamping: past tosses are swamped by the infinitely many tosses that are yet to come.

A sequence of i.i.d. Bernoullis is the simplest possible example of the LLN, but this simple case forms the basis for extremely useful methods in statistics, as the following examples illustrate.

Example 10.2.5 (Monte Carlo integration). Let f be a complicated function whose integral $\int_a^b f(x)dx$ we want to approximate. Assume that $0 \leq f(x) \leq c$ so that we know the integral is finite. On the surface, this problem doesn't involve probability, as $\int_a^b f(x)dx$ is just a number. But where there is no randomness, we can create our own! The technique of Monte Carlo integration uses random samples to obtain approximations of definite integrals when exact integration methods are unavailable.

Let A be the rectangle in the (x, y) -plane given by $a \leq x \leq b$ and $0 \leq y \leq c$. Let B be the region under the curve $y = f(x)$ (and above the x -axis) for $a \leq x \leq b$, so the desired integral is the area of region B . Our strategy will be to take random samples from A , then calculate the proportion of the samples that also fall into the area B . This is depicted in Figure 10.3: points in B are in black, and points not in B are in white.




FIGURE 10.3

Monte Carlo integration. To approximate the area under $f(x)$ from $x = a$ to $x = b$, generate random points in the rectangle $[a, b] \times [0, c]$, and approximate the area under $f(x)$ by the proportion of points falling underneath the curve, multiplied by the total area of the rectangle.

To see why this works, suppose we pick i.i.d. points $(X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n)$ uniformly in the rectangle A . Define indicator r.v.s $I_1, \dots, I_n$ by letting $I_j = 1$ if (X_j, Y_j) is in B and $I_j = 0$ otherwise. Then the I_j are Bernoulli r.v.s whose success probability is precisely the ratio of the area of B to the area of A . Letting $p = E(I_j)$,

$$p = E(I_j) = P(I_j = 1) = \frac{\int_a^b f(x)dx}{c(b-a)}$$

We can estimate p using $\frac{1}{n} \sum_{j=1}^n I_j$, and then estimate the desired integral by

$$\int_a^b f(x)dx \approx c(b-a) \frac{1}{n} \sum_{j=1}^n I_j.$$

Since the I_j are i.i.d. with mean p , it follows from the law of large numbers that with probability 1, the estimate converges to the true value of the integral as the number of points approaches infinity.

Example 10.2.6 (Convergence of empirical CDF). Let $X_1, \dots, X_n$ be i.i.d. random variables with CDF F . For every number x , let $R_n(x)$ count how many of $X_1, \dots, X_n$ are less than or equal to x ; that is,

$$R_n(x) = \sum_{j=1}^n I(X_j \leq x).$$

Since the indicators $I(X_j \leq x)$ are i.i.d. with probability of success $F(x)$, we know $R_n(x)$ is Binomial with parameters n and $F(x)$.

The empirical CDF of $X_1, \dots, X_n$ is defined as

$$\hat{F}_n(x) = \frac{R_n(x)}{n},$$

considered as a function of x . Before we observe $X_1, \dots, X_n$, $\hat{F}_n(x)$ is a random variable for each x . After we observe $X_1, \dots, X_n$, $\hat{F}_n(x)$ crystallizes into a particular value at each x , so $\hat{F}_n$

crystallizes into a particular CDF, which can be used to estimate the true CDF F if the latter is unknown.

For example, suppose $n = 4$ and we observe $X_1 = x_1, X_2 = x_2, X_3 = x_3, X_4 = x_4$. Then the graph of $\frac{R_4(x)}{4}$ starts at 0 and then jumps by $1/4$ every time one of the x_j 's is reached. In other words, $\frac{R_4(x)}{4}$ is the CDF of a discrete random variable taking on values $x_1, \dots, x_4$, each with probability $1/4$. This is illustrated in Figure 10.4.

Now we can ask, what happens to $\hat{F}_n$ as $n \rightarrow \infty$? This is a natural question if we are using $\hat{F}_n$ as an estimate of the true F ; does the approximation fare well in the limit? The law of large numbers provides the answer: for every x , $R_n(x)$ is the sum of n i.i.d. $\text{Bern}(p)$ r.v.s, where $p = F(x)$. So by the SLLN, $\hat{F}_n(x) \rightarrow F(x)$ with probability 1 as $n \rightarrow \infty$.




FIGURE 10.4

Empirical CDF after observing $X_1 = x_1, X_2 = x_2, X_3 = x_3, X_4 = x_4$. The graph jumps by $1/4$ every time one of the x_j 's is reached.

The empirical CDF is commonly used in nonparametric statistics, a branch of statistics that tries to understand a random sample without making strong assumptions about the family of distributions from which it originated. For example, instead of assuming $X_1, \dots, X_n \sim \mathcal{N}(\mu, \sigma^2)$, a nonparametric method would allow $X_1, \dots, X_n \sim F$ for an arbitrary CDF F , then use the empirical CDF as an approximation for F . The law of large numbers is what assures us that this approximation is valid in the limit as we collect more and more samples: at every value of x , the empirical CDF converges to the true CDF.

4.10.3 10.3 Central Limit Theorem

As in the previous section, let $X_1, X_2, X_3, \dots$ be i.i.d. with mean μ and variance σ^2 . The law of large numbers says that as $n \rightarrow \infty$, $\bar{X}_n$ converges to the constant μ (with probability 1). But what is its distribution along the way to becoming a constant? This is addressed by the central limit theorem (CLT), which, as its name suggests, is a limit theorem of central importance in statistics.

The CLT states that for large n , the distribution of $\bar{X}_n$ after standardization approaches a standard Normal distribution. By standardization, we mean that we subtract μ , the mean of $\bar{X}_n$, and divide by $\sigma/\sqrt{n}$, the standard deviation of $\bar{X}_n$.

Theorem 10.3.1 (Central limit theorem). As $n \rightarrow \infty$,

$$\sqrt{n} \left(\frac{\bar{X}_n - \mu}{\sigma} \right) \rightarrow \mathcal{N}(0, 1) \text{ in distribution.}$$

In words, this means that the CDF of the left-hand side converges to Φ , the CDF of the standard Normal distribution.

Proof. We will prove the CLT assuming that the MGF of the X_j exists, though the theorem holds much more generally. Let $M(t) = E(e^{tX_j})$, and without loss of generality let $\mu = 0, \sigma^2 = 1$ (since we end up standardizing $\bar{X}_n$ for the theorem, we might as well standardize the X_j in the first place). Then $M(0) = 1, M'(0) = \mu = 0$, and $M''(0) = \sigma^2 = 1$.

We wish to show that the MGF of $\sqrt{n}\bar{X}_n = (X_1 + \cdots + X_n)/\sqrt{n}$ converges to the MGF of the $\mathcal{N}(0, 1)$ distribution, which is $e^{t^2/2}$. This is a valid strategy because of a theorem that says that if $Z_1, Z_2, \dots$ are r.v.s whose MGFs converge to the MGF of a continuous r.v. Z , then the CDF of Z_n converges to the CDF of Z . (We omit the proof of this result since it requires some difficult analysis. But it should at least seem plausible, in view of the fact that the MGF of an r.v. determines its distribution.)

By properties of MGFs,

$$\begin{aligned} E \left(e^{t(X_1 + \cdots + X_n)/\sqrt{n}} \right) &= E \left(e^{tX_1/\sqrt{n}} \right) E \left(e^{tX_2/\sqrt{n}} \right) \cdots E \left(e^{tX_n/\sqrt{n}} \right) \\ &= \left(M \left(\frac{t}{\sqrt{n}} \right) \right)^n. \end{aligned}$$

Letting $n \rightarrow \infty$, we get the indeterminate form 1^∞ , so instead we should take the limit of the logarithm, $n \log M \left(\frac{t}{\sqrt{n}} \right)$, and then exponentiate at the end. This gives

$$\begin{aligned} \lim_{n \rightarrow \infty} n \log M \left(\frac{t}{\sqrt{n}} \right) &= \lim_{y \rightarrow 0} \frac{\log M(yt)}{y^2} \\ &= \lim_{y \rightarrow 0} \frac{tM'(yt)}{2yM(yt)} \\ &= \frac{t}{2} \lim_{y \rightarrow 0} \frac{M'(yt)}{y} \\ &= \frac{t^2}{2} \lim_{y \rightarrow 0} M''(yt) \\ &= \frac{t^2}{2} \quad \text{by L'Hôpital's rule} \end{aligned}$$

Therefore $\left(M \left(\frac{t}{\sqrt{n}} \right) \right)^n$, the MGF of $\sqrt{n}\bar{X}_n$, approaches $e^{t^2/2}$, the $\mathcal{N}(0, 1)$ MGF.

The CLT is an asymptotic result, telling us about the limiting distribution of $\bar{X}_n$ as $n \rightarrow \infty$, but it also suggests an approximation for the distribution of $\bar{X}_n$ when n is large but finite.

Approximation 10.3.2 (Central limit theorem, approximation form). For large n , the distribution of $\bar{X}_n$ is approximately $\mathcal{N}(\mu, \sigma^2/n)$.

Proof. Change the arrow in the CLT to $\dot{\sim}$, an approximate distribution sign:

$$\sqrt{n} \left(\frac{\bar{X}_n - \mu}{\sigma} \right) \dot{\sim} \mathcal{N}(0, 1)$$

Then by a location-scale transformation,

$$\bar{X}_n \dot{\sim} \mathcal{N}(\mu, \sigma^2/n)$$

Of course, we already knew from properties of expectation and variance that $\bar{X}_n$ has mean μ and variance σ^2/n ; the central limit theorem gives us the additional information that $\bar{X}_n$ is approximately Normal with said mean and variance.

Let's take a moment to admire the generality of this result. The distribution of the individual X_j can be anything in the world, as long as the mean and variance are finite. We could have a discrete distribution like the Binomial, a bounded distribution like the Beta, a skewed

distribution like the Log-Normal, or a distribution with multiple peaks and valleys. No matter what, the act of averaging will cause Normality to emerge. In Figure 10.5 we show histograms of the distribution of $\bar{X}_n$ for 4 different starting distributions and for $n = 1, 5, 30, 100$. As n increases, the distribution of $\bar{X}_n$ starts to look Normal, regardless of the distribution of the X_j .

This does not mean that the distribution of the X_j is irrelevant, however. If the X_j have a highly skewed or multimodal distribution, we may need n to be very large before the Normal approximation becomes accurate; at the other extreme, if the X_j are already i.i.d. Normals, the distribution of $\bar{X}_n$ is exactly $\mathcal{N}(\mu, \sigma^2/n)$ for all n . Since there are no infinite datasets in the real world, the quality of the Normal approximation for finite n is an important consideration.

Example 10.3.3 (Running proportion of Heads, revisited). As in Example 10.2.3, let $X_1, X_2, \dots$ be i.i.d. $\text{Bern}(1/2)$. Before, we used the law of large numbers to conclude that $\bar{X}_n \rightarrow 1/2$ as $n \rightarrow \infty$. Now, using the central limit theorem, we can say more: $E(\bar{X}_n) = 1/2$ and $\text{Var}(\bar{X}_n) = 1/(4n)$, so for large n ,

$$\bar{X}_n \sim \mathcal{N}\left(\frac{1}{2}, \frac{1}{4n}\right)$$

This additional information allows us to quantify what kind of deviation from the mean is typical for a given n . For example, when $n = 100$, $\text{SD}(\bar{X}_n) = 1/20 = 0.05$, so if the Normal approximation is valid, then by the 68 – 95 – 99.7% rule there's a 95% chance that $\bar{X}_n$ will be in the interval $[0.40, 0.60]$.

The CLT says that the sample mean $\bar{X}_n$ is approximately Normal, but since the sum $W_n = X_1 + \dots + X_n = n\bar{X}_n$ is just a scaled version of $\bar{X}_n$, the CLT also implies W_n is approximately Normal. If the X_j have mean μ and variance σ^2 , W_n has mean $n\mu$ and variance $n\sigma^2$. The CLT then states that for large n ,

$$W_n \sim \mathcal{N}(n\mu, n\sigma^2)$$




FIGURE 10.5

Central limit theorem. Histograms of the distribution of $\bar{X}_n$ for different starting distributions of the X_j (indicated by the rows) and increasing values of n (indicated by the columns). Each histogram is based on 10,000 simulated values of $\bar{X}_n$. Regardless of the starting distribution of the X_j , the distribution of $\bar{X}_n$ approaches a Normal distribution as n grows.

This is completely equivalent to the approximation for $\bar{X}_n$, but it can be useful to write it in this form because many of the named distributions we have studied can be considered as a sum of i.i.d. r.v.s. Here are three quick examples.

Example 10.3.4 (Poisson convergence to Normal). Let $Y \sim \text{Pois}(n)$. By Theorem 4.8.1, we can consider Y to be a sum of n i.i.d. $\text{Pois}(1)$ r.v.s. Therefore, for large n ,

$$Y \dot{\sim} \mathcal{N}(n, n).$$

Example 10.3.5 (Gamma convergence to Normal). Let $Y \sim \text{Gamma}(n, \lambda)$. By Theorem 8.4.3, we can consider Y to be a sum of n i.i.d. $\text{Expo}(\lambda)$ r.v.s. Therefore, for large n ,

$$Y \dot{\sim} \mathcal{N}\left(\frac{n}{\lambda}, \frac{n}{\lambda^2}\right)$$

Example 10.3.6 (Binomial convergence to Normal). Let $Y \sim \text{Bin}(n, p)$. By Theorem 3.8.8, we can consider Y to be a sum of n i.i.d. $\text{Bern}(p)$ r.v.s. Therefore, for large n ,

$$Y \dot{\sim} \mathcal{N}(np, np(1-p))$$

This is probably the most widely used Normal approximation in statistics. To account for the discreteness of Y , we write the probability $P(Y = k)$ (which would be exactly 0 under the Normal approximation) as $P(k - 1/2 < Y < k + 1/2)$ (so that it becomes an interval of non-zero width) and apply the Normal approximation to the latter. This is known as the continuity correction, and it yields the following approximation for the PMF of Y :

$$P(Y = k) = P(k - 1/2 < Y < k + 1/2) \approx \Phi\left(\frac{k+1/2-np}{\sqrt{np(1-p)}}\right) - \Phi\left(\frac{k-1/2-np}{\sqrt{np(1-p)}}\right).$$

The Normal approximation to the Binomial distribution is complementary to the Poisson approximation discussed in Chapter 4. The Poisson approximation works best when p is small, while the Normal approximation works best when n is large and p is around $1/2$, so that the distribution of Y is symmetric or nearly so.

We'll conclude with an example that uses both the LLN and the CLT.

Example 10.3.7 (Volatile stock). Each day, a very volatile stock rises 70% or drops 50% in price, with equal probabilities and with different days independent. Let Y_n be the stock price after n days, starting from an initial value of $Y_0 = 100$.

- Explain why $\log Y_n$ is approximately Normal for n large, and state its parameters.
- What happens to $E(Y_n)$ as $n \rightarrow \infty$?
- Use the law of large numbers to find out what happens to Y_n as $n \rightarrow \infty$.

Solution:

(a) We can write $Y_n = Y_0(0.5)^{n-U_n}(1.7)^{U_n}$ where $U_n \sim \text{Bin}(n, \frac{1}{2})$ is the number of times the stock rises in the first n days. This gives

$$\log Y_n = \log Y_0 - n \log 2 + U_n \log 3.4$$

which is a location-scale transformation of U_n . By the CLT, U_n is approximately $\mathcal{N}(\frac{n}{2}, \frac{n}{4})$ for large n , so $\log Y_n$ is approximately Normal with mean

$$E(\log Y_n) = \log 100 - n \log 2 + (\log 3.4) \cdot E(U_n) \approx \log 100 - 0.081n$$

and variance

$$\text{Var}(\log Y_n) = (\log 3.4)^2 \cdot \text{Var}(U_n) \approx 0.374n.$$

(b) We have $E(Y_1) = (170 + 50)/2 = 110$. Similarly,

$$E(Y_{n+1} | Y_n) = \frac{1}{2}(1.7Y_n) + \frac{1}{2}(0.5Y_n) = 1.1Y_n,$$

so

$$E(Y_{n+1}) = E(E(Y_{n+1} | Y_n)) = 1.1E(Y_n).$$

Thus $E(Y_n) = 1.1^n E(Y_0) = 100 \cdot 1.1^n$, which goes to ∞ as $n \rightarrow \infty$.

(c) As in (a), let $U_n \sim \text{Bin}(n, \frac{1}{2})$ be the number of times the stock rises in the first n days. Note that even though $E(Y_n) \rightarrow \infty$, if the stock goes up 70% one day and then drops 50% the next day, then overall it has dropped 15% since $1.7 \cdot 0.5 = 0.85$. So after many days, Y_n will be very small if about half the time the stock rose 70% and about half the time the stock dropped 50%-and the law of large numbers ensures that this will be the case! Writing Y_n in terms of U_n/n in order to apply the LLN, we have

$$Y_n = Y_0(0.5)^{n-U_n}(1.7)^{U_n} = Y_0 \left(\frac{(3.4)^{U_n/n}}{2} \right)^n$$

Since $U_n/n \rightarrow 0.5$ with probability 1, $(3.4)^{U_n/n} \rightarrow \sqrt{3.4} < 2$ with probability 1, so $Y_n \rightarrow 0$ with probability 1.

Paradoxically, $E(Y_n) \rightarrow \infty$ but $Y_n \rightarrow 0$ with probability 1. To gain some intuition for this result, consider the extreme example where a gambler starts with \$100 and each day either quadruples their money or loses their entire fortune, with equal probabilities. Then on average the gambler's wealth doubles each day, which sounds good until one notices that eventually there will be a day when the gambler goes broke. The gambler's actual fortune goes to 0 with probability 1, whereas the expected value goes to infinity due to tiny probabilities of getting extremely large amounts of money, as in the St. Petersburg paradox.

10.3.8 (The evil Cauchy). The central limit theorem requires that the mean and variance of the X_j be finite, and our proof of the WLLN relied on the same conditions. The Cauchy distribution introduced in Example 7.1.25 has no mean or variance, so the Cauchy distribution obeys neither the law of large numbers nor the central limit theorem. It can be shown that the sample mean of n Cauchys is still Cauchy, no matter how large n gets. So the sample mean never approaches a Normal distribution, contrary to the behavior seen in the CLT. There is also no true mean for $\bar{X}_n$ to converge to, so the LLN does not apply either.

4.10.4 10.4 Chi-square and Student- t

We'll round out the chapter by introducing the last two continuous distributions in this book, both of which are closely related to the Normal distribution.

Definition 10.4.1 (Chi-Square distribution). Let $V = Z_1^2 + \cdots + Z_n^2$ where $Z_1, Z_2, \dots, Z_n$ are i.i.d. $\mathcal{N}(0, 1)$. Then V is said to have the Chi-Square distribution with n degrees of freedom. We write this as $V \sim \chi_n^2$.

As it turns out, the χ_n^2 distribution is a special case of the Gamma.

Theorem 10.4.2. The χ_n^2 distribution is the Gamma $(\frac{n}{2}, \frac{1}{2})$ distribution.

Proof. First, we verify that the PDF of $Z_1^2 \sim \chi_1^2$ equals the PDF of the Gamma $(\frac{1}{2}, \frac{1}{2})$: for $x > 0$,

$$F(x) = P(Z_1^2 \leq x) = P(-\sqrt{x} \leq Z_1 \leq \sqrt{x}) = \Phi(\sqrt{x}) - \Phi(-\sqrt{x}) = 2\Phi(\sqrt{x}) - 1,$$

so

$$f(x) = \frac{d}{dx}F(x) = 2\phi(\sqrt{x})\frac{1}{2}x^{-1/2} = \frac{1}{\sqrt{2\pi x}}e^{-x/2}$$

which is indeed the Gamma $(\frac{1}{2}, \frac{1}{2})$ PDF. Then, because $V = Z_1^2 + \cdots + Z_n^2 \sim \chi_n^2$ is the sum of n independent Gamma $(\frac{1}{2}, \frac{1}{2})$ r.v.s, we have $V \sim \text{Gamma}(\frac{n}{2}, \frac{1}{2})$.

From our knowledge of the mean and variance of a Gamma distribution, we have $E(V) = n$ and $\text{Var}(V) = 2n$. We can also obtain the mean and variance using the fact that V is the sum of i.i.d. squared Normals, along with the Normal moments derived in Chapter 6:

$$\begin{aligned} E(V) &= nE(Z_1^2) = n, \\ \text{Var}(V) &= n \text{Var}(Z_1^2) = n(E(Z_1^4) - (EZ_1^2)^2) = n(3 - 1) = 2n. \end{aligned}$$

To get the MGF of the Chi-Square distribution, just plug $n/2$ and $1/2$ into the more general $\text{Gamma}(a, \lambda)$ MGF, which we found in Theorem 8.4.3 to be $(\frac{\lambda}{\lambda-t})^a$ for $t < \lambda$. This gives

$$M_V(t) = \left(\frac{1}{1-2t} \right)^{n/2}, \quad t < 1/2$$

The Chi-Square distribution is important in statistics because it is related to the distribution of the sample variance, which can be used to estimate the true variance of a distribution. When our random variables are i.i.d. Normals, the distribution of the sample variance after appropriate scaling is Chi-Square.

Example 10.4.3 (Distribution of sample variance). For i.i.d. $X_1, \dots, X_n \sim \mathcal{N}(\mu, \sigma^2)$, the sample variance is the r.v.

$$S_n^2 = \frac{1}{n-1} \sum_{j=1}^n (X_j - \bar{X}_n)^2$$

Show that

$$\frac{(n-1)S_n^2}{\sigma^2} \sim \chi_{n-1}^2$$

Solution:

First let's show that $\sum_{j=1}^n (Z_j - \bar{Z}_n)^2 \sim \chi_{n-1}^2$ for standard Normal r.v.s $Z_1, \dots, Z_n$; this is consistent with the more general result we are asked to prove and also serves as a useful stepping stone. Let's start with the following useful identity, which is a special case of the identity from the proof of Theorem 6.3.4:

$$\sum_{j=1}^n Z_j^2 = \sum_{j=1}^n (Z_j - \bar{Z}_n)^2 + n\bar{Z}_n^2$$

Now take the MGF of both sides. By Example 7.5.9, $\sum_{j=1}^n (Z_j - \bar{Z}_n)^2$ and $n\bar{Z}_n^2$ are independent, so the MGF of their sum is the product of the individual MGFs. Also,

$$\sum_{j=1}^n Z_j^2 \sim \chi_n^2 \text{ and } n\bar{Z}_n^2 \sim \chi_1^2$$

so

$$\left(\frac{1}{1-2t} \right)^{n/2} = \left(\text{MGF of } \sum_{j=1}^n (Z_j - \bar{Z}_n)^2 \right) \cdot \left(\frac{1}{1-2t} \right)^{1/2}.$$

This implies

$$\left(\text{MGF of } \sum_{j=1}^n (Z_j - \bar{Z}_n)^2 \right) = \left(\frac{1}{1-2t} \right)^{(n-1)/2}$$

which is the χ_{n-1}^2 MGF. Since the MGF determines the distribution, we have

$$\sum_{j=1}^n (Z_j - \bar{Z}_n)^2 \sim \chi_{n-1}^2$$

For general $X_1, \dots, X_n$, use a location-scale transformation to write $X_j = \mu + \sigma Z_j$ and $\bar{X}_n = \mu + \sigma \bar{Z}_n$. When we express $\sum_{j=1}^n (X_j - \bar{X}_n)^2$ in terms of the Z_j , the μ cancels and the σ comes out squared:

$$\sum_{j=1}^n (X_j - \bar{X}_n)^2 = \sum_{j=1}^n (\mu + \sigma Z_j - (\mu + \sigma \bar{Z}_n))^2 = \sigma^2 \sum_{j=1}^n (Z_j - \bar{Z}_n)^2.$$

All in all,

$$\frac{(n-1)S_n^2}{\sigma^2} = \frac{1}{\sigma^2} \sum_{j=1}^n (X_j - \bar{X}_n)^2 = \frac{1}{\sigma^2} \cdot \sigma^2 \sum_{j=1}^n (Z_j - \bar{Z}_n)^2 \sim \chi_{n-1}^2$$

which is what we wanted. This also implies that $E(S_n^2) = \sigma^2$, which agrees with what we showed in Theorem 6.3.4: the sample variance is unbiased for estimating the true variance.

The Student- t distribution is defined by representing it in terms of a standard Normal r.v. and a χ_n^2 r.v.

Definition 10.4.4 (Student- t distribution). Let

$$T = \frac{Z}{\sqrt{V/n}},$$

where $Z \sim \mathcal{N}(0, 1)$, $V \sim \chi_n^2$, and Z is independent of V . Then T is said to have the Student- t distribution with n degrees of freedom. We write this as $T \sim t_n$. Often "Student- t distribution" is abbreviated to " t distribution".

The Student- t distribution was introduced in 1908 by William Gosset, a Master Brewer at Guinness, while working on quality control for beer. He was required by the company to publish his work under a pseudonym, and he chose the name Student. The t distribution forms the basis for hypothesis testing procedures known as t -tests, which are extremely widely used in practice (we do not introduce the details of t -tests here since they are better left for a course on statistical inference).

The PDF of the Student- t distribution with n degrees of freedom looks similar to that of a standard Normal, except with heavier tails (much heavier if n is small, and not much heavier if n is large). The formula for the PDF is

$$f_T(t) = \frac{\Gamma((n+1)/2)}{\sqrt{n\pi}\Gamma(n/2)} (1 + t^2/n)^{-(n+1)/2}$$

though we will not prove this since the derivation is messy and anyway many of the most important properties of the Student- t distribution are easier to understand by thinking about how we defined it in terms of a Normal and a χ_n^2 r.v., rather than by doing tedious calculations with the PDF. Here are some of these properties.

Theorem 10.4.5 (Student- t properties). The Student- t distribution t_n has the following properties.

1. Symmetry: If $T \sim t_n$, then $-T \sim t_n$ as well.
2. Cauchy as special case: The t_1 distribution is the same as the Cauchy distribution, introduced in Example 7.1.25.
3. Convergence to Normal: As $n \rightarrow \infty$, the t_n distribution approaches the standard Normal distribution.

Proof. In the proof of each property, we appeal to Definition 10.4.4.

1. Express

$$T = \frac{Z}{\sqrt{V/n}},$$

where $Z \sim \mathcal{N}(0, 1)$, $V \sim \chi_n^2$, and Z is independent of V . Then

$$-T = \frac{-Z}{\sqrt{V/n}},$$

where $-Z \sim \mathcal{N}(0, 1)$, so $-T \sim t_n$.

2. Recall that the Cauchy distribution is defined as the distribution of X/Y where X and Y are i.i.d. $\mathcal{N}(0, 1)$. By definition, $T \sim t_1$ can be expressed as $T = Z/\sqrt{V}$, where $\sqrt{V} = \sqrt{Z_1^2} = |Z_1|$ with Z_1 independent of Z . But by symmetry, $Z/|Z_1|$ has the same distribution as Z/Z_1 , and Z/Z_1 is Cauchy. Thus the t_1 and Cauchy distributions are the same.
3. This follows from the SLLN. Consider a sequence of i.i.d. standard Normal r.v.s $Z_1, Z_2, \dots$, and let

$$V_n = Z_1^2 + \dots + Z_n^2.$$

By the SLLN, $V_n/n \rightarrow E(Z_1^2) = 1$ with probability 1. Now let $Z \sim \mathcal{N}(0, 1)$ be independent of all the Z_j , and let

$$T_n = \frac{Z}{\sqrt{V_n/n}}$$

for all n . Then $T_n \sim t_n$ by definition, and since the denominator converges to 1, we have $T_n \rightarrow Z \sim \mathcal{N}(0, 1)$. Therefore, the distribution of the T_n approaches the distribution of Z .

Figure 10.6 plots the Student- t PDF for various values of n , demonstrating all three properties of the above theorem: the PDFs are all symmetric around 0, the PDF for $n = 1$ looks like that of the Cauchy distribution, and as $n \rightarrow \infty$ the heavy tails become lighter, and the PDF approaches the standard Normal PDF.

4.10.5 10.5 Recap

Inequalities and limit theorems are two different ways to handle expectations and probabilities that we don't wish to calculate exactly. Inequalities allow us to obtain lower and/or upper bounds on the unknown value: Cauchy-Schwarz and Jensen give us bounds on expectations,

while Markov, Chebyshev, and Chernoff give us bounds on tail probabilities.




FIGURE 10.6 PDFs of Student- t distribution with (light to dark) $n = 1, 2, 3, 5, 10$ degrees of freedom, as well as the standard Normal PDF (black). As $n \rightarrow \infty$, the Student- t PDFs approach the standard Normal PDF.

Two limit theorems, the law of large numbers and the central limit theorem, describe the behavior of the sample mean $\bar{X}_n$ of i.i.d. $X_1, X_2, \dots$ with mean μ and variance σ^2 . The SLLN says that as $n \rightarrow \infty$, the sample mean $\bar{X}_n$ converges to the true mean μ with probability 1. The CLT says that the distribution of $\bar{X}_n$, after standardization, converges to the standard Normal distribution:

$$\sqrt{n} \left(\frac{\bar{X}_n - \mu}{\sigma} \right) \rightarrow \mathcal{N}(0, 1).$$

This can be translated into an approximation for the distribution of $\bar{X}_n$:

$$\bar{X}_n \sim \mathcal{N}(\mu, \sigma^2/n).$$

Equivalently, we can say that the distribution of the sum $S_n = X_1 + \dots + X_n = n\bar{X}_n$, after standardization, converges to the standard Normal distribution:

$$\frac{S_n - n\mu}{\sigma\sqrt{n}} \rightarrow \mathcal{N}(0, 1).$$

Figure 10.7 illustrates the progression from a distribution to i.i.d. r.v.s with that distribution, from which a sample mean can be formed and studied as an r.v. in its own right. Chebyshev's inequality, the LLN, and the CLT all give important information about the behavior of the sample mean.

The Chi-Square and Student- t distributions are two important named distributions in statistics. The Chi-Square is a special case of the Gamma. The Student- t has a bell-shaped PDF with heavier tails than the Normal, and converges to the standard Normal as the degrees of freedom increase.





FIGURE 10.7

The sample mean $\bar{X}_n$

The sample mean $\bar{X}_n$ of i.i.d. random variables $X_1, \dots, X_n$ is an important quantity in many problems.

Chebyshev's inequality bounds the probability of the sample mean being far from the true mean. The weak law of large numbers, which follows from Chebyshev's inequality, says that for n large, the probability is very high that the sample mean will be very close to the true mean. The central limit theorem says that for n large the distribution of the sample mean will be approximately Normal.

Here, one last time, is the diagram of relationships between the named distributions, updated to include the Chi-Square distribution (as a special case of the Gamma) and Student- t distribution (with the Cauchy as a special case). We have also added arrows to show the convergence of the Poisson, Gamma, and Student- t distributions to Normality; the first two are a consequence of the central limit theorem, and the third is a consequence of the law of large numbers.



Now we see that all the named distributions are connected to one another!

4.10.6 10.6 R

Jensen's Inequality

R makes it easy to compare the expectations of X and $g(X)$ for a given choice of g , and this allows us to verify some special cases of Jensen's inequality. For example, suppose we simulate 10^4 times from the $\text{Expo}(1)$ distribution:

```
x <- rexp(10^4)
```

According to Jensen's inequality, $E(\log X) \leq \log EX$. The former can be approximated by $\text{mean}(\log(x))$ and the latter can be approximated by $\log(\text{mean}(x))$, so compute both:

```
mean(log(x))
log(mean(x))
```

For the $\text{Expo}(1)$ distribution, we find that $\text{mean}(\log(x))$ is approximately -0.6 (the true value is around -0.577), while $\log(\text{mean}(x))$ is approximately 0 (the true value is 0). This indeed suggests $E(\log X) \leq \log EX$. We could also compare $\text{mean}(x^3)$ to $\text{mean}(x)^3$, or $\text{mean}(\sqrt{x})$ to $\sqrt{\text{mean}(x)}$ - the possibilities are endless.

Visualization of the law of Large Numbers

To plot the running proportion of Heads in a sequence of independent fair coin tosses, we first generate the coin tosses themselves:

```
nsim <- 300
p <- 1/2
x <- rbinom(nsim,1,p)
```

Then we compute $\bar{X}_n$ for each value of n and store the results in `xbar`:

```
xbar <- cumsum(x)/(1:nsim)
```

The above line of code performs elementwise division of the two vectors `cumsum(x)` and `1:nsim`. Finally, we plot `xbar` against the number of coin tosses:

```
plot(1:nsim,xbar,type="l",ylim=c(0,1))
```

You should see that the values of `xbar` approach `p`, by the LLN.

Monte Carlo Estimate of π

A famous example of Monte Carlo integration is the Monte Carlo estimate of π . The unit disk $\{(x, y) : x^2 + y^2 \leq 1\}$ is inscribed in the square $[-1, 1] \times [-1, 1]$, which has area 4. If we generate a large number of points that are Uniform on the square, the proportion of points falling inside the disk is approximately equal to the ratio of the disk's area to the square's area, which is $\pi/4$. Thus, to estimate π we can take the proportion of points inside the circle and multiply by 4.

In R, to generate Uniform points on the 2D square, we can independently generate the x -coordinate and the y -coordinate as $\text{Unif}(-1, 1)$ r.v.s, using the results of Example 7.1.23:

```
nsim <- 10^6
```




FIGURE 10.8 Monte Carlo estimate of π : Generate points that are Uniform on the 2D square $[-1, 1] \times [-1, 1]$, which has area 4 . The proportion of points falling in the unit disk is approximately $\pi/4$.

```
x <- runif(nsim,-1,1)
y <- runif(nsim,-1,1)
```

To count the number of points in the disk, we type `sum(x^2 + y^2 < 1)`. The vector `x^2 + y^2 < 1` is an indicator vector whose i th element is 1 if the i th point falls inside the disk and 0 otherwise, so the sum of the elements is the number of points in the disk. To get our estimate of π , we convert the sum into a proportion and multiply by 4 :

```
4*sum(x^2+y^2<1)/nsim}
```

How close was your estimate to the actual value of π ?

Visualizations of the Central Limit Theorem

One way to visualize the central limit theorem for a distribution of interest is to plot the distribution of $\bar{X}_n$ for various values of n , as in Figure 10.5. To do this, we first have to generate i.i.d. $X_1, \dots, X_n$ a bunch of times from our distribution of interest. For example, suppose that our distribution of interest is $\text{Unif}(0, 1)$, and we are interested in the distribution of $\bar{X}_{12}$, i.e., we set $n = 12$. In the following code, we create a matrix of i.i.d. standard Uniforms. The matrix has 12 columns, corresponding to X_1 through X_{12} . Each row of the matrix is a different realization of X_1 through X_{12} .

```
nsim <- 10^4
n <- 12
x <- matrix(runif(n*nsim), nrow=nsim, ncol=n)
```

Now, to obtain realizations of $\bar{X}_{12}$, we simply take the average of each row of the matrix x ; we can do this with the `rowMeans` function:

```
xbar <- rowMeans(x)
```

Finally, we create a histogram:

```
hist(xbar)
```

You should see a histogram that looks approximately Normal. Since the $\text{Unif}(0, 1)$ distribution is symmetric, the CLT kicks in quickly and the Normal approximation for $\bar{X}_n$ works well, even for $n = 12$. Changing `runif` to `rexp`, we see that for $X_j \sim \text{Expo}(1)$, the distribution of $\bar{X}_n$ remains skewed when $n = 12$, so a larger value of n is required before the Normal approximation is decent.

Another neat visualization of the CLT can be found in the `animation` package. This package has a built-in animation of a quincunx or bean machine, invented by the statistician and geneticist Francis Galton to illustrate the Normal distribution. After installing the package, try:

```
library(animation)
quincunx()
```

Can you use the central limit theorem to explain why the histogram produced by a quincunx should look approximately Normal?

Chi-square and Student- t Distributions

Although the Chi-Square is just a special case of the Gamma, it still has its own functions `dchisq`, `pchisq`, and `rchisq` in R: `dchisq(x, n)` and `pchisq(x, n)` return the values of the χ_n^2 PDF and CDF at x , and `rchisq(nsim, n)` generates `nsim` i.i.d. χ_n^2 r.v.s.

The Student- t distribution has functions `dt`, `pt`, and `rt`. To evaluate the PDF or CDF of the t_n distribution at x , we use `dt(x, n)` or `pt(x, n)`. To generate `nsim` i.i.d. r.v.s from the t_n distribution, we use `rt(nsim, n)`. Of course, `dt(x, 1)` is the same as `dcauchy(x)`.

4.10.7 10.7 Exercises

Exercises marked with (S) have detailed solutions at <http://stat110.net>.

Inequalities

1. (S) In a national survey, a random sample of people are chosen and asked whether they support a certain policy. Assume that everyone in the population is equally likely to be surveyed at each step, and that the sampling is with replacement (sampling without replacement is typically more realistic, but with replacement will be a good approximation if the sample size is small compared to the population size). Let n be the sample size, and let $\hat{p}$ and p be the proportion of people who support the policy in the sample and in the entire population, respectively. Show that for every $c > 0$,

$$P(|\hat{p} - p| > c) \leq \frac{1}{4nc^2}$$

2. (S) For i.i.d. r.v.s $X_1, \dots, X_n$ with mean μ and variance σ^2 , give a value of n (as a specific number) that will ensure that there is at least a 99% chance that the sample mean will be within 2 standard deviations of the true mean μ .
3. S Show that for any two positive r.v.s X and Y with neither a constant multiple of the other,

$$E(X/Y)E(Y/X) > 1.$$

4. (S) The famous arithmetic mean-geometric mean inequality says that for any positive numbers $a_1, a_2, \dots, a_n$,

$$\frac{a_1 + a_2 + \dots + a_n}{n} \geq (a_1 a_2 \dots a_n)^{1/n}.$$

Show that this inequality follows from Jensen's inequality, by considering $E \log(X)$ for an r.v. X whose possible values are $a_1, \dots, a_n$ (you should specify the PMF of X ; if you want, you can assume that the a_j are distinct (no repetitions), but be sure to say so if you assume this).

5. (S) Let X be a discrete r.v. whose distinct possible values are $x_0, x_1, \dots$, and let $p_k = P(X = x_k)$. The entropy of X is $H(X) = \sum_{k=0}^{\infty} p_k \log_2(1/p_k)$.

- (a) Find $H(X)$ for $X \sim \text{Geom}(p)$.

Hint: Use properties of logs, and interpret part of the sum as an expected value.

- (b) Let X and Y be i.i.d. discrete r.v.s. Show that $P(X = Y) \geq 2^{-H(X)}$.

Hint: Consider $E(\log_2(W))$, where W is an r.v. taking value p_k with probability p_k .

6. Let X be a random variable with mean μ and variance σ^2 . Show that

$$E(X - \mu)^4 \geq \sigma^4$$

and use this to show that the kurtosis of X is at least -2.

Fill-in-the-blank Inequalities

7. (S) Let X and Y be i.i.d. positive r.v.s, and let $c > 0$. For each part below, fill in the appropriate equality or inequality symbol: write $=$ if the two sides are always equal, $\leq$ if the left-hand side is less than or equal to the right-hand side (but they are not necessarily equal), and similarly for $\geq$. If no relation holds in general, write $?$.

- (a) $E(\log(X))$ _____ $\log(E(X))$
 (b) $E(X) - \sqrt{E(X^2)}$ _____
 (c) $E(\sin^2(X)) + E(\cos^2(X))$ _____ 1
 (d) $E(|X|)$ _____ $\sqrt{E(X^2)}$
 (e) $P(X > c)$ _____ $\frac{E(X^3)}{c^3}$
 (f) $P(X \leq Y)$ _____ $P(X \geq Y)$
 (g) $E(XY)$ _____ $\sqrt{E(X^2)E(Y^2)}$
 (h) $P(X + Y > 10)$ _____ $P(X > 5 \text{ or } Y > 5)$
 (i) $E(\min(X, Y))$ _____ $\min(EX, EY)$
 (j) $E(X/Y)$ _____ $\frac{EX}{EY}$
 (k) $E(X^2(X^2 + 1))$ _____ $E(X^2(Y^2 + 1))$
 (l) $E\left(\frac{X^3}{X^3 + Y^3}\right)$ _____ $E\left(\frac{Y^3}{X^3 + Y^3}\right)$

8. (S) Write the most appropriate of $\leq, \geq, =$, or $?$ in the blank for each part (where "?" means that no relation holds in general).

In (c) through (f), X and Y are i.i.d. (independent identically distributed) positive random variables. Assume that the various expected values exist.

- (a) (probability that a roll of 2 fair dice totals 9) _____ (probability that a roll of 2 fair dice totals 10)
 (b) (probability that at least 65% of 20 fair coin flips land Heads) _____ (probability that at least 65% of 2000 fair coin flips land Heads)
 (c) $E(\sqrt{X})$ _____ $\sqrt{E(X)}$
 (d) $E(\sin X)$ _____ $\sin(EX)$
 (e) $P(X + Y > 4)$ _____ $P(X > 2)P(Y > 2)$
 (f) $E((X + Y)^2)$ _____ $2E(X^2) + 2(EX)^2$

9. Let X and Y be i.i.d. continuous r.v.s. Assume that the various expressions below exist. Write the most appropriate of $\leq, \geq, =$, or $?$ in the blank for each part (where "?" means that no relation holds in general).
- $e^{-E(X)} \text{ \_\_\_\_\_\_ } E(e^{-X})$
 - $P(X > Y + 3) \text{ \_\_\_\_\_\_ } P(Y > X + 3)$
 - $P(X > Y + 3) \text{ \_\_\_\_\_\_ } P(X > Y - 3)$
 - $E(X^4) \text{ \_\_\_\_\_\_ } (E(XY))^2$
 - $\text{Var}(Y) \text{ \_\_\_\_\_\_ } E(\text{Var}(Y | X))$
 - $P(|X + Y| > 3) \text{ \_\_\_\_\_\_ } E|X|$
10. (S) Let X and Y be positive random variables, not necessarily independent. Assume that the various expected values below exist. Write the most appropriate of $\leq, \geq, =$, or $?$ in the blank for each part (where "?" means that no relation holds in general).
- $(E(XY))^2 \text{ \_\_\_\_\_\_ } E(X^2) E(Y^2)$
 - $P(|X + Y| > 2) \text{ \_\_\_\_\_\_ } \frac{1}{10} E((X + Y)^4)$
 - $E(\log(X + 3)) \text{ \_\_\_\_\_\_ } \log(E(X + 3))$
 - $E(X^2 e^X) \text{ \_\_\_\_\_\_ } E(X^2) E(e^X)$
 - $P(X + Y = 2) \text{ \_\_\_\_\_\_ } P(X = 1)P(Y = 1)$
 - $P(X + Y = 2) \text{ \_\_\_\_\_\_ } P(\{X \geq 1\} \cup \{Y \geq 1\})$
11. (S) Let X and Y be positive random variables, not necessarily independent. Assume that the various expected values below exist. Write the most appropriate of $\leq, \geq, =$, or $?$ in the blank for each part (where "?" means that no relation holds in general).
- $E(X^3) \text{ \_\_\_\_\_\_ } \sqrt{E(X^2) E(X^4)}$
 - $P(|X + Y| > 2) \text{ \_\_\_\_\_\_ } \frac{1}{16} E((X + Y)^4)$
 - $E(\sqrt{X + 3}) \text{ \_\_\_\_\_\_ } \sqrt{E(X + 3)}$
 - $E(\sin^2(X)) + E(\cos^2(X)) \text{ \_\_\_\_\_\_ } 1$
 - $E(Y | X + 3) \text{ \_\_\_\_\_\_ } E(Y | X)$
 - $E(E(Y^2 | X)) \text{ \_\_\_\_\_\_ } (EY)^2$
12. (S) Let X and Y be positive random variables, not necessarily independent. Assume that the various expressions below exist. Write the most appropriate of $\leq, \geq, =$, or $?$ in the blank for each part (where "?" means that no relation holds in general).
- $P(X + Y > 2) \text{ \_\_\_\_\_\_ } \frac{EX + EY}{2}$
 - $P(X + Y > 3) \text{ \_\_\_\_\_\_ } P(X > 3)$
 - $E(\cos(X)) \text{ \_\_\_\_\_\_ } \cos(EX)$
 - $E(X^{1/3}) \text{ \_\_\_\_\_\_ } (EX)^{1/3}$
 - $E(X^Y) \text{ \_\_\_\_\_\_ } (EX)^{EY}$
 - $E(E(X | Y) + E(Y | X)) \text{ \_\_\_\_\_\_ } EX + EY$
13. (S) Let X and Y be i.i.d. positive random variables. Assume that the various expressions below exist. Write the most appropriate of $\leq, \geq, =$, or $?$ in the blank for each part (where "?" means that no relation holds in general).
- $E(e^{X+Y}) \text{ \_\_\_\_\_\_ } e^{2E(X)}$
 - $E(X^2 e^X) \text{ \_\_\_\_\_\_ } \sqrt{E(X^4) E(e^{2X})}$
 - $E(X | 3X) \text{ \_\_\_\_\_\_ } E(X | 2X)$
 - $E(X^7 Y) \text{ \_\_\_\_\_\_ } E(X^7 E(Y | X))$
 - $E\left(\frac{X}{Y} + \frac{Y}{X}\right) \text{ \_\_\_\_\_\_ } 2$
 - $P(|X - Y| > 2) \text{ \_\_\_\_\_\_ } \frac{\text{Var}(X)}{2}$
14. (S) Let X and Y be i.i.d. Gamma $(\frac{1}{2}, \frac{1}{2})$, and let $Z \sim \mathcal{N}(0, 1)$ (note that X and Z may be dependent, and Y and Z may be dependent). For (a),(b),(c), write the most appropriate

- of $<$, $>$, $=$, or $?$ in each blank; for (d),(e),(f), write the most appropriate of $\leq$, $\geq$, $=$, or $?$ in each blank.
- (a) $P(X < Y)$ _____ $1/2$
 (b) $P(X = Z^2)$ _____ 1
 (c) $P\left(Z \geq \frac{1}{X^4 + Y^4 + 7}\right)$ _____ 1
 (d) $E\left(\frac{X}{X+Y}\right) E((X+Y)^2)$ _____ $E(X^2) + (E(X))^2$
 (e) $E(X^2 Z^2)$ _____ $\sqrt{E(X^4) E(X^2)}$
 (f) $E((X + 2Y)^4)$ _____ 3^4
15. Let X, Y, Z be i.i.d. $\mathcal{N}(0, 1)$ r.v.s. Write the most appropriate of $\leq$, $\geq$, $=$, or $?$ in each blank (where "?" means that no relation holds in general).
- (a) $P(X^2 + Y^2 + Z^2 > 6)$ _____ $1/2$
 (b) $P(X^2 < 1)$ _____ $2/3$
 (c) $E\left(\frac{X^2}{X^2 + Y^2 + Z^2}\right)$ _____ $1/4$
 (d) $\text{Var}(\Phi(X) + \Phi(Y) + \Phi(Z))$ _____ $1/4$
 (e) $E(e^{-X})$ _____ $E(e^X)$
 (f) $E(|X|e^X)$ _____ $\sqrt{E(e^{2X})}$
16. Let X, Y, Z, W be i.i.d. positive r.v.s with CDF F and $E(X) = 1$. Write the most appropriate of $\leq$, $\geq$, $=$, or $?$ in each blank (where "?" means that no relation holds in general).
- (a) $F(3)$ _____ $2/3$
 (b) $(F(3))^3$ _____ $P(X + Y + Z \leq 9)$
 (c) $E\left(\frac{X^2}{X^2 + Y^2 + Z^2 + W^2}\right)$ _____ $1/4$
 (d) $E(XYZW)$ _____ $E(X^4)$
 (e) $\text{Var}(E(Y | X))$ _____ $\text{Var}(Y)$
 (f) $\text{Cov}(X + Y, X - Y)$ _____ 0
17. Let X, Y, Z be i.i.d. (independent, identically distributed) r.v.s with a continuous distribution. Write the most appropriate of $\leq$, $\geq$, $=$, or $?$ in each blank (where "?" means that no relation holds in general).
- (a) $P(X < Y < Z)$ _____ $1/6$
 (b) $P(X > 1)$ _____ $E(X)$
 (c) $P\left(\sum_{k=0}^{2015} \left(\frac{X^2+1}{X^2+2}\right)^k > 3\right) - P(X^2 > 1)$
 (d) $E(\sqrt{X^2 + Y^2})$ _____ $\sqrt{E(X^2) + E(Y^2)}$
 (e) $\text{Var}(Y^2 | Z)$ _____ $\text{Var}(X^2 | X)$
 (f) $\text{Var}(X - 2Y + 3Z)$ _____ $14 \text{Var}(X)$
18. Let $X, Y, Z \sim \text{Bin}(n, p)$ be i.i.d. (independent and identically distributed). Write the most appropriate of $\leq$, $\geq$, $=$, or $?$ in each blank (where "?" means that no relation holds in general).
- (a) $P(X < Y < Z)$ _____ $1/6$
 (b) $P(X + Y + Z > n)$ _____ $3p$
 (c) $P\left(\sum_{n=0}^{2016} \frac{(X^2+1)^n}{n!} > e^5\right) - P(X > 2)$
 (d) $E(e^X)$ _____ $(pe + 1 - p)^n$
 (e) $\text{Var}(X + Y)$ _____ $n/2$
 (f) $E(X | X + Y = n)$ _____ $n/2$

19. Let $X, Y, Z, W \sim \mathcal{N}(c, c^2)$ be i.i.d., where $c > 0$. Write the most appropriate of $\leq, \geq, =$, or $?$ in each blank (where "?" means that no relation holds in general).
- (a) $P(X + Y \leq Z - W)$ _____ $\Phi(-1)$
 - (b) $P(X^4 - Y^8 \leq Z^4 - W^8)$ _____ $\Phi(-1)$
 - (c) $E(X - c)^2$ _____ c
 - (d) $E(X - c)^3$ _____ c
 - (e) $E\left(\frac{1}{4}X^2 + X\right)$ _____ $\frac{1}{2}c^2 + c$
 - (f) $\log(E(e^X))$ _____ $\frac{1}{2}c^2 + c$
20. Let $X, Y \sim \text{Pois}(\lambda)$ be i.i.d., where $\lambda > 1$. Write the most appropriate of $\leq, \geq, =$, or $?$ in each blank (where "?" means that no relation holds in general).
- (a) $P(X \leq Y)$ _____ $1/2$
 - (b) $P(X + Y \leq 1)$ _____ $3e^{-2\lambda}$
 - (c) $E(e^{X+Y})$ _____ $e^{2\lambda}$
 - (d) $E(X \mid X + Y = 4)$ _____ 2
 - (e) $\text{Var}(X \mid X + Y = 4)$ _____ 1
 - (f) $E(X^2 - Y)$ _____ λ^2

LLN and CLT

21. (S) Let $X_1, X_2, \dots$ be i.i.d. positive random variables with mean 2. Let $Y_1, Y_2, \dots$ be i.i.d. positive random variables with mean 3. Show that

$$\frac{X_1 + X_2 + \dots + X_n}{Y_1 + Y_2 + \dots + Y_n} \rightarrow \frac{2}{3}$$

with probability 1. Does it matter whether the X_i are independent of the Y_j ?

22. (S) Let $U_1, U_2, \dots, U_{60}$ be i.i.d. $\text{Unif}(0, 1)$ and $X = U_1 + U_2 + \dots + U_{60}$.
- (a) Which important distribution is the distribution of X very close to? Specify what the parameters are, and state which theorem justifies your choice.
 - (b) Give a simple but accurate approximation for $P(X > 17)$. Justify briefly.
23. (S) Let $V_n \sim \chi_n^2$ and $T_n \sim t_n$ for all positive integers n .
- (a) Find numbers a_n and b_n such that $a_n(V_n - b_n)$ converges in distribution to $\mathcal{N}(0, 1)$.
 - (b) Show that $T_n^2 / (n + T_n^2)$ has a Beta distribution (without using calculus).
24. (S) Let $T_1, T_2, \dots$ be i.i.d. Student- t r.v.s with $m \geq 3$ degrees of freedom. Find constants a_n and b_n (in terms of m and n) such that $a_n(T_1 + T_2 + \dots + T_n - b_n)$ converges to $\mathcal{N}(0, 1)$ in distribution as $n \rightarrow \infty$.
25. (S) (a) Let $Y = e^X$, with $X \sim \text{Expo}(3)$. Find the mean and variance of Y .
- (b) For $Y_1, \dots, Y_n$ i.i.d. with the same distribution as Y from (a), what is the approximate distribution of the sample mean $\bar{Y}_n = \frac{1}{n} \sum_{j=1}^n Y_j$ when n is large?
26. (S) (a) Explain why the $\text{Pois}(n)$ distribution is approximately Normal if n is a large positive integer (specifying what the parameters of the Normal are).
- (b) Stirling's formula is an amazingly accurate approximation for factorials:

$$n! \approx \sqrt{2\pi n} \left(\frac{n}{e}\right)^n$$

where in fact the ratio of the two sides goes to 1 as $n \rightarrow \infty$. Use (a) to give a quick heuristic derivation of Stirling's formula by using a Normal approximation to the probability that a $\text{Pois}(n)$ r.v. is n , with the continuity correction: first write $P(N = n) = P(n - \frac{1}{2} < N < n + \frac{1}{2})$, where $N \sim \text{Pois}(n)$.

27. (S) (a) Consider i.i.d. $\text{Pois}(\lambda)$ r.v.s $X_1, X_2, \dots$. The MGF of X_j is $M(t) = e^{\lambda(e^t - 1)}$. Find

the MGF $M_n(t)$ of the sample mean $\bar{X}_n = \frac{1}{n} \sum_{j=1}^n X_j$.

(b) Find the limit of $M_n(t)$ as $n \rightarrow \infty$. (You can do this with almost no calculation using a relevant theorem; or you can use (a) and the fact that $e^x \approx 1 + x$ if x is very small.)

28. Let $X_n \sim \text{Pois}(n)$ for all positive integers n . Use MGFs to show that the distribution of the standardized version of X_n converges to a Normal distribution as $n \rightarrow \infty$, without invoking the CLT.

29. An important concept in frequentist statistics is that of a confidence interval (CI). Suppose we observe data X from a distribution with parameter θ . Unlike in Bayesian statistics, θ is treated as a fixed but unknown constant; it is not given a prior distribution. A 95% confidence interval consists of a lower bound $L(X)$ and upper bound $U(X)$ such that

$$P(L(X) < \theta < U(X)) = 0.95$$

for all possible values of θ . Note that in the above statement, $L(X)$ and $U(X)$ are random variables, as they are functions of the r.v. X , whereas θ is a constant. The definition says that the random interval $(L(X), U(X))$ has a 95% chance of containing the true value of θ .

Imagine an army of frequentists all over the world, independently generating 95% CIs. The j th frequentist observes data X_j and makes a confidence interval for the parameter θ_j . Show that if there are n of these frequentists, then the fraction of their intervals which contain the corresponding parameter approaches 0.95 as $n \rightarrow \infty$.

Hint: Consider the indicator r.v. $I_j = I(L(X_j) < \theta_j < U(X_j))$.

30. This problem extends Example 10.3.7 to a more general setting. Again, suppose a very volatile stock rises 70% or drops 50% in price each day, with equal probabilities and with different days independent.

(a) Suppose a hedge fund manager always invests half of her current fortune into the stock each day. Let Y_n be her fortune after n days, starting from an initial fortune of $Y_0 = 100$. What happens to Y_n as $n \rightarrow \infty$?

(b) More generally, suppose the hedge fund manager always invests a fraction α of her current fortune into the stock each day (in Part (a), we took $\alpha = 1/2$). With Y_0 and Y_n defined as in Part (a), find the function $g(\alpha)$ such that

$$\frac{\log Y_n}{n} \rightarrow g(\alpha)$$

with probability 1 as $n \rightarrow \infty$, and prove that $g(\alpha)$ is maximized when $\alpha = 2/7$.

Mixed Practice

31. As in Exercise 36 from Chapter 3, there are n voters in an upcoming election in a certain country, where n is a large, even number. There are two candidates, A and B. Each voter chooses randomly whom to vote for, independently and with equal probabilities.

(a) Use a Normal approximation (with continuity correction) to get an approximation for the probability of a tie, in terms of Φ .

(b) Use a first-order Taylor expansion (linear approximation) to the approximation from Part (a) to show that the probability of a tie is approximately $1/\sqrt{cn}$, where c is a constant (which you should specify).

32. Cassie enters a casino with $X_0 = 1$ dollar and repeatedly plays the following game: with probability $1/3$, the amount of money she has increases by a factor of 3; with probability $2/3$, the amount of money she has decreases by a factor of 3. Let X_n be the amount of money she has after playing this game n times. For example, X_1 is 3 with probability $1/3$ and is 3^{-1} with probability $2/3$.

(a) Compute $E(X_1)$, $E(X_2)$ and, in general, $E(X_n)$.

- (b) What happens to $E(X_n)$ as $n \rightarrow \infty$?
- (c) Let Y_n be the number of times out of the first n games that Cassie triples her money. What happens to Y_n/n as $n \rightarrow \infty$?
- (d) What happens to X_n as $n \rightarrow \infty$?

33. A handy rule of thumb in statistics and life is as follows:

Conditioning Often Makes Things Better.

This problem explores how the above rule of thumb applies to estimating unknown parameters. Let θ be an unknown parameter that we wish to estimate based on data $X_1, X_2, \dots, X_n$ (these are r.v.s before being observed, and then after the experiment they "crystallize" into data). In this problem, θ is viewed as an unknown constant, and is not treated as an r.v. as in the Bayesian approach. Let T_1 be an estimator for θ (this means that T_1 is a function of $X_1, \dots, X_n$ which is being used to estimate θ).

A strategy for improving T_1 (in some problems) is as follows. Suppose that we have an r.v. R such that $T_2 = E(T_1 | R)$ is a function of $X_1, \dots, X_n$ (in general, $E(T_1 | R)$ might involve unknowns such as θ but then it couldn't be used as an estimator). Also suppose that $P(T_1 = T_2) < 1$, and that $E(T_1^2)$ is finite.

(a) Use Jensen's inequality to show that T_2 is better than T_1 in the sense that the mean squared error is less, i.e.,

$$E(T_2 - \theta)^2 < E(T_1 - \theta)^2.$$

Hint: Use Adam's law on the right-hand side.

(b) The bias of an estimator T for θ is defined to be $b(T) = E(T) - \theta$. An important identity in statistics, a form of the bias-variance tradeoff, is that mean squared error is variance plus squared bias:

$$E(T - \theta)^2 = \text{Var}(T) + (b(T))^2$$

Use this identity and Eve's law to give an alternative proof of the result from (a).

(c) Now suppose that $X_1, X_2, \dots$ are i.i.d. with mean θ , and consider the special case $T_1 = X_1, R = \sum_{j=1}^n X_j$. Find T_2 in simplified form, and check that it has lower mean squared error than T_1 for $n \geq 2$. Also, say what happens to T_1 and T_2 as $n \rightarrow \infty$.

34. Each page of an n -page book has a $\text{Pois}(\lambda)$ number of typos, where λ is unknown (but is not treated as an r.v.). Typos on different pages are independent. Thus we have i.i.d. $X_1, \dots, X_n \sim \text{Pois}(\lambda)$, where X_j is the number of typos on page j . Suppose we are interested in estimating the probability θ that a page has no typos:

$$\theta = P(X_j = 0) = e^{-\lambda}$$

(a) Let $\bar{X}_n = \frac{1}{n}(X_1 + \dots + X_n)$. Show that the estimator $T_n = e^{-\bar{X}_n}$ is biased for θ . (That is, show $E(T_n) \neq \theta$. Estimators and bias are defined in the previous problem.)

(b) Show that as $n \rightarrow \infty, T_n \rightarrow \theta$ with probability 1.

(c) Show that $W = \frac{1}{n}(I(X_1 = 0) + \dots + I(X_n = 0))$ is unbiased for θ . Using the fact that $X_1 | (X_1 + \dots + X_n = s) \sim \text{Bin}(s, 1/n)$, find $E(W | X_1 + \dots + X_n)$. Is the estimator $\tilde{W} = E(W | X_1 + \dots + X_n)$ also unbiased for θ ?

(d) Using Eve's law or otherwise, show that $\tilde{W}$ has lower variance than W , and relate this to the previous question.

35. A binary sequence is being generated through some process (random or deterministic). You need to sequentially predict each new number, i.e., you predict whether the next number will be 0 or 1, then observe it, then predict the next number, etc. Each of your predictions can be

based on the entire past history of the sequence.

(a) Suppose for this part that the binary sequence consists of i.i.d. $\text{Bern}(p)$ r.v.s, with p known. What is your optimal strategy (for each prediction, your goal is to maximize the probability of being correct)? What is the probability that you will guess the n th value correctly with this strategy?

(b) Now suppose that the binary sequence consists of i.i.d. $\text{Bern}(p)$ r.v.s, with p unknown. Consider the following strategy: say 1 as your first prediction; after that, say " 1 " if the proportion of 1 's so far is at least $1/2$, and say " 0 " otherwise. Find the limit as $n \rightarrow \infty$ of the probability of guessing the n th value correctly (in terms of p).

(c) Now suppose that you follow the strategy from (b), but that the binary sequence is generated by a nefarious entity who knows your strategy. What can the entity do to make your guesses be wrong as often as possible?

36. (S) Let X and Y be independent standard Normal r.v.s and let $R^2 = X^2 + Y^2$ (where $R > 0$ is the distance from (X, Y) to the origin).

(a) The distribution of R^2 is an example of three of the important distributions we have seen. State which three of these distributions R^2 is an instance of, specifying the parameter values.

(b) Find the PDF of R .

Hint: Start with the PDF $f_W(w)$ of $W = R^2$.

(c) Find $P(X > 2Y + 3)$ in terms of the standard Normal CDF Φ .

(d) Compute $\text{Cov}(R^2, X)$. Are R^2 and X independent?

37. (S) Let $Z_1, \dots, Z_n \sim \mathcal{N}(0, 1)$ be i.i.d.

(a) As a function of Z_1 , create an $\text{Expo}(1)$ r.v. X (your answer can also involve the standard Normal CDF Φ).

(b) Let $Y = e^{-R}$, where $R = \sqrt{Z_1^2 + \dots + Z_n^2}$. Write down (but do not evaluate) an integral for $E(Y)$.

(c) Let $X_1 = 3Z_1 - 2Z_2$ and $X_2 = 4Z_1 + 6Z_2$. Determine whether X_1 and X_2 are independent (be sure to mention which results you're using).

38. (S) Let $X_1, X_2, \dots$ be i.i.d. positive r.v.s. with mean μ , and let $W_n = \frac{X_1}{X_1 + \dots + X_n}$.

(a) Find $E(W_n)$.

Hint: Consider $\frac{X_1}{X_1 + \dots + X_n} + \frac{X_2}{X_1 + \dots + X_n} + \dots + \frac{X_n}{X_1 + \dots + X_n}$.

(b) What random variable does nW_n converge to (with probability 1) as $n \rightarrow \infty$?

(c) For the case that $X_j \sim \text{Expo}(\lambda)$, find the distribution of W_n , preferably without using calculus. (If it is one of the named distributions we have studied, state its name and specify the parameters; otherwise, give the PDF.)

39. Let $X_1, X_2, \dots$ be i.i.d. $\text{Expo}(1)$.

(a) Let $N = \min\{n : X_n \geq 1\}$ be the index of the first X_j to exceed 1. Find the distribution of $N - 1$ (give the name and parameters), and hence find $E(N)$.

(b) Let $M = \min\{n : X_1 + X_2 + \dots + X_n \geq 10\}$ be the number of X_j 's we observe until their sum exceeds 10 for the first time. Find the distribution of $M - 1$ (give the name and parameters), and hence find $E(M)$.

Hint: Consider a Poisson process.

(c) Let $\bar{X}_n = (X_1 + \dots + X_n)/n$. Find the exact distribution of $\bar{X}_n$ (give the name and parameters), as well as the approximate distribution of $\bar{X}_n$ for n large (give the name and parameters).

4.11 Markov Chains

Markov chains were first introduced in 1906 by Andrey Markov (of Markov's inequality), with the goal of showing that the law of large numbers can apply to random variables that are not independent. To see where the Markov model comes from, start by considering an i.i.d.

sequence of random variables $X_0, X_1, \dots, X_n, \dots$ where we think of n as time. This is the setting we worked in throughout Chapter 10, but for modeling real-world phenomena, independence can be an excessively restrictive assumption; it means that the X_n provide absolutely no information about each other. At the other extreme, allowing arbitrary interactions between the X_n makes it very difficult to compute even basic things. A Markov chain is a sequence of r.v.s that exhibits one-step dependence, in a precise sense that we shall soon define. Thus Markov chains are a happy medium between complete independence and complete dependence.

Since their invention, Markov chains have become extremely important in a huge number of fields such as biology, game theory, finance, machine learning, and statistical physics. They are also very widely used for simulations of complex distributions, via algorithms known as Markov chain Monte Carlo (MCMC). In this chapter we will introduce Markov chains and their properties, and in the next chapter we'll look at some examples of MCMC techniques.

4.11.1 11.1 Markov Property and Transition Matrix

Markov chains "live" in both space and time: the set of possible values of the X_n is called the state space, and the index n represents the evolution of the process over time. The state space of a Markov chain can be either discrete or continuous, and time can also be either discrete or continuous (in the continuous-time setting, we would imagine a process X_t defined for all real $t \geq 0$). In this chapter we will focus exclusively on discrete-state, discrete-time Markov chains, with a finite state space. Specifically, we will assume that the X_n take values in a finite set, which we usually take to be $\{1, 2, \dots, M\}$ or $\{0, 1, \dots, M\}$.

Definition 11.1.1 (Markov chain). A sequence of random variables $X_0, X_1, X_2, \dots$ taking values in the state space $\{1, 2, \dots, M\}$ is called a Markov chain if for all $n \geq 0$,

$$P(X_{n+1} = j \mid X_n = i, X_{n-1} = i_{n-1}, \dots, X_0 = i_0) = P(X_{n+1} = j \mid X_n = i).$$

The quantity $P(X_{n+1} = j \mid X_n = i)$ is called the transition probability from state i to state j . In this book, when referring to a Markov chain we will implicitly assume that it is time-homogeneous, which means that the transition probability $P(X_{n+1} = j \mid X_n = i)$ is the same for all times n . But care is needed, since the literature is not consistent about whether to say "time-homogeneous Markov chain" or just "Markov chain".

The above condition is called the Markov property, and it says that given the entire past history $X_0, X_1, X_2, \dots, X_n$, only the most recent term, X_n , matters for predicting X_{n+1} . If we think of time n as the present, times before n as the past, and times after n as the future, the Markov property says that given the present, the past and future are conditionally independent. The Markov property greatly simplifies computations of conditional probability: instead of having to condition on the entire past, we only need to condition on the most recent value.

To describe the dynamics of a Markov chain, we need to know the probabilities of moving from any state to any other state, that is, the probabilities $P(X_{n+1} = j \mid X_n = i)$ on the right-hand side of the Markov property. This information can be encoded in a matrix, called the transition matrix, whose (i, j) entry is the probability of going from state i to state j in one step of the chain.

Definition 11.1.2 (Transition matrix). Let $X_0, X_1, X_2, \dots$ be a Markov chain with state space $\{1, 2, \dots, M\}$, and let $q_{ij} = P(X_{n+1} = j \mid X_n = i)$ be the transition probability from state i to state j . The $M \times M$ matrix $Q = (q_{ij})$ is called the transition matrix of the chain.

Note that Q is a nonnegative matrix in which each row sums to 1. This is because, starting from any state i , the events "move to 1", "move to 2", ..., "move to M " are disjoint, and their union has probability 1 because the chain has to go somewhere.

Example 11.1.3 (Rainy-sunny Markov chain). Suppose that on any given day, the weather can either be rainy or sunny. If today is rainy, then tomorrow will be rainy with probability $1/3$

and sunny with probability $2/3$. If today is sunny, then tomorrow will be rainy with probability $1/2$ and sunny with probability $1/2$. Letting X_n be the weather on day n , $X_0, X_1, X_2, \dots$ is a Markov chain on the state space $\{R, S\}$, where R stands for rainy and S for sunny. We know that the Markov property is satisfied because, from the description of the process, only today's weather matters for predicting tomorrow's.

The transition matrix of the chain is


$$\begin{array}{c} R \quad R \\ R \quad \left(\begin{array}{cc} 1/3 & 2/3 \\ 1/2 & 1/2 \end{array} \right) \\ S \end{array}$$

The first row says that starting from state R , we transition back to state R with probability $1/3$ and transition to state S with probability $2/3$. The second row says that starting from state S , we have a $1/2$ chance of moving to state R and a $1/2$ chance of staying in state S . We could just as well have used

$$\begin{array}{c} S \quad S \\ S \quad \left(\begin{array}{cc} 1/2 & 1/2 \\ 2/3 & 1/3 \end{array} \right) \\ R \end{array}$$

as our transition matrix instead. In general, if there isn't an obvious ordering of the states of a Markov chain (as with the states R and S), we just need to fix an ordering of the states and use it consistently.

The transition probabilities of a Markov chain can also be represented with a diagram. Each state is represented by a circle, and the arrows indicate the possible one-step transitions; we can imagine a particle wandering around from state to state, randomly choosing which arrow to follow. Next to the arrows we write the corresponding transition probabilities.




What if the weather tomorrow depended on the weather today and the weather yesterday? For example, suppose that the weather behaves as above, except that if there have been two consecutive days of rain, then tomorrow will definitely be sunny, and if there have been two consecutive days of sun, then tomorrow will definitely be rainy. Under these new weather dynamics, the X_n no longer form a Markov chain, as the Markov property is violated: conditional on today's weather, yesterday's weather can still provide useful information for predicting tomorrow's weather.

However, by enlarging the state space, we can create a new Markov chain: let $Y_n = (X_{n-1}, X_n)$ for $n \geq 1$. Then $Y_1, Y_2, \dots$ is a Markov chain on the state space $\{(R, R), (R, S), (S, R), (S, S)\}$. You can verify that the new transition matrix is

$$\begin{array}{c} (R, R) \\ (R, R) \quad \left(\begin{array}{cc|cc} 0 & 1 & (S, R) & (S, S) \\ 0 & 0 & 1/2 & 0 \\ (R, S) & 1/3 & 2/3 & 0 & 0 \\ (S, R) & 0 & 0 & 1 & 0 \\ (S, S) \end{array} \right) \end{array}$$

and that its corresponding graphical representation is given in the following figure.



Similarly, we could handle third-order or fourth-order dependencies in the weather by further enlarging the state space to make the Markov property hold. $\square$

Once we have the transition matrix Q of a Markov chain, we can work out the transition probabilities for longer timescales.

Definition 11.1.4 (n -step transition probability). The n -step transition probability from i to j is the probability of being at j exactly n steps after being at i . We denote this by $q_{ij}^{(n)}$:

$$q_{ij}^{(n)} = P(X_n = j \mid X_0 = i).$$

Note that

$$q_{ij}^{(2)} = \sum_k q_{ik} q_{kj}$$

since to get from i to j in two steps, the chain must go from i to some intermediary state k , and then from k to j ; these transitions are independent because of the Markov property. Since the right-hand side is the (i, j) entry of Q^2 by definition of matrix multiplication, we conclude that the matrix Q^2 gives the two-step transition probabilities. By induction, the n th power of the transition matrix gives the n -step transition probabilities:

$$q_{ij}^{(n)} \text{ is the } (i, j) \text{ entry of } Q^n.$$

Example 11.1.5 (Transition matrix of 4-state Markov chain). Consider the 4-state Markov chain depicted in Figure 11.1. When no probabilities are written over the arrows, as in this case, it means all arrows originating from a given state are equally likely. For example, there are 3 arrows originating from state 1, so the transitions $1 \rightarrow 3$, $1 \rightarrow 2$, and $1 \rightarrow 1$ all have probability $1/3$. Therefore the transition matrix of the chain is

$$Q = \begin{pmatrix} 1/3 & 1/3 & 1/3 & 0 \\ 0 & 0 & 1/2 & 1/2 \\ 0 & 1 & 0 & 0 \\ 1/2 & 0 & 0 & 1/2 \end{pmatrix}.$$




FIGURE 11.1
A 4-state Markov chain.

To compute the probability that the chain is in state 3 after 5 steps, starting at state 1, we would look at the $(1, 3)$ entry of Q^5 . Here, using a computer to find Q^5 ,

$$Q^5 = \begin{pmatrix} 853/3888 & 509/1944 & 52/243 & 395/1296 \\ 173/864 & 85/432 & 31/108 & 91/288 \\ 37/144 & 29/72 & 1/9 & 11/48 \\ 499/2592 & 395/1296 & 71/324 & 245/864 \end{pmatrix},$$

so $q_{13}^{(5)} = 52/243$. $\square$

Using the language of Chapter 7, the transition matrix Q encodes the conditional distribution of X_1 given the initial state of the chain. Specifically, the i th row of Q is the conditional PMF of X_1 given $X_0 = i$, displayed as a row vector. Similarly, the i th row of Q^n is the conditional PMF of X_n given $X_0 = i$.

To get the marginal distributions of $X_0, X_1, \dots$, we need to specify not only the transition matrix, but also the initial conditions of the chain. The initial state X_0 can be specified deterministically, or randomly according to some distribution. Let $(t_1, t_2, \dots, t_M)$ be the PMF of X_0 displayed as a vector, that is, $t_i = P(X_0 = i)$. Then the marginal distribution of the chain at any time can be computed from the transition matrix, averaging over all the states using LOTP.

Proposition 11.1.6 (Marginal distribution of X_n). Define $\mathbf{t} = (t_1, t_2, \dots, t_M)$ by $t_i = P(X_0 = i)$, and view $\mathbf{t}$ as a row vector. Then the marginal distribution of X_n is given by the vector $\mathbf{t}Q^n$. That is, the j th component of $\mathbf{t}Q^n$ is $P(X_n = j)$.

Proof. By the law of total probability, conditioning on X_0 , the probability that the chain is in state j after n steps is

$$P(X_n = j) = \sum_{i=1}^M P(X_0 = i) P(X_n = j | X_0 = i) = \sum_{i=1}^M t_i q_{ij}^{(n)},$$

which is the j th component of $\mathbf{t}Q^n$ by definition of matrix multiplication. $\square$

Example 11.1.7 (Marginal distributions of 4 -state Markov chain). Again consider the 4 -state Markov chain shown in Figure 11.1. Suppose that the initial conditions are $\mathbf{t} = (1/4, 1/4, 1/4, 1/4)$, meaning that the chain has equal probability of starting in each of the four states. Let X_n be the position of the chain at time n . Then the marginal distribution of X_1 is

$$\begin{aligned} \mathbf{t}Q &= \begin{pmatrix} 1/4 & 1/4 & 1/4 & 1/4 \end{pmatrix} \begin{pmatrix} 1/3 & 1/3 & 1/3 & 0 \\ 0 & 0 & 1/2 & 1/2 \\ 0 & 1 & 0 & 0 \\ 1/2 & 0 & 0 & 1/2 \end{pmatrix} \\ &= \begin{pmatrix} 5/24 & 1/3 & 5/24 & 1/4 \end{pmatrix} \end{aligned}$$

The marginal distribution of X_5 is

$$\begin{aligned} \mathbf{t}Q^5 &= \begin{pmatrix} 1/4 & 1/4 & 1/4 & 1/4 \end{pmatrix} \begin{pmatrix} 853/3888 & 509/1944 & 52/243 & 395/1296 \\ 173/864 & 85/432 & 31/108 & 91/288 \\ 37/144 & 29/72 & 1/9 & 11/48 \\ 499/2592 & 395/1296 & 71/324 & 245/864 \end{pmatrix} \\ &= \begin{pmatrix} 3379/15552 & 2267/7776 & 101/486 & 1469/5184 \end{pmatrix} \end{aligned}$$

We used a computer to perform the matrix multiplication.

4.11.2 11.2 Classification of States

In this section we introduce terminology for describing the various characteristics of a Markov chain. The states of a Markov chain can be classified as recurrent or transient, depending on whether they are visited over and over again in the long run or are eventually abandoned. States can also be classified according to their period, which is a positive integer summarizing the amount of time that can elapse between successive visits to a state. These characteristics are important because they determine the long-run behavior of the Markov chain, which we will study in Section 11.3.

The concepts of recurrence and transience are best illustrated with a concrete example. In the Markov chain shown on the left of Figure 11.2 (previously featured in Example 11.1.5), a particle moving around between states will continue to spend time in all 4 states in the long run, since it is possible to get from any state to any other state. In contrast, consider the chain on the right of Figure 11.2, and let the particle start at state 1. For a while, the chain may linger in the triangle formed by states 1, 2, and 3, but eventually it will reach state 4, and from there it can never return to states 1, 2, or 3. It will then wander around between states 4, 5, and 6 forever. States 1, 2, and 3 are transient and states 4, 5, and 6 are recurrent.

In general, these concepts are defined as follows.




FIGURE 11.2

Left: 4-state Markov chain with all states recurrent. Right: 6-state Markov chain with states 1, 2, and 3 transient.

Definition 11.2.1 (Recurrent and transient states). State i of a Markov chain is recurrent if starting from i , the probability is 1 that the chain will eventually return to i . Otherwise, the state is transient, which means that if the chain starts from i , there is a positive probability of never returning to i .

In fact, although the definition of a transient state only requires that there be a positive probability of never returning to the state, we can say something stronger: as long as there is a positive probability of leaving i forever, the chain eventually will leave i forever. Moreover, we can find the distribution of the number of returns to the state.

Proposition 11.2.2 (Number of returns to transient state is Geometric). Let i be a transient state of a Markov chain. Suppose the probability of never returning to i , starting from i , is a positive number $p > 0$. Then, starting from i , the number of times that the chain returns to i before leaving forever is distributed $\text{Geom}(p)$.

The proof is by the story of the Geometric distribution: each time that the chain is at i , we have a Bernoulli trial which results in "failure" if the chain eventually returns to i and "success" if the chain leaves i forever; these trials are independent by the Markov property. The number of returns to state i is the number of failures before the first success, which matches the story of the Geometric distribution. And since a Geometric random variable always takes finite values, this proposition tells us that after some finite number of visits, the chain will leave state i forever.

If the number of states is not too large, one way to classify states as recurrent or transient is to draw a diagram of the Markov chain and use the same kind of reasoning that we used when analyzing the chains in Figure 11.2. A special case where we can immediately conclude all states are recurrent is when the chain is irreducible, meaning that it is possible to get from any state to any other state.


Definition 11.2.3 (Irreducible and reducible chain). A Markov chain with transition matrix Q is irreducible if for any two states i and j , it is possible to go from i to j in a finite number of steps (with positive probability). That is, for any states i, j there is some positive integer n such that the (i, j) entry of Q^n is positive. A Markov chain that is not irreducible is called reducible.

Proposition 11.2.4 (Irreducible implies all states recurrent). In an irreducible Markov chain with a finite state space, all states are recurrent.

Proof. It is clear that at least one state must be recurrent; if all states were transient, the chain would eventually leave all states forever and have nowhere to go! So assume without loss of generality that state 1 is recurrent, and consider any other state i . We know that $q_{1i}^{(n)}$ is positive for some n , by definition of irreducibility. Thus, every time the chain is at state 1, there is a positive probability that after n more steps it will be at state i .

Since the chain visits state 1 infinitely often, we know the chain will eventually reach state i from state 1 ; think of each visit to state 1 as starting a trial, where "success" is defined as reaching state i in at most n steps. From state i , the chain will return to state 1 because state 1 is recurrent, and by the same logic, it will eventually reach state i again. By induction, the chain will visit state i infinitely often. Since i was arbitrary, we conclude that all states are recurrent.

The converse of the proposition is false; it is possible to have a reducible Markov chain whose states are all recurrent. An example is given by the Markov chain below, which consists of two "islands" of states.




11.2.5. Note that recurrence or transience is a property of each state in a Markov chain, while irreducibility or reducibility is a property of the chain as a whole.

Here are two familiar problems from earlier chapters, viewed through the lens of Markov chains. For each, we'll identify the recurrent and transient states.


Example 11.2.6 (Gambler's ruin as a Markov chain). In the gambler's ruin problem (Example 2.7.3), two gamblers, A and B, start with i and $N - i$ dollars respectively, making a sequence of bets for \$1. In each round, player A has probability p of winning and probability $q = 1 - p$ of losing. Let X_n be the wealth of gambler A at time n . Then $X_0, X_1, \dots$ is a Markov chain on the state space $\{0, 1, \dots, N\}$. By

design, $X_0 = i$. Once the Markov chain reaches 0 or N , signifying bankruptcy for player A or player B, the chain stays in that state forever. A diagram of the chain is shown below.



We proved in Chapter 2 that the probability that either A or B goes bankrupt is 1, so for any starting state i other than 0 or N , the Markov chain will eventually be absorbed into state 0 or N , never returning to i . Therefore, for this Markov chain, states 0 and N are recurrent, and all other states are transient. The chain is reducible because from state 0 it is only possible to go to state 0, and from state N it is only possible to go to state N . $\square$

Example 11.2.7 (Coupon collector as a Markov chain). In the coupon collector problem (Example 4.3.12), there are C types of coupons (or toys), which we collect one by one, sampling with replacement from the C coupon types each time. Let X_n be the number of distinct coupon types in our collection after n attempts. Then $X_0, X_1, \dots$ is a Markov chain on the state space $\{0, 1, \dots, C\}$. By design, $X_0 = 0$. This chain is depicted below.



With the exception of state C , we can never return to a state after leaving it; the number of coupon types in the collection can only increase with time. Thus all states are transient except for C , which is recurrent. The chain is reducible since, e.g., it's impossible to go from state 2 back to state 1. $\square$

Another way to classify states is according to their periods. The period of a state summarizes how much time can elapse between successive visits to the state.

Definition 11.2.8 (Period of a state, periodic and aperiodic chain). The period of a state

i in a Markov chain is the greatest common divisor (gcd) of the possible numbers of steps it can take to return to i when starting at i . That is, the period of i is the greatest common divisor of numbers n such that the (i, i) entry of Q^n is positive. (The period of i is ∞ if it's impossible ever to return to i after starting at i .) A state is called aperiodic if its period equals 1, and periodic otherwise. The chain itself is called aperiodic if all its states are aperiodic, and periodic otherwise.

For example, let's consider again the two Markov chains from Figure 11.2, shown again in Figure 11.3. We first consider the 6-state chain on the right. Starting from state 1, it is possible to be back at state 1 after 3 steps, 6 steps, 9 steps, etc., but it




FIGURE 11.3

Left: an aperiodic Markov chain. Right: a periodic Markov chain in which states 1, 2, and 3 have period 3.

is not possible to be back at state 1 after any number of steps that is not a multiple of 3. Therefore, state 1 has period 3. Similarly, states 2 and 3 also have period 3. On the other hand, states 4, 5, and 6 have period 1, but the chain is periodic since at least one state does not have period 1. By contrast, in the chain on the left all states are aperiodic, so that chain is aperiodic.

In the gambler's ruin chain from Example 11.2.6, each state has period 2 except for 0 and N , which have period 1. In the coupon collector chain, each state has period 1 except for state 0, which has period ∞ because it's impossible to return to state 0. So neither of these chains is aperiodic.

Checking whether an irreducible chain is aperiodic is often much easier than it might seem: the next proposition shows that we only need to calculate the period of one state, rather than searching state by state for a state whose period is not 1.

Proposition 11.2.9 (Periods in an irreducible chain). In an irreducible Markov chain, all states have the same period.

4.11.3 11.3 Stationary Distribution

The concepts of recurrence and transience are important for understanding the longrun behavior of a Markov chain. At first, the chain may spend time in transient states. Eventually though, the chain will spend all its time in recurrent states. But what fraction of the time will it spend in each of the recurrent states? This question is answered by the stationary distribution of the chain, also known as the steadystate distribution. We will learn in this section that for irreducible and aperiodic

Markov chains, the stationary distribution describes the long-run behavior of the chain, regardless of its initial conditions. It will tell us both the long-run probability of being in any particular state, and the long-run proportion of time that the chain spends in that state.

Definition 11.3.1 (Stationary distribution). A row vector $\mathbf{s} = (s_1, \dots, s_M)$ such that $s_i \geq 0$ and $\sum_i s_i = 1$ is a stationary distribution for a Markov chain with transition matrix Q if

$$\sum_i s_i q_{ij} = s_j$$

for all j . This system of linear equations can be written as one matrix equation:

$$\mathbf{s}Q = \mathbf{s}$$

Recall that if $\mathbf{s}$ is the distribution of X_0 , then $\mathbf{s}Q$ is the marginal distribution of X_1 . Thus the equation $\mathbf{s}Q = \mathbf{s}$ means that if X_0 has distribution $\mathbf{s}$, then X_1 also has distribution $\mathbf{s}$. But then X_2 also has distribution $\mathbf{s}$, as does X_3 , etc. That is, a Markov chain whose initial distribution is the stationary distribution $\mathbf{s}$ will stay in the stationary distribution forever.

One way to think about the stationary distribution of a Markov chain intuitively is to imagine a large number of particles, each independently bouncing from state to state according to the transition probabilities. After a while, the system of particles will approach an equilibrium where, at each time period, the number of particles leaving a state will be counterbalanced by the number of particles entering that state, and this will be true for all states. At this equilibrium, the system as a whole will appear to be stationary, and the proportion of particles in each state will be given by the stationary distribution. We will explore this perspective on stationary distributions more after Definition 11.4.1.

11.3.2 (Stationary distribution is marginal, not conditional). When a Markov chain is at the stationary distribution, the unconditional PMF of X_n equals $\mathbf{s}$ for all n , but the conditional PMF of X_n given $X_{n-1} = i$ is still encoded by the i th row of the transition matrix Q .

If a Markov chain starts at the stationary distribution, then all of the X_n are identically distributed (since they have the same marginal distribution $\mathbf{s}$), but they are not necessarily independent, since the conditional distribution of X_n given $X_{n-1} = i$ is, in general, different from the marginal distribution of X_n .

11.3.3 (Sympathetic magic). If a Markov chain starts at the stationary distribution, then the marginal distributions of the X_n are all equal. This is not the same as saying that the X_n themselves are all equal; confusing the random variables X_n with their distributions is an example of sympathetic magic.

For very small Markov chains, we may solve for the stationary distribution by hand, using the definition. The next example illustrates this for a two-state chain.

Example 11.3.4 (Stationary distribution for a two-state chain). Let

$$Q = \begin{pmatrix} 1/3 & 2/3 \\ 1/2 & 1/2 \end{pmatrix}$$

The stationary distribution is of the form $\mathbf{s} = (s, 1 - s)$, and we must solve for s in

$$\begin{pmatrix} s & 1 - s \end{pmatrix} \begin{pmatrix} 1/3 & 2/3 \\ 1/2 & 1/2 \end{pmatrix} = \begin{pmatrix} s & 1 - s \end{pmatrix},$$

which is equivalent to

$$\begin{aligned} \frac{1}{3}s + \frac{1}{2}(1 - s) &= s \\ \frac{2}{3}s + \frac{1}{2}(1 - s) &= 1 - s \end{aligned}$$

The only solution to these equations is $s = 3/7$, so $(3/7, 4/7)$ is the unique stationary distribution of this Markov chain.

More generally, suppose that $q_{12} = a$ and $q_{21} = b$, where $0 < a < 1$ and $0 < b < 1$. Then the transition matrix is

$$Q = \begin{pmatrix} 1-a & a \\ b & 1-b \end{pmatrix}$$

Writing $\mathbf{s} = (s_1, s_2)$, the equation $\mathbf{s}Q = \mathbf{s}$ becomes the linear system

$$\begin{aligned} (1-a)s_1 + bs_2 &= s_1 \\ as_1 + (1-b)s_2 &= s_2 \end{aligned}$$

Both equations in this system simplify to

$$as_1 = bs_2.$$

Plugging in $s_2 = 1 - s_1$, it follows that the unique solution to this system is

$$\mathbf{s} = \left(\frac{b}{a+b}, \frac{a}{a+b} \right).$$

In short, $\mathbf{s} \propto (b, a)$, i.e., $\mathbf{s}$ is a constant times (b, a) . The constant is whatever it needs to be to make the components of $\mathbf{s}$ sum to 1. As a check, note that for the Q given earlier with specific numbers, this result says that the stationary distribution should be proportional to $(b, a) = (1/2, 2/3)$, which, multiplying by 6 to clear the denominators, is equivalent to being proportional to $(3, 4)$. The stationary distribution we solved for earlier was $(3/7, 4/7)$, which is indeed proportional to $(3, 4)$.

In linear algebra terminology, the equation $\mathbf{s}Q = \mathbf{s}$ says that $\mathbf{s}$ is a left eigenvector of Q with eigenvalue 1 (see Section A. 3 of the math appendix). To get the usual kind of eigenvector (a right eigenvector), take transposes: $Q'\mathbf{s}' = \mathbf{s}'$, where the $'$ symbol denotes taking the transpose.

11.3.1 Existence and Uniqueness

Does a stationary distribution always exist, and is it unique? It turns out that for a finite state space, a stationary distribution always exists. Furthermore, in irreducible Markov chains, the stationary distribution is unique.

Theorem 11.3.5 (Existence and uniqueness of stationary distribution). For any irreducible Markov chain, there exists a unique stationary distribution. In this distribution, every state has positive probability.

The theorem is a consequence of a result from linear algebra called the PerronFrobenius theorem, which is stated in Section A. 3 of the math appendix.

The 4-state chain on the left of Figure 11.3 is irreducible: in terms of the picture, it is possible to go from anywhere to anywhere following the arrows; in terms of the transition matrix Q , all the entries of Q^5 are positive. Therefore, by Theorem 11.3.5, the chain has a unique stationary distribution.

On the other hand, the gambler's ruin chain from Example 11.2.6 is reducible, so the theorem does not apply. It turns out that this chain has infinitely many stationary distributions. In the long run, the chain will either wind up at state 0 (and stay there forever) or wind up at state N (and stay there forever). This suggests, and it is easy to check, that the degenerate distributions $\mathbf{s} = (1, 0, \dots, 0)$ and $\mathbf{t} = (0, 0, \dots, 1)$ are both stationary distributions. It follows that any weighted combination $p\mathbf{s} + (1-p)\mathbf{t}$, where $0 \leq p \leq 1$, is also a stationary distribution, since it sums to 1 and

$$(p\mathbf{s} + (1-p)\mathbf{t})Q = p\mathbf{s}Q + (1-p)\mathbf{t}Q = p\mathbf{s} + (1-p)\mathbf{t},$$

where, as usual, Q is the transition matrix.

11.3.2 Convergence

We have already informally stated that the stationary distribution describes the long-run behavior of the chain, in the sense that if we run the chain for a long time, the marginal distribution of X_n converges to the stationary distribution $\mathbf{s}$. The next theorem states that this is true as long as the chain is both irreducible and aperiodic. Then, regardless of the chain's initial conditions, the PMF of X_n will converge to the stationary distribution as $n \rightarrow \infty$. This relates the concept of stationarity to the long-run behavior of a Markov chain. The proof is omitted.

Theorem 11.3.6 (Convergence to stationary distribution). Let $X_0, X_1, \dots$ be an irreducible, aperiodic Markov chain with stationary distribution $\mathbf{s}$ and transition matrix Q . Then $P(X_n = i)$ converges to s_i as $n \rightarrow \infty$. In terms of the transition matrix, Q^n converges to a matrix in which each row is $\mathbf{s}$.

Therefore, after a large number of steps, the probability that the chain is in state i is close to the stationary probability s_i , regardless of the chain's initial conditions. This makes irreducible, aperiodic chains especially nice to work with. Irreducibility means that for each (i, j) there is some power Q^m where the (i, j) entry is positive, but if we also assume aperiodicity it turns out that we can find a value of m that works for all i, j . More precisely, a chain is irreducible and aperiodic if and only if some power Q^m is positive in all entries.

Intuitively, the extra condition of aperiodicity is needed in order to rule out chains that just go around in circles, such as the chain in the following example, or chains where, say, some states are accessible only after an even number of steps while other states are accessible only after an odd number of steps.

Example 11.3.7 (Periodic chain). The figure below shows a periodic Markov chain where each state has period 5.




FIGURE 11.4

A periodic chain.

The transition matrix of the chain is

$$Q = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

It can be verified without much difficulty that $\mathbf{s} = (1/5, 1/5, 1/5, 1/5, 1/5)$ is a stationary distribution of this chain, and by Theorem 11.3.5, $\mathbf{s}$ is unique. However, suppose the chain

starts at $X_0 = 1$. Then the PMF of X_n assigns probability 1 to the state $(n \bmod 5) + 1$ and 0 to all other states.¹ In particular, it does not converge to $\mathbf{s}$ as $n \rightarrow \infty$. Nor does Q^n converge to a matrix in which each row is $\mathbf{s}$: the chain's transitions are deterministic, so Q^n always consists of 0's and 1's. $\square$

Lastly, the stationary distribution tells us the average time between visits to any particular state.

Theorem 11.3.8 (Expected time to return). Let $X_0, X_1, \dots$ be an irreducible Markov chain with stationary distribution $\mathbf{s}$. Let r_i be the expected time it takes the chain to return to i , given that it starts at i . Then $s_i = 1/r_i$.

Here is how the theorems apply to the two-state chain from Example 11.3.4.

Example 11.3.9 (Long-run behavior of a two-state chain). In the long run, the chain in Example 11.3.4 will spend $3/7$ of its time in state 1 and $4/7$ of its time in state 2. Starting at state 1, it will take an average of $7/3$ steps to return to state 1. The powers of the transition matrix converge to a matrix where each row is the stationary distribution:

$$\begin{pmatrix} 1/3 & 2/3 \\ 1/2 & 1/2 \end{pmatrix}^n \rightarrow \begin{pmatrix} 3/7 & 4/7 \\ 3/7 & 4/7 \end{pmatrix} \text{ as } n \rightarrow \infty.$$

11.3.3 Google Pagerank

We next consider a vastly larger example of a stationary distribution, for a Markov chain on a state space with billions of interconnected nodes: the World Wide Web. The next example explains how the founders of Google modeled web-surfing as a Markov chain, and then used its stationary distribution to rank the relevance of webpages. For years Google described the resulting method, known as PageRank, as "the heart of our software".

Suppose you are interested in a certain topic, say chess, so you use a search engine to look for useful webpages with information about chess. There are millions of webpages that mention the word "chess", so a key issue a search engine needs to deal with is what order to show the search results in. It would be a disaster to have to wade through thousands of garbage pages that mention "chess" before finding informative content.

In the early days of the web, various approaches to this ranking problem were used. For example, some search engines employed people to manually decide which pages were most useful, like a museum curator. But aside from being subjective and expensive, this quickly became infeasible as the web grew. Others focused on the number of times the search term was mentioned on the site. But a page that mentions "chess" over and over again could easily be less useful than a concise reference page or a page about chess that doesn't repeatedly mention the word. Furthermore, this method is very open to abuse: a spam page could boost its ranking just by including a long list of words repeated over and over again.

Both of the above methods ignore the structure of the web: which pages link to which other pages? Taking the link structure into account led to dramatic improvements in search engines. As a first attempt, one could rank a page based on how many other pages link to it. That is, if Page A links to Page B, we consider it a vote for B, and we rank pages based on how many votes they have.

But this is again very open to abuse: a spam page could boost its ranking by creating thousands of other spam pages linking to it. And though it may seem democratic for each page to have equal voting power, an incoming link from a reliable page is more meaningful than a link from an uninformative page. Google PageRank, which was introduced in 1998 by Sergey Brin and the aptly named Larry Page, ranks

⁰¹ The notation $n \bmod 5$ denotes the remainder when n is divided by 5.

the importance of a page not only by how many pages link to it, but also by the importance of those pages.

Consider the web as a directed network-which is what it is. Each page on the web is a node, and links between nodes represent links between pages. For example, suppose for simplicity that the web only has 4 pages, connected as shown in the figure below.




FIGURE 11.5
It's a small web after all.

Imagine someone randomly surfing the web, starting at some page and then randomly clicking links to go from one page to the next (with equal probabilities for all links on the current page). The idea of PageRank is to measure the importance of a page by the long-run fraction of time spent at that page.

Of course, some pages may have no outgoing links at all, such as page 4 above. When the web surfer encounters such a page, rather than despairing they open up a new browser window and visits a uniformly random page. Thus a page with no links is converted to a page that links to every page, including itself. For the example above, the resulting transition matrix is

$$Q = \begin{pmatrix} 0 & 1/2 & 0 & 1/2 \\ 1/2 & 0 & 1/2 & 0 \\ 0 & 0 & 0 & 1 \\ 1/4 & 1/4 & 1/4 & 1/4 \end{pmatrix}.$$

In general, let M be the number of pages on the web, let Q be the M by M transition matrix of the chain described above, and let $\mathbf{s}$ be the stationary distribution (assuming it exists and is unique). Think of s_j as a measure of how important page j is. Intuitively, the equation

$$s_j = \sum_i s_i q_{ij}$$

says that the score of page j should be based not only on how many other pages link to it, but on their scores. Furthermore, the "voting power" of a page gets diluted if it has a lot of outgoing links: it counts for more if page i 's only link is to page j (so that $q_{ij} = 1$) than if page i has thousands of links, one of which happens to be to page j .

It is not clear that a unique stationary distribution exists for this chain, since it may not be irreducible and aperiodic. Even if it is irreducible and aperiodic, convergence to the stationary

distribution could be very slow since the web is so immense. To address these issues, suppose that before each move, the web surfer flips a coin with probability α of Heads. If Heads, the web surfer clicks a random link from the current page; if Tails, the web surfer teleports to a uniformly random page. The resulting chain has the Google transition matrix

$$G = \alpha Q + (1 - \alpha) \frac{J}{M},$$

where J is the M by M matrix of all 1's. Note that the row sums of G are 1 and that all entries are positive, so G is a valid transition matrix for an irreducible, aperiodic Markov chain. This means there is a unique stationary distribution $\mathbf{s}$, called PageRank, and the chain will converge to it! The choice of α is an important consideration; choosing α close to 1 makes sense to respect the structure of the web as much as possible, but there is a tradeoff since it turns out that smaller values of α make the chain converge much faster. As a compromise, the original recommendation of Brin and Page was $\alpha = 0.85$.

PageRank is conceptually nice, but computing it sounds extremely difficult, considering that $\mathbf{s}G = \mathbf{s}$ could be a system of 100 billion equations in 100 billion unknowns. Instead of thinking of this as a massive algebra problem, we can use the Markov chain interpretation: for any starting distribution $\mathbf{t}$, $\mathbf{t}G^n \rightarrow \mathbf{s}$ as $n \rightarrow \infty$. And $\mathbf{t}G$ is easier to compute than it might seem at first:

$$\mathbf{t}G = \alpha(\mathbf{t}Q) + \frac{1 - \alpha}{M}(\mathbf{t}J),$$

where computing the first term isn't too hard since Q is very sparse (mostly 0's) and computing the second term is easy since $\mathbf{t}J$ is a vector of all 1's. Then $\mathbf{t}G$ becomes the new $\mathbf{t}$, and we can compute $\mathbf{t}G^2 = (\mathbf{t}G)G$, etc., until the sequence appears to have converged (though it is hard to know that it has converged). This gives an approximation to PageRank, and has an intuitive interpretation as the distribution of where the web surfer is after a large number of steps.

4.11.4 11.4 Reversibility

We have seen that the stationary distribution of a Markov chain is extremely useful for understanding its long-run behavior. Unfortunately, in general it may be computationally difficult to find the stationary distribution when the state space is large. This section addresses an important special case where working with eigenvalue equations for large matrices can be avoided.

Definition 11.4.1 (Reversibility). Let $Q = (q_{ij})$ be the transition matrix of a Markov chain. Suppose there is $\mathbf{s} = (s_1, \dots, s_M)$ with $s_i \geq 0$, $\sum_i s_i = 1$, such that

$$s_i q_{ij} = s_j q_{ji}$$

for all states i and j . This equation is called the reversibility or detailed balance condition, and we say that the chain is reversible with respect to $\mathbf{s}$ if it holds.

The term "reversible" comes from the fact that a reversible chain, started according to its stationary distribution, behaves in the same way regardless of whether time is run forwards or backwards. If you record a video of a reversible chain, started according to its stationary distribution, and then show the video to a friend, either in the normal way or with time reversed, your friend will not be able to determine from watching the video whether time is running forwards or backwards.

As discussed after Definition 11.3.1, we can think about the stationary distribution of a Markov chain intuitively in terms of a system consisting of a large number of particles independently bouncing around according to the transition probabilities. In the long run, the

proportion of particles in any state j is the stationary probability of state j , and the flow of particles out of state j is counterbalanced by the flow of particles into state j . To see this in more detail, let n be the number of particles and $\mathbf{s}$ be a probability vector such that s_j is the current proportion of particles at state j . By definition, $\mathbf{s}$ is the stationary distribution of the chain if and only if

$$s_j = \sum_i s_i q_{ij} = s_j q_{jj} + \sum_{i:i \neq j} s_i q_{ij}$$

for all states j . This equation can be rewritten as

$$ns_j(1 - q_{jj}) = \sum_{i:i \neq j} ns_i q_{ij}$$

The left-hand side is the approximate number of particles that will exit from state j on the next step, since there are ns_j particles at state j , each of which will stay at j with probability q_{jj} and leave with probability $1 - q_{jj}$. The right-hand side is the approximate number of particles that will enter state j on the next step, since for each $i \neq j$ there are ns_i particles at state i , each of which will enter state j with probability q_{ij} . So there is a balance between particles leaving state j and particles entering state j .

The reversibility condition imposes a much more stringent form of balance, in which for each pair of states i, j with $i \neq j$, the flow of particles from state i to state j is counterbalanced by the flow of particles from state j to state i . To see this, write the reversibility equation for states i and j as

$$ns_i q_{ij} = ns_j q_{ji}$$

The left-hand side is the approximate number of particles that will go from state i to state j on the next step, since there are ns_i particles at state i , each of which has probability q_{ij} of going to state j . Similarly, the right-hand side is the approximate number of particles that will go from state j to state i on the next step. So reversibility says that there is a balance between particles going from state i to state j , and particles going from state j to state i .

Given a transition matrix, if we can find a probability vector $\mathbf{s}$ that satisfies the reversibility condition, then $\mathbf{s}$ is automatically a stationary distribution. This should not be surprising, in light of the above discussion.

Proposition 11.4.2 (Reversible implies stationary). Suppose that $Q = (q_{ij})$ is the transition matrix of a Markov chain that is reversible with respect to a nonnegative vector $\mathbf{s} = (s_1, \dots, s_M)$ whose components sum to 1. Then $\mathbf{s}$ is a stationary distribution of the chain.

Proof. We have

$$\sum_i s_i q_{ij} = \sum_i s_j q_{ji} = s_j \sum_i q_{ji} = s_j,$$

where the last equality is because each row sum of Q is 1. So $\mathbf{s}$ is stationary.

This is a powerful result because it is often easier to verify the reversibility condition than it is to solve the entire system of equations $\mathbf{s}Q = \mathbf{s}$. However, in general we may not know in advance whether it is possible to find $\mathbf{s}$ satisfying the reversibility condition, and even when it is possible, it may take a lot of effort to find an $\mathbf{s}$ that works. In the remainder of this section, we look at three types of Markov chains where it is possible to find an $\mathbf{s}$ that satisfies the reversibility condition. Such Markov chains are called reversible.

First, if Q is a symmetric matrix, then the stationary distribution is uniform over the state space: $\mathbf{s} = (1/M, 1/M, \dots, 1/M)$. It is easy to see that if $q_{ij} = q_{ji}$, then the reversibility condition $s_i q_{ij} = s_j q_{ji}$ is satisfied when $s_i = s_j$ for all i and j .


This is a special case of a more general fact, stated below: if each column of Q sums to 1, then the stationary distribution is uniform over the state space.

Proposition 11.4.3. If each column of the transition matrix Q sums to 1, then the uniform distribution over all states, $(1/M, 1/M, \dots, 1/M)$, is a stationary distribution. (A nonnegative matrix such that the row sums and the column sums are all equal to 1 is called a doubly stochastic matrix.)

Proof. Assuming each column sums to 1, the row vector $\mathbf{v} = (1, 1, \dots, 1)$ satisfies $\mathbf{v}Q = \mathbf{v}$. It follows that $(1/M, 1/M, \dots, 1/M)$ is stationary.

Second, if the Markov chain is a random walk on an undirected network, then there is a simple formula for the stationary distribution.

Example 11.4.4 (Random walk on an undirected network). A network is a collection of nodes joined by edges; the network is undirected if edges can be traversed in either direction, meaning there are no one-way streets. Suppose a wanderer randomly traverses the edges of an undirected network. From a node i , the wanderer randomly picks any of the edges at i , with equal probabilities, and then traverses the chosen edge. For example, in the network shown below, from node 3 the wanderer goes to node 1 or node 2, with probability $1/2$ each.



The degree of a node is the number of edges attached to it, and the degree sequence of a network with nodes $1, 2, \dots, n$ is the vector $(d_1, \dots, d_n)$ listing all the degrees, where d_j is the degree of node j . An edge from a node to itself is allowed (such an edge is called a self-loop), and counts 1 toward the degree of that node.

For example, the network above has degree sequence $\mathbf{d} = (4, 3, 2, 3, 2)$. Note that

$$d_i q_{ij} = d_j q_{ji}$$

for all i, j , since q_{ij} is $1/d_i$ if $\{i, j\}$ is an edge and 0 otherwise, for $i \neq j$. Therefore, by Proposition 11.4.2, the stationary distribution is proportional to the degree sequence. Intuitively, the nodes with the highest degrees are the most well-connected, so it makes sense that

the chain spends the most time in these states in the long run. In the example above, this says that


$$s = \left(\frac{4}{14}, \frac{3}{14}, \frac{2}{14}, \frac{3}{14}, \frac{2}{14} \right)$$

is the stationary distribution for the random walk.

Exercise 20 explores random walk on a weighted undirected network; each edge has a weight assigned to it, and the wanderer chooses where to go from i with probabilities proportional to the weights on the available edges. It turns out that this is a reversible Markov chain. More surprisingly, every reversible Markov chain can be represented as random walk on a weighted undirected network! $\square$

Here is a concrete example of a random walk on an undirected network.

Example 11.4.5 (Knight on a chessboard). Consider a knight randomly moving around on a 4×4 chessboard.



The 16 squares are labeled in a grid, e.g., the knight is currently at the square B3, and the upper left square is A4. Each move of the knight is an L-shaped jump: the knight moves two squares horizontally followed by one square vertically, or vice versa. For example, from B3 the knight can move to A1, C1, D2, or D4; from A4 it can move to B2 or C3. Note that from a light square, the knight always moves to a dark square and vice versa.

Suppose that at each step, the knight moves randomly, with each possibility equally likely. This creates a Markov chain where the states are the 16 squares. Compute the stationary distribution of the chain.

Solution:

There are only three types of squares on the board: 4 center squares, 4 corner squares (such as A4), and 8 edge squares (such as B4; exclude corner squares from being considered edge squares). We can consider the board to be an undirected network where two squares are connected by an edge if they are accessible via a single knight's move. Then a center square has degree 4, a corner square has degree 2, and an edge square has degree 3, so their stationary probabilities are $4a$, $2a$, $3a$ respectively for some a .


To find a , count the number of squares of each type to get $4a \cdot 4 + 2a \cdot 4 + 3a \cdot 8 = 1$, giving $a = 1/48$. Thus, each center square has stationary probability $4/48 = 1/12$, each corner

square has stationary probability $2/48 = 1/24$, and each edge square has stationary probability $3/48 = 1/16$. $\square$

Third and finally, if in each time period a Markov chain can only move one step to the left, one step to the right, or stay in place, then it is called a birth-death chain. All birth-death chains are reversible.

Example 11.4.6 (Birth-death chain). A birth-death chain on states $\{1, 2, \dots, M\}$ is a Markov chain with transition matrix $Q = (q_{ij})$ such that $q_{ij} > 0$ if $|i - j| = 1$ and $q_{ij} = 0$ if $|i - j| \geq 2$. This says it's possible to go one step to the left and possible to go one step to the right (except at boundaries) but impossible to jump further in one step. The name stems from applications to the growth or decline of a population, where a step to the right is thought of as a birth and a step to the left is thought of as a death in the population.

For example, the chain shown below is a birth-death chain if the labeled transitions have positive probabilities, except for the loops from a state to itself, which are allowed to have 0 probability.



We will now show that any birth-death chain is reversible, and construct the stationary distribution. Let s_1 be a positive number, to be specified later. Since we want $s_1 q_{12} = s_2 q_{21}$, let

$$s_2 = s_1 q_{12} / q_{21}.$$

Then since we want $s_2 q_{23} = s_3 q_{32}$, let

$$s_3 = s_2 q_{23} / q_{32} = s_1 q_{12} q_{23} / (q_{32} q_{21}).$$


Continuing in this way, let

$$s_j = \frac{s_1 q_{12} q_{23} \dots q_{j-1,j}}{q_{j,j-1} q_{j-1,j-2} \dots q_{21}},$$


for all states j with $2 \leq j \leq M$. Choose s_1 so that the s_j sum to 1. Then the chain is reversible with respect to $\mathbf{s}$, since $q_{ij} = q_{ji} = 0$ if $|i - j| \geq 2$ and by construction $s_i q_{ij} = s_j q_{ji}$ if $|i - j| = 1$. Thus, $\mathbf{s}$ is the stationary distribution. $\square$

The Ehrenfest chain is a birth-death chain that can be used as a simple model for the diffusion of gas molecules. The stationary distribution turns out to be a Binomial distribution.

Example 11.4.7 (Ehrenfest). There are two containers with a total of M distinguishable particles. Transitions are made by choosing a random particle and moving it from its current container into the other container. Initially, all of the particles are in the second container. Let X_n be the number of particles in the first container at time n , so $X_0 = 0$ and the transition from X_n to X_{n+1} is done as described above. This is a periodic Markov chain with state space $\{0, 1, \dots, M\}$.



$$X_n$$



$$M - X_n$$

As illustrated above, the first container has X_n particles and the second has $M - X_n$ particles. We will use the reversibility condition to show that $\mathbf{s} = (s_0, s_1, \dots, s_M)$ with

$$s_i = \binom{M}{i} \left(\frac{1}{2}\right)^M$$

is the stationary distribution. Note that this is the $\text{Bin}(M, 1/2)$ PMF.

Let s_i be as claimed, and check that $s_i q_{ij} = s_j q_{ji}$. If $j = i + 1$ (with $i < M$), then

$$\begin{aligned} s_i q_{ij} &= \binom{M}{i} \left(\frac{1}{2}\right)^M \frac{M-i}{M} = \frac{M!}{(M-i)!i!} \left(\frac{1}{2}\right)^M \frac{M-i}{M} = \binom{M-1}{i} \left(\frac{1}{2}\right)^M, \\ s_j q_{ji} &= \binom{M}{j} \left(\frac{1}{2}\right)^M \frac{j}{M} = \frac{M!}{(M-j)!j!} \left(\frac{1}{2}\right)^M \frac{j}{M} = \binom{M-1}{j-1} \left(\frac{1}{2}\right)^M = s_i q_{ij} \end{aligned}$$

By a similar calculation, if $j = i - 1$ (with $i > 0$), then $s_i q_{ij} = s_j q_{ji}$. For all other values of i and j , $q_{ij} = q_{ji} = 0$. Therefore, $\mathbf{s}$ is stationary.

The Binomial was a natural guess for the stationary distribution because after running the Markov chain for a long time, each particle is about equally likely to be in either container. However, the PMF does not converge to a Binomial since the chain has period 2, with X_n guaranteed to be even when n is even, and odd when n is odd.

Happily, it turns out that another interpretation of the stationary distribution remains valid here: s_i is the long-run proportion of time that the chain spends in state i . More precisely, letting I_k be the indicator of the chain being in state i at time k , it can be shown that

$$\frac{1}{n} \sum_{k=0}^{n-1} I_k \rightarrow s_i$$

as $n \rightarrow \infty$, with probability 1. $\square$

4.11.5 11.5 Recap

A Markov chain is a sequence of r.v.s $X_0, X_1, X_2, \dots$ satisfying the Markov property, which states that given the present, the past and future are conditionally independent:

$$P(X_{n+1} = j \mid X_n = i, X_{n-1} = i_{n-1}, \dots, X_0 = i_0) = P(X_{n+1} = j \mid X_n = i) = q_{ij}.$$

The transition matrix $Q = (q_{ij})$ gives the probabilities of moving from any state to any other state in one step. The i th row of the transition matrix is the conditional PMF of X_{n+1} given $X_n = i$. The n th power of the transition matrix gives the n step transition probabilities. If we specify initial conditions $s_i = P(X_0 = i)$ and let $\mathbf{s} = (s_1, \dots, s_M)$, then the marginal PMF of X_n is $\mathbf{s}Q^n$.

States of a Markov chain can be classified as recurrent or transient: recurrent if the chain will return to the state over and over, and transient if it will eventually leave forever. States can also be classified according to their periods; the period of state i is the greatest common divisor of the numbers of steps it can take to return to i , starting from i . A chain is irreducible if it is possible to get from any state to any state in a finite number of steps, and aperiodic if each state has period 1.

A stationary distribution for a finite Markov chain is a PMF such that $\mathbf{s}Q = \mathbf{s}$. Under various conditions, the stationary distribution of a finite Markov chain exists and is unique, and the PMF of X_n converges to $\mathbf{s}$ as $n \rightarrow \infty$. If state i has stationary probability s_i , then the expected time for the chain to return to i , starting from i , is $r_i = 1/s_i$.

If a PMF $\mathbf{s}$ satisfies the reversibility condition $s_i q_{ij} = s_j q_{ji}$ for all i and j , it guarantees that $\mathbf{s}$ is a stationary distribution of the Markov chain with transition matrix $Q = (q_{ij})$. Markov chains for which there exists $\mathbf{s}$ satisfying the reversibility condition are called reversible. We discussed three types of reversible chains:

1. If the transition matrix is symmetric, then the stationary distribution is uniform over all states.

2. If the chain is a random walk on an undirected network, then the stationary distribution is proportional to the degree sequence, i.e.,

$$s_j = \frac{d_j}{\sum_i d_i}$$

3. If the chain is a birth-death chain, then the stationary distribution satisfies

$$s_j = \frac{s_1 q_{12} q_{23} \dots q_{j-1,j}}{q_{j,j-1} q_{j-1,j-2} \dots q_{21}}$$

for $j > 1$, where s_1 is solved for at the end to make $s_1 + \dots + s_M = 1$.




FIGURE 11.6 Given a transition matrix Q and a distribution $\mathbf{t}$ over the states, we can generate a Markov chain $X_0, X_1, \dots$ by choosing X_0 according to $\mathbf{t}$ and then running the chain according to the transition probabilities. An important event is $X_n = i$, the event that the chain is visiting state i at time n . We can then find the PMF of X_n in terms of Q and $\mathbf{t}$, and (under conditions discussed in the chapter) the PMF will converge to the stationary distribution $\mathbf{s}$. If instead we start the chain out according to $\mathbf{s}$, then the chain will stay stationary forever.

Figure 11.6 compares two ways to run a Markov chain with transition matrix Q : choosing the initial state according to an arbitrary distribution $\mathbf{t}$ over the states, or choosing the initial state according to the stationary distribution $\mathbf{s}$. In the former case, the exact PMF after n steps can be found in terms of Q and $\mathbf{t}$, and the PMF converges to $\mathbf{s}$ (under some very general conditions discussed in this chapter). In the latter case, the chain is stationary forever.

4.11.6 11.6 R

Matrix Calculations

Let's do some calculations for the 4-state Markov chain in Example 11.1.5, as an example of working with transition matrices in R. First, we need to specify the transition matrix Q . This is done with the matrix command: we type in the entries of the matrix, row by row, as a long vector, and then we tell R the number of rows and columns in the matrix (nrow and ncol), as well as the fact that we typed in the entries by row (byrow=TRUE):

```
Q <- matrix(c(1/3,1/3,1/3,0,
              0,0,1/2,1/2,
              0,1,0,0,
              1/2,0,0,1/2),nrow=4,ncol=4,byrow=TRUE)
```

To obtain higher order transition probabilities, we can multiply Q by itself repeatedly. The matrix multiplication command in R is `%*%` (not just `*`). So


```
Q2 <- Q %*% Q
Q3 <- Q2 %*% Q
Q4 <- Q2 %*% Q2
Q5 <- Q3 %*% Q2
```

produces Q^2 through Q^5 . If we want to know the probability of going from state 3 to state 4 in exactly 5 steps, we can extract the (3, 4) entry of Q^5 :

```
Q5 [3,4]
```

This gives 0.229 , agreeing with the value 11/48 shown in Example 11.1.5.

To compute a power Q^n without directly doing repeated matrix multiplications, we can use the command `Q%^n` after installing and loading the `expm` package. For example, `Q%^42` yields Q^{42} . By exploring the behavior of Q^n as n grows, we can see Theorem 11.3.6 in action (and get a sense of how long it takes for the chain to get very close to its stationary distribution).

In particular, for n large each row of Q^n is approximately (0.214, 0.286, 0.214, 0.286), so this is approximately the stationary distribution. Another way to obtain the stationary distribution numerically is to use

```
eigen(t(Q))
```

to compute the eigenvalues and eigenvectors of the transpose of Q ; then the eigenvector corresponding to the eigenvalue 1 can be selected and normalized so that the components sum to 1 .

Gambler's Ruin

To simulate from the gambler's ruin chain from Example 11.2.6, we start by deciding the total amount of money N , the probability p of gambler A winning a given round, and the number of time periods $nsim$ that we wish to simulate.

```
N <- 10
p <- 1/2
nsim <- 80
```

Next, we allocate a vector of length $nsim$ called x , which will store the values of the Markov chain. For the initial condition, we set the first entry of x equal to 5 ; this gives both gamblers \$5 to start with.

```
x <- rep(0,nsim)
x[1] <- 5
```

Now we are ready to simulate the subsequent values of the Markov chain. This is achieved with the following block of code, which we will explain step by step.

```
for (i in 2:nsim){
  if (x[i-1]==0 || x[i-1]==N){
    x[i] <- x[i-1]
  }
  else{
    x[i] <- x[i-1] + sample(c(1,-1), 1, prob=c(p,1-p))
  }
}
```

The first line and the outer set of braces constitute a for loop: `for (i in 2:nsim)` means that all the code inside the for loop will be executed over and over, with the value of `i` set to 2, then set to 3, then set to 4, all the way until `i` reaches the value `nsim`. Each pass through the loop represents one step of the Markov chain.

Inside the for loop, we first check to see whether the chain is already at one of the endpoints, 0 or `N`; we do this with an if statement. If the chain is already at 0 or `N`, then we set its new value equal to its previous value, since the chain is not allowed to escape 0 or `N`. Otherwise, if the chain is not at 0 or `N`, it is free to move left or right. We use the `sample` command to move to the right 1 unit or to the left 1 unit, with probabilities p and $1 - p$, respectively.

To see what path was taken by the Markov chain during our simulation, we can plot the values of `x` as a function of time:

```
plot(x,type='l',ylim=c(0,N))
```

You should see a path that starts at 5 and bounces up and down before being absorbed into state 0 or state `N`.

Simulating From a Finite-state Markov Chain

With a few modifications, we can simulate from an arbitrary Markov chain on a finite state space. For concreteness, we will illustrate how to simulate from the 4-state Markov chain in Example 11.1.5.

As above, we can type

```
Q <- matrix(c(1/3,1/3,1/3,0,
             0,0,1/2,1/2,
             0,1,0,0
             1/2,0,0,1/2),nrow = 4,ncol = 4,byrow = TRUE)
```

to specify the transition matrix Q .

Next, we choose the number of states and the number of time periods to simulate, we allocate space for the results of the simulation, and we choose initial conditions for the chain. In this example, `x[1] <- sample(1 : M, 1)` says the initial distribution of the chain is uniform over all states.

```
M <- nrow(Q)
nsim <- 10^4
x <- rep(0,nsim)
x[1] <- sample(1:M,1)
```

For the simulation itself, we again use `sample` to choose a number from 1 to `M`. At time i , the chain was previously at state $x[i - 1]$, so we must use row $x[i - 1]$ of the transition matrix to determine the probabilities of sampling 1, 2, ..., `M`. The notation $Q[x[i - 1],]$ denotes row $x[i - 1]$ of the matrix Q .

```
for (i in 2:nsim){
  x[i] <- sample(M, 1, prob=Q[x[i-1],])
}
```

Since we set `nsim` to a large number, it may be reasonable to believe that the chain is close to stationarity during the latter portion of the simulation. To check this, we eliminate the first half of the simulations to give the chain time to reach stationarity:

```
x <- x[-(1:(nsim/2))]
```

We then use the `table` command to calculate the number of times the chain visited each state; dividing by `length(x)` converts the counts into proportions. The result is an approximation to the stationary distribution.

```
table(x)/length(x)
```


For comparison, the true stationary distribution of the chain is approximately $(0.214, 0.286, 0.214, 0.286)$. Is this close to what you obtained via simulation?

4.11.7 11.7 Exercises

Exercises marked with (S) have detailed solutions at <http://stat110.net>.

Markov Property

- (S) Let $X_0, X_1, X_2, \dots$ be a Markov chain. Show that $X_0, X_2, X_4, X_6, \dots$ is also a Markov chain, and explain why this makes sense intuitively.
- (S) Let $X_0, X_1, X_2, \dots$ be an irreducible Markov chain with state space $\{1, 2, \dots, M\}$, $M \geq 3$, transition matrix $Q = (q_{ij})$, and stationary distribution $\mathbf{s} = (s_1, \dots, s_M)$. Let the initial state X_0 follow the stationary distribution, i.e., $P(X_0 = i) = s_i$.
 - On average, how many of $X_0, X_1, \dots, X_9$ equal 3? (In terms of $\mathbf{s}$; simplify.)
 - Let $Y_n = (X_n - 1)(X_n - 2)$. For $M = 3$, find an example of Q (the transition matrix for the original chain $X_0, X_1, \dots$) where $Y_0, Y_1, \dots$ is Markov, and another example of Q where $Y_0, Y_1, \dots$ is not Markov. In your examples, make $q_{ii} > 0$ for at least one i and make sure it is possible to get from any state to any other state eventually.
- A Markov chain has two states, A and B , with transitions as follows:



Suppose we do not get to observe this Markov chain, which we'll call $X_0, X_1, X_2, \dots$. Instead, whenever the chain transitions from A back to A , we observe a 0, and whenever it changes states, we observe a 1. Let the sequence of 0's and 1's be called $Y_0, Y_1, Y_2, \dots$. For example, if the X chain starts out as

$A, A, B, A, B, A, A, \dots$

then the Y chain starts out as

$0, 1, 1, 1, 1, 0, \dots$

- Show that $Y_0, Y_1, Y_2, \dots$ is not a Markov chain.
- In Example 11.1.3, we dealt with a violation of the Markov property by enlarging the state space to incorporate second-order dependence. Show that such a trick will not work for $Y_0, Y_1, Y_2, \dots$. That is, no matter how large m is,

$$Z_n = (Y_{n-m+1}, Y_{n-m+2}, \dots, Y_n), \quad n = m-1, m, \dots$$

is still not a Markov chain.


4. There are three blocks, floating in a sea of lava. Label the blocks 1, 2, 3, from left to right. Mark the Kangaroo, a video game character, is standing on block 1. To reach safety, he must get to block 3. He can't jump directly from block 1 to block 3; his only hope is to jump from block 1 to block 2, then jump from block 2 to block 3. Each time he jumps, he has probability $1/2$ of success and probability $1/2$ of "dying" by falling into the lava. If he "dies", he starts again at block 1.

Let J be the total number of jumps that Mark will make in order to get to block 3.


- Find $E(J)$, using the Markov property and conditional expectation.
- Explain how this problem relates to a coin tossing problem from Chapter 9.


Stationary Distribution

5. (S) Consider the Markov chain shown below, where $0 < p < 1$ and the labels on the arrows indicate transition probabilities.



- Write down the transition matrix Q for this chain.
 - Find the stationary distribution of the chain.
 - What is the limit of Q^n as $n \rightarrow \infty$?
6. (S) Consider the Markov chain shown below, where the state space is $\{1, 2, 3, 4\}$ and the labels to the right of arrows indicate transition probabilities.






- (a) Write down the transition matrix Q for this chain.
 - (b) Which states (if any) are recurrent? Which states (if any) are transient?
 - (c) Find two different stationary distributions for the chain.
7. (S) Daenerys has three dragons: Drogon, Rhaegal, and Viserion. Each dragon independently explores the world in search of tasty morsels. Let X_n, Y_n, Z_n be the locations at time n of Drogon, Rhaegal, Viserion respectively, where time is assumed to be discrete and the number of possible locations is a finite number M . Their paths $X_0, X_1, X_2, \dots$; $Y_0, Y_1, Y_2, \dots$; and $Z_0, Z_1, Z_2, \dots$ are independent Markov chains with the same stationary distribution $\mathbf{s}$. Each dragon starts out at a random location generated according to the stationary distribution.
- (a) Let state 0 be home (so s_0 is the stationary probability of the home state). Find the expected number of times that Drogon is at home, up to time 24, i.e., the expected number of how many of $X_0, X_1, \dots, X_{24}$ are state 0 (in terms of s_0).
 - (b) If we want to track all 3 dragons simultaneously, we need to consider the vector of positions, (X_n, Y_n, Z_n) . There are M^3 possible values for this vector; assume that each is assigned a number from 1 to M^3 , e.g., if $M = 2$ we could encode the states $(0, 0, 0), (0, 0, 1), (0, 1, 0), \dots, (1, 1, 1)$ as $1, 2, 3, \dots, 8$ respectively. Let W_n be the number between 1 and M^3 representing (X_n, Y_n, Z_n) . Determine whether $W_0, W_1, \dots$ is a Markov chain.
 - (c) Given that all 3 dragons start at home at time 0, find the expected time it will take for all 3 to be at home again at the same time.
8. Consider the following Markov chain with $52! \approx 8 \times 10^{67}$ states. The states are the possible orderings of a standard 52-card deck. To run one step of the chain, pick 2 different cards from the deck, with all pairs equally likely, and swap the 2 cards. Find the stationary distribution of the chain.

Reversibility

9. (S) A Markov chain $X_0, X_1, \dots$ with state space $\{-3, -2, -1, 0, 1, 2, 3\}$ proceeds as follows. The chain starts at $X_0 = 0$. If X_n is not an endpoint (-3 or 3), then X_{n+1} is $X_n - 1$ or $X_n + 1$, each with probability $1/2$. Otherwise, the chain gets reflected off the endpoint, i.e., from 3 it always goes to 2 and from -3 it always goes to -2 . A diagram of the chain is shown below.



- (a) Is $|X_0|, |X_1|, |X_2|, \dots$ also a Markov chain? Explain.

Hint: For both (a) and (b), think about whether the past and future are conditionally independent given the present; don't do calculations with a 7 by 7 transition matrix!

- (b) Let sgn be the sign function: $\text{sgn}(x) = 1$ if $x > 0$, $\text{sgn}(x) = -1$ if $x < 0$, and $\text{sgn}(0) = 0$. Is $\text{sgn}(X_0), \text{sgn}(X_1), \text{sgn}(X_2), \dots$ a Markov chain? Explain.

- (c) Find the stationary distribution of the chain $X_0, X_1, X_2, \dots$.


- (d) Find a simple way to modify some of the transition probabilities q_{ij} for $i \in \{-3, 3\}$ to make the stationary distribution of the modified chain uniform over the states.

10. (S) Let G be an undirected network with nodes labeled $1, 2, \dots, M$ (edges from a node to itself are not allowed), where $M \geq 2$ and random walk on this network is irreducible. Let d_j be the degree of node j for each j . Create a Markov chain on the state space $1, 2, \dots, M$, with transitions as follows. From state i , generate a proposal j by choosing a uniformly random j such that there is an edge between i and j in G ; then go to j with probability $\min(d_i/d_j, 1)$, and stay at i otherwise.

- (a) Find the transition probability q_{ij} from i to j for this chain, for all states i, j .

- (b) Find the stationary distribution of this chain.

11. (S) (a) Consider a Markov chain on the state space $\{1, 2, \dots, 7\}$ with the states arranged in a "circle" as shown below, and transitions given by moving one step clockwise or counterclockwise with equal probabilities. For example, from state 6 , the chain moves to state 7 or state 5 with probability $1/2$ each; from state 7 , the chain moves to state 1 or state 6 with probability $1/2$ each. The chain starts at state 1 .



Find the stationary distribution of this chain.

- (b) Consider a new chain obtained by "unfolding the circle". Now the states are arranged as shown below. From state 1 the chain always goes to state 2 , and from state 7 the chain always goes to state 6 . Find the new stationary distribution.




12. (S) Let X_n be the price of a certain stock at the start of the n th day, and assume that $X_0, X_1, X_2, \dots$ follows a Markov chain with transition matrix Q . (Assume for simplicity that the stock price can never go below 0 or above a certain upper bound, and that it is always rounded to the nearest dollar.)

(a) A lazy investor only looks at the stock once a year, observing the values on days $0, 365, 2 \cdot 365, 3 \cdot 365, \dots$. So the investor observes $Y_0, Y_1, \dots$, where Y_n is the price after n years (which is $365n$ days; you can ignore leap years). Is $Y_0, Y_1, \dots$ also a Markov chain? Explain why or why not; if so, what is its transition matrix?

(b) The stock price is always an integer between \$0 and \$28. From each day to the next, the stock goes up or down by \$1 or \$2, all with equal probabilities (except for days when the stock is at or near a boundary, i.e., at \$0, \$1, \$27, or \$28).

If the stock is at \$0, it goes up to \$1 or \$2 on the next day (after receiving government bailout money). If the stock is at \$28, it goes down to \$27 or \$26 the next day. If the stock is at \$1, it either goes up to \$2 or \$3, or down to \$0 (with equal probabilities); similarly, if the stock is at \$27 it either goes up to \$28, or down to \$26 or \$25. Find the stationary distribution of the chain.

13. (S) In chess, the king can move one square at a time in any direction (horizontally, vertically, or diagonally).



For example, in the diagram, from the current position the king can move to any of 8 possible squares. A king is wandering around on an otherwise empty 8×8 chessboard, where for each move all possibilities are equally likely. Find the stationary distribution of this chain (of course, don't list out a vector of length 64 explicitly! Classify the 64 squares into types and say what the stationary probability is for a square of each type).

14. A chess piece is wandering around on an otherwise vacant 8×8 chessboard. At each move, the piece (a king, queen, rook, bishop, or knight) chooses uniformly at random where to go, among the legal choices (according to the rules of chess, which you should look up if you are unfamiliar with them).

(a) For each of these cases, determine whether the Markov chain is irreducible, and whether it is aperiodic.

Hint for the knight: Note that a knight's move always goes from a light square to a dark square or vice versa. A knight's tour is a sequence of knight moves on a chessboard such that

the knight visits each square exactly once. Many knight's tours exist.

(b) Suppose for this part that the piece is a rook, with initial position chosen uniformly at random. Find the distribution of where the rook is after n moves.

(c) Now suppose that the piece is a king, with initial position chosen deterministically to be the upper left corner square. Determine the expected number of moves it takes the king to return to that square, fully simplified, preferably in at most 140 characters.


(d) The stationary distribution for the random walk of the king from the previous part is not uniform over the 64 squares of the chessboard. A recipe for modifying the chain to obtain a uniform stationary distribution is as follows. Label the squares as $1, 2, \dots, 64$, and let d_i be the number of legal moves from square i . Suppose the king is currently at square i . The next move of the chain is determined as follows:

Step 1: Generate a proposal square j by picking uniformly at random among the legal moves from i .

Step 2: Flip a coin with probability $\min(d_i/d_j, 1)$ of Heads. If the coin lands Heads, go to j . Otherwise, stay at i .

Show that this modified chain has a stationary distribution that is uniform over the 64 squares.

15. (S) Find the stationary distribution of the Markov chain shown below, without using matrices. The number above each arrow is the corresponding transition probability.



16. There are two urns with a total of $2N$ distinguishable balls. Initially, the first urn has N white balls and the second urn has N black balls. At each stage, we pick a ball at random from each urn and interchange them. Let X_n be the number of black balls in the first urn at time n . This is a Markov chain on the state space $\{0, 1, \dots, N\}$.

(a) Give the transition probabilities of the chain.

(b) Show that $(s_0, s_1, \dots, s_N)$ where

$$s_i = \frac{\binom{N}{i} \binom{N}{N-i}}{\binom{2N}{N}}$$

is the stationary distribution, by verifying the reversibility condition.

17. Find the stationary distribution of a Markov chain $X_0, X_1, X_2, \dots$ on the state space $\{0, 1, \dots, 110\}$, with transition probabilities given by

$$P(X_{n+1} = j \mid X_n = 0) = p, \text{ for } j = 1, 2, \dots, 110;$$


$$P(X_{n+1} = 0 \mid X_n = 0) = 1 - 110p;$$

$$P(X_{n+1} = j \mid X_n = j) = 1 - r, \text{ for } j = 1, 2, \dots, 110;$$

$$P(X_{n+1} = 0 \mid X_n = j) = r, \text{ for } j = 1, 2, \dots, 110,$$

where p and r are constants with $0 < p < \frac{1}{110}$ and $0 < r < 1$.

18. Determine whether the Markov chain shown below is reversible, and find the stationary distribution of the chain. The label to the left of an arrow gives the corresponding transition probability.



19. Nausicaa Distribution sells distribution plushies on Etsy. They have two different photos of the Evil Cauchy plushie but do not know which is more effective in getting a customer to purchase an Evil Cauchy plushie. Each visitor to their website is shown one of the two photos (call them Photo A and Photo B), and then the visitor either does buy an Evil Cauchy ("success") or does not buy one ("failure").

Let a and b be the probabilities of success when Photo A is shown and when Photo B is shown, respectively. Even though the Evil Cauchy is irresistible, suppose that $0 < a < 1$ and $0 < b < 1$. Suppose that the following strategy is followed (note that the strategy can be followed without knowing a and b). Show the first visitor Photo A. If that visitor buys an Evil Cauchy, continue with Photo A for the next visitor; otherwise, switch to Photo B. Similarly, if the n th visitor is a "success" then show the $(n + 1)$ st visitor the same photo, and otherwise switch to the other photo.

(a) Show how to represent the resulting process as a Markov chain, drawing a diagram and giving the transition matrix. The states are A1, B1, A0, B0 (use this order for the transition matrix and stationary distribution), where, for example, being at state A1 means that the current visitor was shown Photo A and was a success.

(b) Determine whether this chain is reversible.

Hint: First think about which transition probabilities are zero and which are nonzero.

(c) Show that the stationary distribution is proportional to $(\frac{a}{1-a}, \frac{b}{1-b}, 1, 1)$, and find the stationary distribution.

(d) Show that for $a \neq b$, the stationary probability of success for each visitor is strictly better than the success probability that would be obtained by independently, randomly choosing (with equal probabilities) which photo to show to each visitor.

20. This exercise considers random walk on a weighted undirected network. Suppose that an undirected network is given, where each edge (i, j) has a nonnegative weight w_{ij} assigned to it (we allow $i = j$ as a possibility). We assume that $w_{ij} = w_{ji}$ since the edge from i to j is considered the same as the edge from j to i . To simplify notation, define $w_{ij} = 0$ whenever (i, j) is not an edge.

When at node i , the next step is determined by choosing an edge attached to i with probabilities proportional to the weights. For example, if the walk is at node 1 and there are 3 possible edges coming out from node 1, with weights 7, 1, 4, then the first of these 3 edges is traversed with probability $7/12$, the second is traversed with probability $1/12$, and the third is traversed with probability $4/12$. If all the weights equal 1, then the process reduces to the kind of random walk on a network that we studied earlier.

(a) Let $v_i = \sum_j w_{ij}$ for all nodes i . Show that the stationary distribution of node i is proportional to v_i .

(b) Show that every reversible Markov chain can be represented as a random walk on a weighted

undirected network. That is, given the transition matrix Q of a reversible Markov chain, show that we can choose the weights w_{ij} so that the random walk defined above is a Markov chain with transition matrix Q . Be sure to check that $w_{ij} = w_{ji}$.

Hint: Let $w_{ij} = s_i q_{ij}$, with $\mathbf{s}$ the stationary distribution and q_{ij} the (i, j) entry of Q .

Mixed Practice

21. (S) A cat and a mouse move independently back and forth between two rooms. At each time step, the cat moves from the current room to the other room with probability 0.8. Starting from room 1, the mouse moves to Room 2 with probability 0.3 (and remains otherwise). Starting from room 2, the mouse moves to room 1 with probability 0.6 (and remains otherwise).
- (a) Find the stationary distributions of the cat chain and of the mouse chain.
- (b) Note that there are 4 possible (cat, mouse) states: both in room 1, cat in room 1 and mouse in room 2, cat in room 2 and mouse in room 1, and both in room 2. Number these cases 1, 2, 3, 4, respectively, and let Z_n be the number representing the (cat, mouse) state at time n . Is $Z_0, Z_1, Z_2, \dots$ a Markov chain?
- (c) Now suppose that the cat will eat the mouse if they are in the same room. We wish to know the expected time (number of steps taken) until the cat eats the mouse for two initial configurations: when the cat starts in room 1 and the mouse starts in room 2, and vice versa. Set up a system of two linear equations in two unknowns whose solution is the desired values.
22. (a) Alice and Bob are wandering around randomly, independently of each other, in a house with M rooms, labeled 1, 2, ..., M . Let d_i be the number of doors in room i (leading to other rooms, not leading outside). At each step, Alice moves to another room by choosing randomly which door to go through (with equal probabilities). Bob does the same, independently. The Markov chain they each follow is irreducible and aperiodic. Let A_n and B_n be Alice's room and Bob's room at time n , respectively, for $n = 0, 1, 2, \dots$. Find $\lim_{n \rightarrow \infty} P(A_n = i, B_n = j)$.
- (b) With setup as in (a), let p_{ij} be the transition probability for going from room i to room j . Let t_{ik} be the expected first time at which Alice and Bob are in the same room, if Alice starts in room i and Bob starts in room k (note that $t_{ik} = 0$ for $i = k$). Provide a system of linear equations, the solution of which would yield t_{ik} for all rooms i, k .
- Hint: Condition on Alice's first move and Bob's first move.

23. Let $\{X_n\}$ be a Markov chain on states $\{0, 1, 2\}$ with transition matrix


$$Q = \begin{pmatrix} 0.8 & 0.2 & 0 \\ 0 & 0.8 & 0.2 \\ 0 & 0 & 1 \end{pmatrix}$$

The chain starts at $X_0 = 0$. Let T be the time it takes to reach state 2:

$$T = \min \{n : X_n = 2\}.$$

By drawing the Markov chain and telling a story, find $E(T)$ and $\text{Var}(T)$.

24. Consider the following Markov chain on the state space $\{1, 2, 3, 4, 5, 6\}$.



(a) Suppose the chain starts at state 1. Find the distribution of the number of times that the chain returns to state 1.

(b) In the long run, what fraction of the time does the chain spend in state 3? Explain briefly.

25. Let Q be the transition matrix of a Markov chain on the state space $\{1, 2, \dots, M\}$, such that state M is an absorbing state, i.e., from state M the chain can never leave. Suppose that from any other state, it is possible to reach M (in some number of steps).

(a) Which states are recurrent, and which are transient? Explain.

(b) What is the limit of Q^n as $n \rightarrow \infty$?

(c) For $i, j \in \{1, 2, \dots, M-1\}$, find the probability that the chain is at state j at time n , given that the chain is at state i at time 0 (your answer should be in terms of Q).

(d) For $i, j \in \{1, 2, \dots, M-1\}$, find the expected number of times that the chain is at state j up to (and including) time n , given that the chain is at state i at time 0 (in terms of Q).

(e) Let R be the $(M-1) \times (M-1)$ matrix obtained from Q by deleting the last row and the last column of Q . Show that the (i, j) entry of $(I - R)^{-1}$ is the expected number of times that the chain is at state j before absorption, given that it starts out at state i .

Hint: We have $I + R + R^2 + \dots = (I - R)^{-1}$, analogously to a geometric series. Also, if we partition Q as

$$Q = \left(\begin{array}{c|c} R & B \\ \hline 0 & 1 \end{array} \right)$$

where B is a $(M-1) \times 1$ matrix and 0 is the $1 \times (M-1)$ zero matrix, then

$$Q^k = \left(\begin{array}{c|c} R^k & B_k \\ \hline 0 & 1 \end{array} \right)$$

for some $(M-1) \times 1$ matrix B_k .


26. In the game called Chutes and Ladders, players try to be first to reach a certain destination on a board. The board is a grid of squares, numbered from 1 to the number of squares. The

board has some "chutes" and some "ladders", each of which connects a pair of squares. Here we will consider the one player version of the game (this can be extended to the multi-player version without too much trouble, since with more than one player, the players simply take turns independently until one reaches the destination).

On each turn, the player rolls a fair die, which determines how many squares forward to move on the grid, e.g., if the player is at square 5 and rolls a 3, then they advance to square 8. If the resulting square is the base of a ladder, the player gets to climb the ladder, instantly arriving at a more advanced square. If the resulting square is the top of a chute, the player instantly slides down to the bottom of the chute.

This game can be viewed naturally as a Markov chain: given where the player currently is on the board, the past history does not matter for computing, for example, the probability of winning within the next 3 moves.

Consider a simplified version of Chutes and Ladders, played on the 3×3 board shown below. The player starts out at square 1, and wants to get to square 9. On each move, a fair coin is flipped, and the player gets to advance 1 square if the coin lands Heads and 2 squares if the coin lands Tails. However, there are 2 ladders (shown as upward-pointing arrows) and 2 chutes (shown as downward-pointing arrows) on the board.



(a) Explain why, despite the fact that there are 9 squares, we can represent the game using the 5×5 transition matrix

$$Q = \begin{pmatrix} 0 & 0.5 & 0.5 & 0 & 0 \\ 0 & 0 & 0.5 & 0.5 & 0 \\ 0.5 & 0 & 0 & 0.5 & 0 \\ 0.5 & 0 & 0 & 0 & 0.5 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

(b) Find the mean and variance for the number of times the player will visit square 7, without using matrices or any messy calculations.

The remaining parts of this problem require matrix calculations that are best done on a computer. You can use whatever computing environment you want, but here is some information for how to do it in R. In any case, you should state what environment you used and include your code. To create the transition matrix in R, you can use the following commands:

```
a <- 0.5
Q <- matrix(c(0, 0, a, a, 0, a, 0, 0, 0, 0, a, a, 0, 0, 0, 0, a, a, 0, 0, 0, 0, a, 1), nrow = 5)
```

Some useful R commands for matrices are in Appendix B.2. In particular, `diag(n)` gives the

$n \times n$ identity matrix, solve (A) gives the inverse A^{-1} , and $A \% * \% B$ gives the product AB (note that $A * B$ does not do ordinary matrix multiplication). Matrix powers are not built into R, but you can compute A^k using $A \% \wedge \% k$ after installing and loading the `expm` package.

(c) Find the median duration of the game (where duration is the number of coin flips).

Hint: Relate the CDF of the duration to powers of Q .

(d) Find the mean duration of the game (with duration defined as above).

Hint: Relate the duration to the total amount of time spent in transient states, and apply Part (e) of the previous problem.

4.12 12 Markov Chain Monte Carlo

We have seen throughout this book that simulation is a powerful technique in probability. If you can't convince your friend that it is a good idea to switch doors in the Monty Hall problem, in one second you can simulate playing the game a few thousand times and your friend will just see that switching succeeds about $2/3$ of the time. If you're unsure how to calculate the mean and variance of an r.v. X but you know how to generate i.i.d. draws $X_1, X_2, \dots, X_n$ from that distribution, you can approximate the true mean and true variance using the sample mean and sample variance of the simulated draws:

$$E(X) \approx \frac{1}{n} (X_1 + \dots + X_n) = \bar{X}_n,$$
$$\text{Var}(X) \approx \frac{1}{n-1} \sum_{j=1}^n (X_j - \bar{X}_n)^2.$$

The law of large numbers tells us that these approximations will be good if n is large. We can get better and better approximations by increasing n , just by running the computer for a longer time (rather than having to struggle with a possibly intractable sum or integral). As discussed in Chapter 10, this simulation approach, where we generate random values to approximate a quantity, is called a Monte Carlo method.

A major limitation of the Monte Carlo idea above is that we need to know how to generate $X_1, X_2, \dots, X_n$ (hopefully efficiently, since we want to be able to make n large). For example, suppose that we want to simulate random draws from the continuous distribution with PDF f given by

$$f(x) \propto x^{3.1}(1-x)^{4.2}$$

for $0 < x < 1$ (and 0 otherwise). Staring at a density function does not immediately suggest how to get a random variable with that density. In this case we recognize the PDF of the $\text{Beta}(4.1, 5.2)$ distribution, so assuming we had access to a $\text{Unif}(0, 1)$ r.v., we could theoretically use universality of the Uniform. Unfortunately, finding the CDF of a Beta distribution is difficult, let alone the inverse of the CDF. See Example 12.1.4 for more about the Beta simulation problem.

Of course, the distributions that arise in real-world applications are often much more complicated than the Beta. For the Beta distribution, we know an expression for the normalizing constant in terms of the gamma function. But for many other distributions that are of scientific interest, the normalizing constant in the PDF or PMF is unknown, and beyond the reach of the fastest available computers and the fanciest available mathematical techniques.

This chapter introduces Markov chain Monte Carlo (MCMC), a powerful collection of algorithms that enable us to simulate from complicated distributions using Markov chains. The

development of MCMC has revolutionized statistics and scientific computation by vastly expanding the range of possible distributions that we can simulate from, including joint distributions in high dimensions. The basic idea is to build your own Markov chain so that the distribution of interest is the stationary distribution of the chain.

In the previous chapter, we looked at Markov chains whose transition matrix Q was specified, and we tried to find the stationary distribution $\mathbf{s}$ of the chain. In this chapter we do the reverse: starting with a distribution $\mathbf{s}$ that we want to simulate, we will engineer a Markov chain whose stationary distribution is $\mathbf{s}$. If we then run this engineered Markov chain for a very long time, the distribution of the chain will approach our desired $\mathbf{s}$.

But is it possible to create a transition matrix Q with the stationary distribution we desire? Even if it is possible, is solving this problem any easier than the original problem of how to simulate random draws from that distribution? In amazing generality, MCMC shows that it is possible to create a Markov chain with the desired stationary distribution in an easy-to-describe way, without having to know the normalizing constant for the distribution!

MCMC is now being applied to a large number of problems in the biological, natural, and physical sciences, and many different MCMC algorithms have been developed. Here we will introduce two of the most important and most widely used MCMC algorithms: the Metropolis-Hastings algorithm and Gibbs sampling. MCMC is an enormous and growing area of statistical computing; see Brooks, Gelman, Jones, and Meng [2] for much more about its theory, methods, and applications.

4.12.1 12.1 Metropolis-hastings

The Metropolis-Hastings algorithm is a general recipe that lets us start with any irreducible Markov chain on the state space of interest and then modify it into a new Markov chain that has the desired stationary distribution. This modification consists of introducing some selectiveness in the original chain: moves are proposed according to the original chain, but the proposal may or may not be accepted. For example, suppose the original chain is at a state called "Boston" and is about to transition to "San Francisco". Then for the new chain, we either accept the proposal and go to San Francisco, or we turn down the proposal and remain in Boston as the next step. With a careful choice of the probability of accepting the proposal, this simple modification guarantees that the new chain has the desired stationary distribution.

Algorithm 12.1.1 (Metropolis-Hastings). Let $\mathbf{s} = (s_1, \dots, s_M)$ be a desired stationary distribution on state space $\{1, \dots, M\}$. Assume that $s_i > 0$ for all i (if not, just delete any states i with $s_i = 0$ from the state space). Suppose that $P = (p_{ij})$ is the transition matrix for a Markov chain on state space $\{1, \dots, M\}$. Intuitively, P is a Markov chain that we know how to run but that doesn't have the desired stationary distribution.

Our goal is to modify P to construct a Markov chain $X_0, X_1, \dots$ with stationary distribution $\mathbf{s}$. We will give a Metropolis-Hastings algorithm for this. Start at any state X_0 (chosen randomly or deterministically), and suppose that the new chain is currently at X_n . To make one move of the new chain, do the following.

1. If $X_n = i$, propose a new state j using the transition probabilities in the i th row of the original transition matrix P .
2. Compute the acceptance probability

$$a_{ij} = \min \left(\frac{s_j p_{ji}}{s_i p_{ij}}, 1 \right)$$

3. Flip a coin that lands Heads with probability a_{ij} .

4. If the coin lands Heads, accept the proposal (i.e., go to j), setting $X_{n+1} = j$. Otherwise, reject the proposal (i.e., stay at i), setting $X_{n+1} = i$.

That is, the Metropolis-Hastings chain uses the original transition probabilities p_{ij} to propose where to go next, then accepts the proposal with probability a_{ij} , staying in its current state in the event of a rejection. An especially nice aspect of this algorithm is that the normalizing constant for $\mathbf{s}$ does not need to be known, since it cancels out in s_j/s_i anyway. For example, in some problems we may want the stationary distribution to be uniform over all states (i.e., $\mathbf{s} = (1/M, 1/M, \dots, 1/M)$), but the number of states M is large and unknown, and it would be a very hard counting problem to find M . Fortunately, $s_j/s_i = 1$ regardless of M , so we can simply say $\mathbf{s} \propto (1, 1, \dots, 1)$, and we can calculate a_{ij} without having to know M .

The p_{ij} in the denominator in a_{ij} will never be 0 when the algorithm is run, since if $p_{ij} = 0$ then the original chain will never propose going from i to j . Also, if $p_{ii} > 0$ it is possible that the proposal j will equal the current state i ; in that case, the chain stays at i regardless of whether the proposal is accepted. (Rejecting the proposal of staying at i but staying there anyway is like a child who just got grounded saying "yes, I will stay in my room, but not because you told me to!")

We will now show that the Metropolis-Hastings chain is reversible with stationary distribution $\mathbf{s}$.

Proof. Let Q be the transition matrix of the Metropolis-Hastings chain. We just need to check the reversibility condition $s_i q_{ij} = s_j q_{ji}$ for all i and j . This is clear for $i = j$, so assume $i \neq j$. If $q_{ij} > 0$, then $p_{ij} > 0$ (the chain can't get from i to j if it can't even propose going from i to j) and $p_{ji} > 0$ (otherwise the acceptance probability would be 0). Conversely, if $p_{ij} > 0$ and $p_{ji} > 0$, then $q_{ji} > 0$. So q_{ij} and q_{ji} are either both zero or both nonzero. We can assume they are both nonzero. Then

$$q_{ij} = p_{ij} a_{ij}$$

since, starting at i , the only way to get to j is first to propose doing so and then to accept the proposal. First consider the case $s_j p_{ji} \leq s_i p_{ij}$. We have

$$a_{ij} = \frac{s_j p_{ji}}{s_i p_{ij}}, a_{ji} = 1,$$

so

$$s_i q_{ij} = s_i p_{ij} a_{ij} = s_i p_{ij} \frac{s_j p_{ji}}{s_i p_{ij}} = s_j p_{ji} = s_j p_{ji} a_{ji} = s_j q_{ji}$$

Symmetrically, if $s_j p_{ji} > s_i p_{ij}$, we again have $s_i q_{ij} = s_j q_{ji}$, by switching the roles of i and j in the preceding calculation. Since the reversibility condition holds, $\mathbf{s}$ is the stationary distribution of the chain with transition matrix Q .

12.1.2. The Metropolis-Hastings algorithm is an extremely general way to construct a Markov chain with a desired stationary distribution. In the above formulation, both $\mathbf{s}$ and P were very general, and nothing was stipulated about their being related (aside from being on the same state space). In practice, however, the choice of the proposal distribution is extremely important since it can make an enormous difference in how quickly the chain converges to its stationary distribution.

How to choose a good proposal distribution is a complicated topic and will not be discussed in detail here. Intuitively, a proposal distribution with a very low acceptance rate will be slow to converge (since the chain will rarely move anywhere). But a high acceptance rate may not be ideal either, since it may indicate that the chain tends to make small, timid proposals. In a large state space, such a chain will take a very long time to explore the entire space.

Here are a couple examples of how the Metropolis-Hastings algorithm can be used to simulate from distributions.

Example 12.1.3 (Zipf distribution simulation). Let $M \geq 2$ be an integer. An r.v. X has the Zipf distribution with parameter $a > 0$ if its PMF is

$$P(X = k) = \frac{1/k^a}{\sum_{j=1}^M (1/j^a)},$$

for $k = 1, 2, \dots, M$ (and 0 otherwise). This distribution is widely used in linguistics for studying frequencies of words.

Create a Markov chain $X_0, X_1, \dots$ whose stationary distribution is the Zipf distribution, and such that $|X_{n+1} - X_n| \leq 1$ for all n . Your answer should provide a simple, precise description of how each move of the chain is obtained, i.e., how to transition from X_n to X_{n+1} for each n .

Solution:

We can use the Metropolis-Hastings algorithm, after coming up with a proposal distribution. There are many possible proposal distributions, but one simple choice is the following random walk on $\{1, 2, \dots, M\}$. From state i with $i \neq 1, i \neq M$, move to state $i - 1$ or $i + 1$, with probability $1/2$ each. From state 1, stay there or move to state 2, with probability $1/2$ each. From state M , stay there or move to state $M - 1$, with probability $1/2$ each. This chain is shown below.



FIGURE 12.1 Proposal chain for Zipf distribution simulation.

Let P be the transition matrix of this chain. The stationary distribution for P is uniform because P is a symmetric matrix, so Proposition 11.4.3 applies. MetropolisHastings lets us transmogrify P into a chain whose stationary distribution is Zipf's distribution.

Let X_0 be any starting state, and generate a chain $X_0, X_1, \dots$ as follows. If the chain is currently at state i , then:

1. Generate a proposal state j according to the proposal chain P .
2. Accept the proposal with probability $\min(i^a/j^a, 1)$. If the proposal is accepted, go to j ; otherwise, stay at i .

This chain is easy to implement and a move requires very little computation; note that the normalizing constant $\sum_{j=1}^M (1/j^a)$ is not needed to run the chain. $\square$

Example 12.1.4 (Beta simulation). Let us now return to the Beta simulation problem introduced at the beginning of the chapter. Suppose that we want to generate $W \sim \text{Beta}(a, b)$, but we don't know about the `rbeta` command in R. Instead, what we have available are i.i.d. $\text{Unif}(0, 1)$ r.v.s.

- (a) How can we generate W exactly if a and b are positive integers, using a story and universality of the Uniform?
- (b) How can we generate W which is approximately $\text{Beta}(a, b)$ if a and b are any positive real numbers, with the help of a Markov chain on the state space $(0, 1)$?

Solution:

- (a) Applying universality of the Uniform directly for the Beta is hard, so let's first use the bank-post office story: if $X \sim \text{Gamma}(a, 1)$ and $Y \sim \text{Gamma}(b, 1)$ are

independent, then $X/(X + Y) \sim \text{Beta}(a, b)$. So if we can simulate the Gamma distribution, then we can simulate the Beta distribution!

To simulate $X \sim \text{Gamma}(a, 1)$, we can use $X_1 + X_2 + \cdots + X_a$ with the X_j i.i.d. $\text{Expo}(1)$; similarly, we can simulate $Y \sim \text{Gamma}(b, 1)$ as the sum of b i.i.d. $\text{Expo}(1)$ r.v.s. Lastly, by taking the inverse of the $\text{Expo}(1)$ CDF and applying universality of the Uniform, $-\log(1 - U) \sim \text{Expo}(1)$ for $U \sim \text{Unif}(0, 1)$, so we can easily construct as many $\text{Expo}(1)$ r.v.s as we want.

(b) Let's use the Metropolis-Hastings algorithm. We have only introduced Metropolis-Hastings for finite state spaces, but the ideas are analogous for an infinite state space. A simple proposal chain we have available consists of independent $\text{Unif}(0, 1)$ r.v.s. That is, the proposed state on the interval $(0, 1)$ is always a fresh $\text{Unif}(0, 1)$, independent of the current state. The resulting Metropolis-Hastings chain is called an independence sampler.

Let W_0 be any starting state, and generate a chain $W_0, W_1, \dots$ as follows. If the chain is currently at state w , then:

1. Generate a proposal u by drawing a $\text{Unif}(0, 1)$ r.v.
2. Accept the proposal with probability $\min\left(\frac{u^{a-1}(1-u)^{b-1}}{w^{a-1}(1-w)^{b-1}}, 1\right)$. If the proposal is accepted, go to u ; otherwise, stay at w .

Again, the normalizing constant was not needed in order to run the chain. In obtaining the acceptance probability, the $\text{Beta}(a, b)$ PDF plays the role of s since it's the desired stationary distribution, and the $\text{Unif}(0, 1)$ PDF plays the role of p_{ij} (and p_{ji}) since the proposals are $\text{Unif}(0, 1)$ r.v.s, independent of the current state.

Running the Markov chain, we have that $W_n, W_{n+1}, W_{n+2}, \dots$ are approximately $\text{Beta}(a, b)$ for n large. Note that these are correlated r.v.s., not i.i.d. draws.

12.1.5 (MCMC produces correlated samples). A major question in running a Markov chain $X_0, X_1, \dots$ for a Monte Carlo computation is how long to run it. In part, this is because it is usually hard to know how close the chain's distribution at time n will be to the stationary distribution. Another issue is that $X_0, X_1, \dots$ are correlated in general. Some chains tend to get stuck in certain regions of the state space, rather than exploring the whole space. If a chain can get stuck easily, then X_n may be highly positively correlated with X_{n+1} . The autocorrelation at lag k is the correlation between X_n and the value k steps later, X_{n+k} , in the limit as n grows. It is desirable for the autocorrelation at lag k to approach 0 rapidly as k increases. High autocorrelation tends to mean high variances for Monte Carlo approximations.

Analysis of how long to run a chain and finding diagnostics for whether the chain has been run long enough are active research areas. Some general advice is to run your chains for a very large number of steps and to try chains from diverse starting points to see how stable the results are.

Metropolis-Hastings is often useful even for immense state spaces. It can even be useful for problems that may not sound at first like they have anything to do with simulating a distribution, such as code-breaking.

Example 12.1.6 (Code-breaking). Markov chains have recently been applied to code-breaking; this example will introduce one way in which this can be done. (For further information about such applications, see Diaconis [6] and Chen and Rosenthal [3].) A substitution cipher is a permutation g of the letters from a to z , where a message is enciphered by replacing each letter α by $g(\alpha)$. For example, if g is the permutation given by

```
abcdefghijklmnopqrstuvwxyz
zyxwvutsrqponmlkjihgfedcba
```

where the second row lists the values $g(a), g(b), \dots, g(z)$, then we would encipher the word "statistics" as "hgzgrhgrxh". (We could also include capital letters, spaces, and punctuation marks if desired.) The state space is all $26! \approx 4 \cdot 10^{26}$ permutations of the letters a through z . This is an extremely large space: if we had to try decoding a text using each of these permutations, and could handle one permutation per nanosecond, it would still take over 12 billion years to work through all the permutations. So a brute-force investigation that goes through each permutation one by one is infeasible; instead, we will look at random permutations. (a) Consider the Markov chain that picks two different random coordinates between 1 and 26 and swaps those entries of the 2 nd row, e.g., if we pick 7 and 20 , then

```
abcdefghijklmnopqrstuvwxyz
zyxwvutsrqponmlkjihgfedcba
```

becomes

```
abcdefghijklmnopqrstuvwxyz
zyxwvugsrqponmlkjihgfedcba
```

Find the probability of going from a permutation g to a permutation h in one step (for all g, h), and find the stationary distribution of this chain.

(b) Suppose we have a system that assigns a positive "score" $s(g)$ to each permutation g . Intuitively, this could be a measure of how likely it would be to get the observed enciphered text, given that g was the cipher used. Use the MetropolisHastings algorithm to construct a Markov chain whose stationary distribution is proportional to the list of all scores $s(g)$.

Solution:

(a) The probability of going from g to h in one step is 0 unless h can be obtained from g by swapping 2 entries of the second row. Assuming that h can be obtained in this way, the probability is $\frac{1}{\binom{26}{2}}$, since there are $\binom{26}{2}$ such swaps, all equally likely.

This Markov chain is irreducible, since by performing enough swaps we can get from any permutation to any other permutation. (Imagine rearranging a deck of cards by swapping cards two at a time; it is possible to reorder the cards in any desired configuration by doing this enough times.) Note that $p(g, h) = p(h, g)$, where $p(g, h)$ is the transition probability of going from g to h . Since the transition matrix is symmetric, the stationary distribution is uniform over all $26!$ permutations of the letters a through z .

(b) For our proposal chain, we'll use the chain from (a). Starting from any state g , generate a proposal h using the chain from (a). Flip a coin with probability $\min(s(h)/s(g), 1)$ of Heads. If Heads, go to h ; if Tails, stay at g .

To prove this has the desired stationary distribution, we can appeal to our general proof of Algorithm 12.1.1 or check the reversibility condition directly. For practice, we'll do the latter: we need $s(g)q(g, h) = s(h)q(h, g)$ for all g and h , where $q(g, h)$ is the transition probability from g to h in the modified chain. If $g = h$ or $q(g, h) = 0$, then the equation clearly holds, so assume $g \neq h$ and $q(g, h) \neq 0$. Let $p(g, h)$ be the transition probability from (a) (which is the probability of proposing h when at g). First consider the case that $s(g) \leq s(h)$. Then $q(g, h) = p(g, h)$ and

$$q(h, g) = p(h, g) \frac{s(g)}{s(h)} = p(g, h) \frac{s(g)}{s(h)} = q(g, h) \frac{s(g)}{s(h)}$$

so $s(g)q(g, h) = s(h)q(h, g)$. Now consider the case that $s(h) < s(g)$. By a symmetric argument (reversing the roles of g and h), we again have $s(g)q(g, h) = s(h)q(h, g)$. Thus, the stationary probability of g is proportional to its score $s(g)$.

In other words, using the Metropolis-Hastings algorithm, we started with a Markov chain that was equally likely to visit all of the ciphers in the long run and created a Markov chain whose stationary distribution sorts the ciphers according to their scores, visiting the most promising ciphers the most often in the long run.

Here is another MCMC example with an immense state space and which at first sight might not seem to have much connection with simulating a distribution. This example points to the fact that MCMC can be used not only for sampling but also for optimization.

Example 12.1.7 (Knapsack problem). Bilbo the Burglar finds m treasures in Smaug's Lair. Bilbo is deciding which treasures to steal (or justly reclaim, depending on one's point of view); he can't take everything all at once, since the maximum weight he can carry is w pounds. Label the treasures from 1 to m , and suppose that the j th treasure is worth g_j gold pieces and weighs w_j pounds. So Bilbo must choose a vector $x = (x_1, \dots, x_m)$, where x_j is 1 if he steals the j th treasure and 0 otherwise, such that the total weight of the treasures j with $x_j = 1$ is at most w . Let C be the space of all such vectors, so C consists of all binary vectors $(x_1, \dots, x_m)$ with $\sum_{j=1}^m x_j w_j \leq w$.

Bilbo wishes to maximize the total worth of the treasure he takes. Finding an optimal solution is an extremely difficult problem, known as the knapsack problem, which has a long history in computer science. A brute force solution would be

completely infeasible in general. Bilbo decides instead to explore the space C using MCMC—luckily, he has a laptop running R with him.

(a) Consider the following Markov chain. Start at $(0, 0, \dots, 0)$. One move of the chain is as follows. Suppose the current state is $x = (x_1, \dots, x_m)$. Choose a uniformly random J in $\{1, 2, \dots, m\}$, and obtain y from x by replacing x_J with $1 - x_J$ (i.e., toggle whether that treasure will be taken). If y is not in C , stay at x ; if y is in C , move to y . Show that the uniform distribution over C is stationary for this chain.

(b) Show that the chain from (a) is irreducible, and that it may or may not be aperiodic (depending on $w, w_1, \dots, w_m$).

(c) The chain from (a) is a useful way to get approximately uniform solutions, but Bilbo is more interested in finding solutions where the value (in gold pieces) is high. In this part, the goal is to construct a Markov chain with a stationary distribution that puts much higher probability on any particular high-value solution than on any particular low-value solution. Specifically, suppose that we want to simulate from the distribution

$$s(x) \propto e^{\beta V(x)},$$

where $V(x) = \sum_{j=1}^m x_j g_j$ is the value of x in gold pieces and β is a positive constant. The idea behind this distribution is to give exponentially more probability to each high-value solution than to each low-value solution. Create a Markov chain whose stationary distribution is as desired.

Solution:

(a) The transition matrix is symmetric since for $x \neq y$, the transition probabilities from x to y and from y to x are either both 0 or both $1/m$. So the stationary distribution is uniform over C .

(b) We can go from any $x \in C$ to $(0, 0, \dots, 0)$ by dropping treasures one at a time. We can go from $(0, 0, \dots, 0)$ to any $y \in C$ by picking up treasures one at a time. Combining these, we can go from anywhere to anywhere, so the chain is irreducible.

To study periodicity, let's look at some simple cases. First consider the simple case where $w_1 + \dots + w_m < w$, i.e., Bilbo can carry all the treasure at the same time. Then all binary vectors of length m are allowed. So the period of $(0, 0, \dots, 0)$ is 2 since, starting at that state,

Bilbo needs to pick up and then put down a treasure in order to get back to that state. In fact, if Bilbo starts at $(0, 0, \dots, 0)$, after any odd number of moves he will be carrying an odd number of treasures.

Now consider the case where $w_1 > w$, i.e., the first treasure is too heavy for Bilbo. From any $x \in C$, there is a $1/m$ chance that the chain will try to pick up the first treasure, and if that happens, the chain will stay at x . So the period of each state is 1.

(c) We can apply Metropolis-Hastings using the chain from (a) to make proposals. Start at $(0, 0, \dots, 0)$. Suppose the current state is $x = (x_1, \dots, x_m)$. Then:

1. Choose a uniformly random J in $\{1, 2, \dots, m\}$, and obtain y from x by replacing x_J with $1 - x_J$.
2. If y is not in C , stay at x . If y is in C , flip a coin that lands Heads with probability $\min(1, e^{\beta(V(y) - V(x))})$. If the coin lands Heads, go to y ; otherwise, stay at x .

This chain will converge to the desired stationary distribution. But how should β be chosen? If β is very large, then the best solutions are given very high probability, but the chain may be very slow to converge to the stationary distribution since it can easily get stuck in local modes: the chain may find itself in a state which, while not globally optimal, is still better than the other states that can be reached in one step, and then the probability of rejecting proposals to go elsewhere may be very high. On the other hand, if β is close to 0, then it's easy for the chain to explore the space, but there isn't as much incentive for the chain to uncover good solutions.

An optimization technique called simulated annealing avoids having to choose one value of β . Instead, one specifies a sequence of β values, such that β gradually increases over time. At first, β is small and the space C can be explored broadly. As β gets larger and larger, the stationary distribution becomes more and more concentrated on the best solution or solutions. The name "simulated annealing" comes from an analogy with the annealing of metals, a process in which a metal is heated to a high temperature and then gradually cooled until it reaches a very strong, stable state; β corresponds to the reciprocal of temperature.

As mentioned in Example 12.1.4, the Metropolis-Hastings algorithm can also be applied in a continuous state space, using PDFs instead of PMFs. This is extremely useful in Bayesian inference, where we often want to study the posterior distribution of an unknown parameter. This posterior distribution may be very complicated to work with analytically, and may have an unknown normalizing constant.

The MCMC approach is to obtain a large number of draws from a Markov chain whose stationary distribution is the posterior distribution. We can then use these draws to approximate the true posterior distribution. For example, we can estimate the posterior mean using the sample mean of these draws, and the posterior median using the sample median of the draws.

Gelman et al. [10] and McElreath [18] provide extensive introductions to Bayesian thinking and Bayesian data analysis, with wide ranges of applications and emphasis on statistical modeling and simulation. Various supplementary materials for these two books can be found at <http://www.stat.columbia.edu/~gelman/book> and <https://xcelab.net/rm/statistical-rethinking> respectively.

Example 12.1.8 (Normal-Normal conjugacy). Let $Y \mid \theta \sim \mathcal{N}(\theta, \sigma^2)$, where σ^2 is known but θ is unknown. Using the Bayesian framework, we treat θ as a random variable, with prior given by $\theta \sim \mathcal{N}(\mu, \tau^2)$ for some known constants μ and τ^2 . That is, we have the two-level model

$$\begin{aligned}\theta &\sim \mathcal{N}(\mu, \tau^2) \\ Y \mid \theta &\sim \mathcal{N}(\theta, \sigma^2)\end{aligned}$$

Describe how to use the Metropolis-Hastings algorithm to find the posterior mean and variance of θ after observing the value of Y .

Solution:

After observing $Y = y$, we can update our prior uncertainty for θ using Bayes' rule. Because we are interested in the posterior distribution of θ , any terms not depending on θ can be treated as part of the normalizing constant. Thus,

$$f_{\theta|Y}(\theta | y) \propto f_{Y|\theta}(y | \theta) f_{\theta}(\theta) \propto e^{-\frac{1}{2\sigma^2}(y-\theta)^2} e^{-\frac{1}{2\tau^2}(\theta-\mu)^2}$$

Since we have a quadratic function of θ in the exponent, we recognize the posterior PDF of θ as a Normal PDF. The posterior distribution stays in the Normal family, which tells us that the Normal is the conjugate prior of the Normal. In fact, by completing the square (a rather tedious calculation which we shall omit), we can obtain an explicit formula for the posterior distribution of θ :

$$\theta \Big| Y = y \sim \mathcal{N} \left(\frac{\frac{1}{\sigma^2}}{\frac{1}{\sigma^2} + \frac{1}{\tau^2}} y + \frac{\frac{1}{\tau^2}}{\frac{1}{\sigma^2} + \frac{1}{\tau^2}} \mu, \frac{1}{\frac{1}{\sigma^2} + \frac{1}{\tau^2}} \right)$$

Let's try to make sense of this formula.

- It says that the posterior mean of θ , $E(\theta | Y = y)$, is a weighted average of the prior mean μ and the observed data y . The weights are determined by how certain we are about θ before getting the data and how precisely the data are measured. If we are already very sure about θ even before getting the data, then τ^2 will be small and $1/\tau^2$ will be large, which will give a lot of weight to the prior mean μ . On the other hand, if the data are very precise, then σ^2 will be small and $1/\sigma^2$ will be large, which will give a lot of weight to the data y .
- For the posterior variance, if we define precision to be the reciprocal of variance, then the result simply says that the posterior precision of θ is the sum of the prior precision $1/\tau^2$ and the data precision $1/\sigma^2$.

This is all well and good, but let's suppose we didn't know how to complete the square, or that we wanted to check our calculations for specific values of y, σ^2, μ , and τ^2 . We can do this by simulating from the posterior distribution of θ , using the Metropolis-Hastings algorithm to construct a Markov chain whose stationary distribution is $f_{\theta|Y}(\theta | y)$. The same method can also be applied to a wide variety of distributions that are far more complicated than the Normal to work with analytically. A Metropolis-Hastings algorithm for generating $\theta_0, \theta_1, \dots$ is as follows.

1. If $\theta_n = x$, propose a new state x' according to some transition rule. One way to do this in a continuous state space is to generate a Normal r.v. ϵ_n with mean 0 and add it onto the current state to get the proposed state: in other words, we generate $\epsilon_n \sim \mathcal{N}(0, d^2)$ for some constant d , and then set $x' = x + \epsilon_n$. This is the analog of a transition matrix for a continuous state space. The only additional detail is deciding d ; in practice, we try to choose a moderate value that is neither too large nor too small.
2. The acceptance probability is

$$a(x, x') = \min \left(\frac{s(x') p(x', x)}{s(x) p(x, x')}, 1 \right),$$

where s is the desired stationary PDF (this was a PMF in the discrete case) and $p(x, x')$ is the probability density of proposing x' from x (this was p_{ij} in the discrete case).

In this problem, we want the stationary PDF to be $f_{\theta|Y}$, so we'll use that for s . As for $p(x, x')$, proposing x' from x is the same as having $\epsilon_n = x' - x$, so we evaluate the PDF of ϵ_n at $x' - x$ to get

$$p(x, x') = \frac{1}{\sqrt{2\pi}d} e^{-\frac{1}{2d^2}(x'-x)^2}$$

However, since $p(x', x) = p(x, x')$, these terms cancel from the acceptance probability, leaving us with

$$a(x, x') = \min \left(\frac{f_{\theta|Y}(x' | y)}{f_{\theta|Y}(x | y)}, 1 \right).$$

Once again, the normalizing constant cancels in the numerator and denominator of the acceptance probability.

3. Flip a coin that lands Heads with probability $a(x, x')$, independently of the Markov chain.

4. If the coin lands Heads, accept the proposal and set $\theta_{n+1} = x'$. Otherwise, stay in place and set $\theta_{n+1} = x$.

We ran the algorithm for 10^4 iterations with the settings $Y = 3, \mu = 0, \sigma^2 = 1, \tau^2 = 4$, and $d = 1$. Figure 12.2 shows a histogram of the resulting draws from the posterior distribution of θ . The posterior distribution indeed looks like a Normal curve. We can estimate the posterior mean and variance using the sample mean and sample variance. For the draws we obtained, the sample mean is 2.4 and the sample variance is 0.8. These are in close agreement with the theoretical values:

$$E(\theta | Y = 3) = \frac{\frac{1}{\sigma^2}}{\frac{1}{\sigma^2} + \frac{1}{\tau^2}} y + \frac{\frac{1}{\tau^2}}{\frac{1}{\sigma^2} + \frac{1}{\tau^2}} \mu = \frac{1}{1 + \frac{1}{4}} \cdot 3 + \frac{\frac{1}{4}}{1 + \frac{1}{4}} \cdot 0 = 2.8,$$

$$\text{Var}(\theta | Y = 3) = \frac{1}{\frac{1}{\sigma^2} + \frac{1}{\tau^2}} = \frac{1}{1 + \frac{1}{4}} = 0.8.$$

The posterior mean is closer to the observed data than to the prior mean, which makes sense because τ^2 is larger than σ^2 , corresponding to a relatively high level of prior uncertainty. Using the code provided in the R section of this chapter, you can see how the posterior distribution changes for different values of y, μ, σ^2 and τ^2 .




FIGURE 12.2 Histogram of 10^4 draws from the posterior distribution of θ given $Y = 3$, obtained using Metropolis-Hastings with $\mu = 0, \sigma^2 = 1$, and $\tau^2 = 4$. The sample mean is 2.4 and the sample variance is 0.8, in agreement with the theoretical values.




FIGURE 12.3
Trace plots of θ_n
as a function of the iteration number n , for $d = 100, 1, 0.01$.

To help diagnose whether our Markov chain is adequately exploring the state space, we can make a trace plot, which is a plot of the samples θ_n as a function of n . Figure 12.3 shows three trace plots corresponding to three different choices for the standard deviation d of the proposals, namely $d = 100$, $d = 1$, and $d = 0.01$. The trace plot for $d = 100$ has numerous flat regions where the chain is staying in place. This indicates that d is too large, so the proposals are often rejected. On the other hand, $d = 0.01$ is too small; we can see from the trace plot that the chain takes tiny steps and is unable to venture very far from its starting point. The trace plot for $d = 1$ is just right, exhibiting neither the low acceptance rate of the $d = 100$ chain, nor the restricted mobility of the $d = 0.01$ chain.

In this example, the posterior distribution of θ is available analytically, so we used MCMC for illustrative purposes, to show that the results obtained by MCMC agree with their theoretical counterparts. But the same technique also applies in problems where the prior isn't conjugate and the posterior isn't a named distribution.

4.12.2 12.2 Gibbs Sampling

Gibbs sampling is an MCMC algorithm for obtaining approximate draws from a joint distribution, based on sampling from conditional distributions one at a time: at each stage, one variable is updated (keeping all the other variables fixed) by drawing from the conditional distribution of that variable given all the other variables. This approach is especially useful in problems where these conditional distributions are pleasant to work with.

First we will run through how the Gibbs sampler works in the bivariate case, where the desired stationary distribution is the joint PMF of discrete r.v.s X and Y . There are several forms of Gibbs samplers, depending on the order in which updates are done. We will introduce two major kinds of Gibbs sampler: systematic scan, in which the updates sweep through the components in a deterministic order, and random scan, in which a randomly chosen component is updated at each stage.

Algorithm 12.2.1 (Systematic scan Gibbs sampler). Let X and Y be discrete r.v.s with joint PMF $p_{X,Y}(x, y) = P(X = x, Y = y)$. We wish to construct a two-dimensional Markov chain (X_n, Y_n) whose stationary distribution is $p_{X,Y}$. The systematic scan Gibbs sampler proceeds by updating the X -component and the Y -component in alternation. If the current state is $(X_n, Y_n) = (x_n, y_n)$, then we update the X -component while holding the Y -component fixed, and then update the Y -component while holding the X -component fixed:

1. Draw x_{n+1} from the conditional distribution of X given $Y = y_n$, and set $X_{n+1} = x_{n+1}$.
2. Draw y_{n+1} from the conditional distribution of Y given $X = x_{n+1}$, and set $Y_{n+1} = y_{n+1}$.

Repeating steps 1 and 2 over and over, the stationary distribution of the chain $(X_0, Y_0), (X_1, Y_1), (X_2, Y_2), \dots$ is $p_{X,Y}$.

Algorithm 12.2.2 (Random scan Gibbs sampler). As above, let X and Y be discrete r.v.s with joint PMF $p_{X,Y}(x, y)$. We wish to construct a two-dimensional

Markov chain (X_n, Y_n) whose stationary distribution is $p_{X,Y}$. Each move of the random scan Gibbs sampler picks a uniformly random component and updates it, according to the conditional distribution given the other component:

1. Choose which component to update, with equal probabilities.
2. If the X -component was chosen, draw a value x_{n+1} from the conditional distribution of X given $Y = y_n$, and set $X_{n+1} = x_{n+1}, Y_{n+1} = y_n$. Similarly, if the Y -component was chosen, draw a value y_{n+1} from the conditional distribution of Y given $X = x_n$, and set $X_{n+1} = x_n, Y_{n+1} = y_{n+1}$.

Repeating steps 1 and 2 over and over, the stationary distribution of the chain $(X_0, Y_0), (X_1, Y_1), (X_2, Y_2), \dots$ is $p_{X,Y}$.

Gibbs sampling generalizes naturally to higher dimensions. If we want to sample from a d -dimensional joint distribution, the Markov chain we construct will be a sequence of d -dimensional random vectors. At each stage, we choose one component of the vector to update, and we draw from the conditional distribution of that component given the most recent values of the other components. We can either cycle through the components of the vector in a systematic order, or choose a random component to update each time.

The Gibbs sampler is less flexible than the Metropolis-Hastings algorithm in the sense that we don't get to choose a proposal distribution; this also makes it simpler in the sense that we don't have to choose a proposal distribution. The flavors of Gibbs and Metropolis-Hastings are rather different, in that Gibbs emphasizes conditional distributions while Metropolis-Hastings emphasizes acceptance probabilities. But the algorithms are closely connected, as we show below.

Theorem 12.2.3 (Random scan Gibbs as Metropolis-Hastings). The random scan Gibbs sampler is a special case of the Metropolis-Hastings algorithm, in which the proposal is always accepted. In particular, it follows that the stationary distribution of the random scan Gibbs sampler is as desired.

Proof. We will show this in two dimensions, but the proof is similar in any dimension. Let X and Y be discrete r.v.s whose joint PMF is the desired stationary distribution. Let's work out what the Metropolis-Hastings algorithm says to do, using the following proposal distribution: from (x, y) , randomly update one coordinate by running one move of the random scan Gibbs sampler.

To simplify notation, write

$$P(X = x, Y = y) = p(x, y), P(Y = y \mid X = x) = p(y \mid x), P(X = x \mid Y = y) = p(x \mid y).$$

More formally, we should write $p_{Y|X}(y \mid x)$ instead of $p(y \mid x)$, to avoid issues like wondering what $p(5 \mid 3)$ means. But writing $p(y \mid x)$ is more compact and does not create ambiguity in this proof.

Let's compute the Metropolis-Hastings acceptance probability for going from (x, y) to (x', y') . The states (x, y) and (x', y') must be equal in at least one component, since the proposal says to update only one component. Suppose that $x = x'$ (the case $y = y'$ can be handled symmetrically). Then the acceptance probability is

$$\frac{p(x, y') p(y \mid x)^{\frac{1}{2}}}{p(x, y) p(y' \mid x)^{\frac{1}{2}}} = \frac{p(x) p(y' \mid x) p(y \mid x)}{p(x) p(y \mid x) p(y' \mid x)} = 1.$$

Thus, this Metropolis-Hastings algorithm always accepts the proposal! So it's just running the random scan Gibbs sampler without modifying it. $\square$

Let's study some concrete examples of Gibbs samplers.

Example 12.2.4 (Graph coloring). Let G be a network (also called a graph): there are n nodes, and for each pair of distinct nodes, there either is or isn't an edge joining them. We have a set of k colors, e.g., if $k = 7$, the color set may be {red, orange, yellow, green, blue, indigo, violet}. A k -coloring of the network is an assignment of a color to each node, such that two nodes joined by an edge can't be the same color. For example, a 3-coloring of a network is illustrated below. Graph coloring is an important topic in computer science, with wide-ranging applications such as task scheduling and the game of Sudoku.




FIGURE 12.4
A 3-coloring of a network.

Suppose that it is possible to k -color G . Form a Markov chain on the space of all k -colorings of G , with transitions as follows: starting with a k -coloring of G , pick a uniformly random node, figure out what the legal colors are for that node, and then repaint that node with a uniformly random legal color (note that this random color may be the same as the current color). Show that this Markov chain is reversible, and find its stationary distribution.

Solution:

Let C be the set of all k -colorings of G , and let q_{ij} be the transition probability of going from i to j for any k -colorings i and j in C . We will show that $q_{ij} = q_{ji}$, which implies that the stationary distribution is uniform on C .

For any k -coloring i and node v , let $L(i, v)$ be the number of legal colorings for node v , keeping the colors of all other nodes the same as they are in i . If k -colorings i and j differ at more than one node, then $q_{ij} = 0 = q_{ji}$. If $i = j$, then obviously $q_{ij} = q_{ji}$. If i and j differ at exactly one node v , then $L(i, v) = L(j, v)$, so

$$q_{ij} = \frac{1}{n} \frac{1}{L(i, v)} = \frac{1}{n} \frac{1}{L(j, v)} = q_{ji}.$$

So the transition matrix is symmetric, which shows that the stationary distribution is uniform over the state space.

How is this an example of Gibbs sampling? Think of each node in the graph as a discrete r.v. that can take on k possible values. These nodes have a joint distribution, and the constraint

that connected nodes cannot have the same color imposes a complicated dependence structure between nodes.

We would like to sample a random k -coloring of the entire graph; that is, we want to draw from the joint distribution of all the nodes. Since this is difficult, we instead condition on all but one node. If the joint distribution is to be uniform over all legal graphs, then the conditional distribution of one node given all the others is uniform over its legal colors. Thus, at each stage of the algorithm, we are drawing from the conditional distribution of one node given all the others: we're running a random scan Gibbs sampler!

Example 12.2.5 (Darwin's finches). When Charles Darwin visited the Galápagos Islands, he kept a record of the finch species he observed on each island. Table 12.1 summarizes Darwin's data, with each row corresponding to a species and each column to an island. The presence of a 1 in entry (i, j) of the table indicates that species i was observed on island j .

		Island																
Species	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	Total
1	0	0	1	1	1	1	1	1	1	1	0	1	1	1	1	1	1	14
2	1	1	1	1	1	1	1	1	1	1	0	1	0	1	1	0	0	13
3	1	1	1	1	1	1	1	1	1	1	1	1	0	1	1	0	0	14
4	0	0	1	1	1	0	0	1	0	1	0	1	1	0	1	1	1	10
5	1	1	1	0	1	1	1	1	1	1	0	1	0	1	1	0	0	12
6	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	0	0	2
7	0	0	1	1	1	1	1	1	1	0	0	1	0	1	1	0	0	10
8	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	1
9	0	0	1	1	1	1	1	1	1	1	0	1	0	0	1	0	0	10
10	0	0	1	1	1	1	1	1	1	1	0	1	0	1	1	0	0	11
11	0	0	1	1	1	0	1	1	0	1	0	0	0	0	0	0	0	6
12	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2
13	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	17
Total	4	4	11	10	10	8	9	10	8	9	3	10	4	7	9	3	3	122

TABLE 12.1

Presence of 13 finch species (rows) on 17 islands (columns). A value of 1 in entry (i, j) indicates that species i was observed on island j . Data are from Sanderson [22].

Given these data, we might be interested in knowing whether the pattern of 0's and 1's observed in the table is anomalous in some way. For example, does there appear to be dependence between the rows and columns? Do some pairs of species frequently occur together on the same islands, more often than one would expect by chance? These patterns may shed light on the dynamics of inter-species cooperation or competition. One way to test for such patterns is by looking at a lot of random tables with the same row and column sums as the observed table, to see how the observed table compares to the random ones. This is a common technique in statistics known as a goodness-of-fit test.

But how do we generate random tables with the same row and column sums as Table 12.1? The number of tables satisfying these constraints is impossible to enumerate. MCMC comes to the rescue: we'll create a Markov chain on the space of all tables with these row and column sums, whose stationary distribution is uniform over all such tables.

To construct the Markov chain, we need a way to transition from one table to another without changing the row or column sums. Starting from the observed table, randomly select two rows and two columns. If the four entries at their intersection have one of the following

two patterns:

$$\begin{array}{cc} 0 & 1 \\ 1 & 0 \end{array}$$

$$\text{or} \quad \begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array}$$

then switch to the opposite pattern with probability $1/2$; otherwise stay in place. For example, if we selected rows 1 and 3 and columns 1 and 17, we would switch the four entries at their intersection from $\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}$ to $\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}$ with probability $1/2$. This is a symmetric transition rule (for all tables t and t' , the transition probability from t to t' equals the transition probability from t' to t), the transitions never alter the row or column sums, and it can be shown that the Markov chain defined in this way is irreducible. Therefore the stationary distribution is uniform over all tables with the given row and column sums, as desired.

To interpret this procedure as a Gibbs sampler, consider conditioning on all entries in the table besides the four entries at the intersection of rows 1 and 3 and columns 1 and 17. If the stationary distribution is to be uniform over all tables with the given row and column sums, then the conditional distribution of these four entries must be uniform over all configurations that don't change the row and column sums, namely

$\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}$ and $\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}$. Thus, at each stage, we are drawing from the conditional distribution of four entries given all the rest.

As with Metropolis-Hastings, Gibbs sampling also applies to continuous distributions, replacing conditional PMFs with conditional PDFs.

Example 12.2.6 (Chicken-egg with unknown parameters). A chicken lays N eggs, where $N \sim \text{Pois}(\lambda)$. Each egg hatches with probability p , where p is unknown; we let $p \sim \text{Beta}(a, b)$. The constants λ, a, b are known.

Here's the catch: we don't get to observe N . Instead, we only observe the number of eggs that hatch, X . Describe how to use Gibbs sampling to find $E(p \mid X = x)$, the posterior mean of p after observing x hatched eggs.

Solution:

By the chicken-egg story, the distribution of X given p is $\text{Pois}(\lambda p)$. The posterior PDF of p is proportional to

$$f(p \mid X = x) \propto P(X = x \mid p)f(p) \propto e^{-\lambda p}(\lambda p)^x p^{a-1} q^{b-1}$$

where we have dropped all terms not depending on p .

This isn't a named distribution, so it might appear as though we're stuck, but we can get ourselves out of this rut by thinking conditionally. What do we wish we knew? The total number of eggs! Conditional on observing $N = n$ and knowing the true value of p , the distribution of X would be $\text{Bin}(n, p)$. By conditioning on the total number of eggs, we recover Beta-Binomial conjugacy between p and X . This allows us to write down the posterior distribution right away using Story 8.3.3:

$$p \mid X = x, N = n \sim \text{Beta}(x + a, n - x + b).$$

The fact that conditioning on N makes matters so much nicer inspires us to use Gibbs sampling to tackle the problem. We alternate between sampling from p conditional on N and sampling from N conditional on p , as described below. Throughout, we must also condition on

$X = x$, since we seek to learn about the posterior distributions of the parameters conditional on the evidence.

We make an initial guess for p and N , then iterate the following steps:

1. Conditional on $N = n$ and $X = x$, draw a new guess for p from the $\text{Beta}(x + a, n - x + b)$ distribution.
2. Conditional on p and $X = x$, the number of unhatched eggs is $Y \sim \text{Pois}(\lambda(1 - p))$ by the chicken-egg story, so we can draw Y from the $\text{Pois}(\lambda(1 - p))$ distribution and set the new guess for N to be $N = x + Y$.

After many iterations, we have draws for both p and N . If we want, we can ignore the draws of N , since N was merely a device to help us sample p . But for fun, we'll plot both: Figure 12.5 shows histograms of the posterior draws of p and N when $\lambda = 10, a = b = 1$ (corresponding to a $\text{Unif}(0, 1)$ prior on p), and we observe $X = 7$ hatched eggs.

As for the posterior mean $E(p \mid X = x)$ originally asked for in the problem, we can take the sample mean of the draws of p to get a good approximation. In this case, the sample mean is 0.68. Using the code provided in the R section of this chapter, you can try changing the values of λ, a, b , and x to see how the histograms and the posterior mean are affected. The key strategy for this problem was to add the unobserved number of eggs N to the model, so that we would have pleasant conditional distributions and could use Gibbs sampling conveniently.




FIGURE 12.5

Histograms of 10^4

\$ draws from the posterior distributions of p and N , where $\lambda = 10, a = 1, b = 1$, and we observe $X = 7$. }

4.12.3 12.3 Recap

Markov chain Monte Carlo allows us to sample from complicated distributions using Markov chains. MCMC has been applied in a very wide variety of problems in recent years. The main idea behind MCMC algorithms is to construct a Markov chain whose stationary distribution is the distribution we wish to sample from. After running the Markov chain for a long time, the values that the Markov chain takes on can serve as draws from the desired distribution.

The two MCMC algorithms discussed in this chapter are Metropolis-Hastings and Gibbs sampling. The Metropolis-Hastings algorithm uses any irreducible Markov chain on the state space to generate proposals, then accepts or rejects those proposals so as to produce a modified Markov chain with the desired stationary distribution. Moreover, the resulting chain is reversible. The choice of the proposal distribution is extremely important in practice, since a bad proposal distribution may result in very slow convergence to the stationary distribution.

Gibbs sampling is a method for drawing from a d -dimensional joint distribution by updating the components of a d -dimensional Markov chain one at a time, conditional on all other components. This can be done through a systematic scan, which deterministically cycles through the components in a fixed order, or a random scan, which randomly chooses which component to update at each stage.

4.12.4 12.4 R

Metropolis-hastings

Here's how to implement the Metropolis-Hastings algorithm for Example 12.1.8, the Normal-Normal model. First, we choose our observed value of Y and decide on values for the constants σ , μ , and τ :

```
y <- 3
sigma <- 1
mu <- 0
tau <- 2
```

We also need to choose the standard deviation of the proposals for step 1 of the algorithm, as explained in Example 12.1.8; for this problem, we let $d = 1$. We set the number of iterations to run, and we allocate a vector `theta` of length 10^4 which we will fill with our simulated draws:

```
d <- 1
niter <- 10^4
theta <- rep(0,niter)
```

Now for the main loop. We initialize θ to the observed value y , then run the algorithm described in Example 12.1.8:

```
theta[1] <- y
for (i in 2:niter){
  theta.p <- theta[i-1] + rnorm(1,0,d)
  r <- dnorm(y,theta.p,sigma) * dnorm(theta.p,mu,tau) /
      (dnorm(y,theta[i-1],sigma) * dnorm(theta[i-1],mu,tau))
  flip <- rbinom(1,1,min(r,1))
  theta[i] <- if(flip==1) theta.p else theta[i-1]
}
```

Let's step through each line inside the loop. The proposed value of θ is `theta.p`, which equals the previous value of θ plus a Normal random variable with mean 0 and standard deviation `d` (recall that `rnorm` takes the standard deviation and not the variance as input). The ratio r is

$$\frac{f_{\theta|Y}(x' | y)}{f_{\theta|Y}(x | y)} = \frac{e^{-\frac{1}{2\sigma^2}(y-x')^2} e^{-\frac{1}{2\tau^2}(x'-\mu)^2}}{e^{-\frac{1}{2\sigma^2}(y-x)^2} e^{-\frac{1}{2\tau^2}(x-\mu)^2}},$$

where `theta.p` is playing the role of x' and `theta [ i - 1 ]` is playing the role of x . The coin flip to determine whether to accept or reject the proposal is `flip`, which is a coin flip with probability $\min(r, 1)$ of Heads (encoding Heads as 1 and Tails as 0). Finally, we set `theta[i]` equal to the proposed value if the coin flip lands Heads, and keep it at the previous value otherwise.

The vector `theta` now contains all of our simulation draws. We typically discard some of the initial draws to give the chain some time to approach the stationary distribution. The following

line of code discards the first half of the draws:

```
theta <- theta[-(1:(niter/2))]
```

To see what the remaining draws look like, we can create a histogram using `hist(theta)`. We can also compute summary statistics such as `mean(theta)` and `var(theta)`, which give us the sample mean and sample variance.

Gibbs

Now let's implement Gibbs sampling for Example 12.2.6, the chicken-egg story with unknown hatching probability and invisible unhatched eggs. The first step is to decide on our observed value of X , as well as the constants λ, a, b :

```
x <- 7
lambda <- 10
a <- 1
b <- 1
```

Next we decide how many iterations to run, and we allocate space for our results, creating two vectors p and N of length 10^4 which we will fill with our simulated draws:

```
niter <- 10^4
p <- rep(0,niter)
N <- rep(0,niter)
```

Finally, we're ready to run the Gibbs sampler. We initialize p and N to the values 0.5 and $2x$, respectively, and then we run the algorithm as explained in Example 12.2.6:

```
p[1] <- 0.5
N[1] <- 2*x
for (i in 2:niter){
  p[i] <- rbeta(1,x+a,N[i-1]-x+b)
  N[i] <- x + rpois(1,lambda*(1-p[i-1]))
}
```

Again, we discard the initial draws:

```
p <- p[-(1:(niter/2))]
N <- N[-(1:(niter/2))]
```

To see what the remaining draws look like, we can make histograms using `hist(p)` and `hist(N)`, which is how we created Figure 12.5. We can also compute summary statistics such as `mean(p)` or `median(p)`.

4.12.5 12.5 Exercises

1. Let $p(x, y)$ be the joint PMF of two discrete r.v.s X and Y . Using shorthand notation as we used with the Gibbs sampler, let $p(x)$ and $p(y)$ be the marginal PMFs of X and Y , and $p(x | y)$ and $p(y | x)$ be the conditional PMFs of X given Y and Y given X . Suppose the support of Y is the same as the support of the conditional distribution of $Y | X$.
 - (a) Use the identity $p(x)p(y | x) = p(y)p(x | y)$ to find an expression for the marginal PMF $p(y)$ in terms of the conditional PMFs $p(x | y)$ and $p(y | x)$.
Hint: Rewrite the identity as $p(x)/p(y) = p(x | y)/p(y | x)$ and take a sum.
 - (b) Explain why the two conditional distributions $p(x | y)$ and $p(y | x)$ determine the joint distribution $p(x, y)$, and how this fact relates to the Gibbs sampler.

2. We have a network G with n nodes and some edges. Each node of G can either be vacant or occupied. We want to place particles on the nodes of G in such a way that the particles are not too crowded. Thus, define a feasible configuration as a placement of particles such that each node is occupied by at most one particle, and no neighbor of an occupied node is occupied.

Construct a Markov chain whose stationary distribution is uniform over all feasible configurations. Clearly specify the transition rule of your Markov chain, and explain why its stationary distribution is uniform.

3. This problem considers an application of MCMC techniques to image analysis. Imagine a 2D image consisting of an $L \times L$ grid of black-or-white pixels. Let Y_j be the indicator of the j th pixel being white, for $j = 1, \dots, L^2$. Viewing the pixels as nodes in a network, the neighbors of a pixel are the pixels immediately above, below, to the left, and to the right (except for boundary cases).

Let $i \sim j$ stand for " i and j are neighbors". A commonly used model for the joint PMF of $\mathbf{Y} = (Y_1, \dots, Y_{L^2})$ is

$$P(\mathbf{Y} = \mathbf{y}) \propto \exp \left(\beta \sum_{(i,j): i \sim j} I(y_i = y_j) \right)$$

If β is positive, this says that neighboring pixels prefer to have the same color. The normalizing constant of this joint PMF is a sum over all 2^{L^2} possible configurations, so it may be very computationally difficult to obtain. This motivates the use of MCMC to simulate from the model.

(a) Suppose that we wish to simulate random draws from the joint PMF of $\mathbf{Y}$, for a particular known value of β . Explain how we can do this using Gibbs sampling, cycling through the pixels one by one in a fixed order.

(b) Now provide a Metropolis-Hastings algorithm for this problem, based on a proposal of picking a uniformly random site and toggling its value.

4.13 13 Poisson Processes

Poisson processes serve as a simple model for events occurring in time or space: in one dimension, cars passing by a highway checkpoint; in two dimensions, flowers in a meadow; in three dimensions, stars in a region of the galaxy. Poisson processes are a primary building block for more complicated processes in time and space, which are the focus of a branch of statistics called spatial statistics.

Poisson processes are also useful in probability, as they tie together many of the named distributions and provide us with insightful story proofs for results that might otherwise be tedious to show. This leads to a new problem-solving strategy, which we will demonstrate: even when a problem makes no mention of Poisson processes, there are sometimes ways to pretend the r.v.s are coming from a Poisson process so that we can use convenient Poisson process properties.

In this chapter, we'll review the already familiar definition of the 1D Poisson process, derive and discuss three important properties of the 1D Poisson process, and then extend these properties to Poisson processes in higher dimensions.

4.13.1 13.1 Poisson Processes in one Dimension

In Section 5.6, we defined a one-dimensional Poisson process and showed that the interarrival times are i.i.d. Exponentials. In Chapter 8 we showed that the time of the j th arrival (relative

to some fixed starting time) is Gamma. Let's review these results, with notation that will help us generalize to higher dimensions.

Definition 13.1.1 (1D Poisson process). A sequence of arrivals in continuous time is a Poisson process with rate λ if the following conditions hold:

1. The number of arrivals in an interval of length t is distributed $\text{Pois}(\lambda t)$.
2. The numbers of arrivals in disjoint time intervals are independent.

Usually we will assume that the timeline starts at $t = 0$, in which case we have a Poisson process on $(0, \infty)$, but we can use the same conditions to define a Poisson process on $(-\infty, \infty)$ if we want the timeline to be infinite in both directions.

Consider a Poisson process on $(0, \infty)$. Departing slightly from our notation in previous chapters, let $N(t)$ be the number of arrivals in $(0, t]$. Then the number of arrivals in $(t_1, t_2]$ is $N(t_2) - N(t_1)$, for $0 < t_1 < t_2$. Let T_j be the time of the j th arrival. Since $T_1 > t$ is the same event as $N(t) = 0$ (by the count-time duality, as discussed in Section 5.6),

$$P(T_1 > t) = P(N(t) = 0) = e^{-\lambda t}$$

so $T_1 \sim \text{Expo}(\lambda)$. Next let's condition on the first arrival time T_1 and look at the additional time $T_2 - T_1$ until the second arrival. Then

$$(T_2 - T_1) \mid T_1 \sim \text{Expo}(\lambda)$$

by the same argument as above, since we have a fresh Poisson process starting at T_1 . Since the conditional distribution of $(T_2 - T_1) \mid T_1$ does not depend on T_1 , we have that $T_2 - T_1$ is independent of T_1 , so $T_2 - T_1 \sim \text{Expo}(\lambda)$ unconditionally too.

Continuing in this way, the interarrival times $T_j - T_{j-1}$ are independent, with

$$T_j - T_{j-1} \sim \text{Expo}(\lambda)$$

So a Poisson process can be described dually as a process in which arrival counts are Poissons or a process in which interarrival times are Exponentials. Also note that since T_j is the sum of j i.i.d. $\text{Expo}(\lambda)$ r.v.s,

$$T_j \sim \text{Gamma}(j, \lambda).$$

The connection with the Exponential gives us a simple way to generate n arrivals from a Poisson process.

Story 13.1.2 (Generative story for 1D Poisson process). To generate n arrivals from a Poisson process on $(0, \infty)$ with rate λ :

1. Generate n i.i.d. $\text{Expo}(\lambda)$ r.v.s $X_1, \dots, X_n$.
2. For $j = 1, \dots, n$, set $T_j = X_1 + \dots + X_j$.

Then we can take $T_1, \dots, T_n$ to be the arrival times.

Figure 13.1 depicts three realizations of Poisson processes with rates 1, 2, and 5, respectively, plotted up to time 10. In all three cases, we can see that despite the interarrival times being i.i.d., the arrivals are not evenly spaced. Rather, there is a lot of variability in the interarrival times, which produces clumps of arrivals. This phenomenon is known as Poisson clumping. It might seem like an amazing coincidence to observe a cluster of several arrivals that are close together in time, but Poisson clumping says that having such clusters is common with Poisson processes.

$\lambda = 1$


 $\lambda = 2$

 $\lambda = 5$


FIGURE 13.1

Simulated Poisson process in one dimension, for $\lambda = 1, 2, 5$. The arrivals are far from evenly spaced from each other, and in fact they sometimes clump together,

4.13.2 13.2 Conditioning, Superposition, and Thinning

The three most important properties to understand about the Poisson process are conditioning, superposition, and thinning. These correspond to properties we've already seen about the Poisson distribution, so they should already be plausible.

13.2.1 Conditioning

What happens when we take a Poisson process and condition on the total number of events in an interval? Our first result is that conditional on the total number of events in an interval, the number of events in a fixed subinterval is Binomial. This follows from Theorem 4.8.2, where we showed that we can get from the Poisson to the Binomial by conditioning.

Theorem 13.2.1 (Conditional counts). Let $(N(t) : t > 0)$ be a Poisson process with rate λ , and $t_1 < t_2$. The conditional distribution of $N(t_1)$ given $N(t_2) = n$ is

$$N(t_1) \mid N(t_2) = n \sim \text{Bin} \left(n, \frac{t_1}{t_2} \right).$$

Proof. Figure 13.2 illustrates the setup. The claim is that conditional on having a total of n arrivals in $(0, t_2]$, the number of arrivals in $(0, t_1]$ is Binomial, with n trials and success probability proportional to t_1 .

Since $(0, t_1]$ and $(t_1, t_2]$ are disjoint, $N(t_1)$ is independent of $N(t_2) - N(t_1)$. The first is distributed $\text{Pois}(\lambda t_1)$, the second is distributed $\text{Pois}(\lambda(t_2 - t_1))$, and their sum is




FIGURE 13.2

Conditioning. Given that n arrivals occurred in $(0, t_2]$, the conditional distribution of the number of arrivals in $(0, t_1]$ is Binomial with parameters n and t_1/t_2 . the number of arrivals in $(0, t_2]$, namely $N(t_2)$. By Theorem 4.8.2,

$$N(t_1) \mid N(t_2) = n \sim \text{Bin} \left(n, \frac{\lambda t_1}{\lambda t_1 + \lambda(t_2 - t_1)} \right),$$

which is precisely what we wanted to show. $\square$

Carrying the idea of conditioning on the number of arrivals further, we have the following striking result: in a Poisson process, given that $N(t) = n$, the arrival times are distributed as if we threw down n i.i.d. $\text{Unif}(0, t)$ points.

First let's look at a simple case, where there has been only one arrival.

Proposition 13.2.2. In a Poisson process of rate λ , conditional on $N(t) = 1$, the first arrival time T_1 has the $\text{Unif}(0, t)$ distribution.

Proof. Let $0 < s < t$. By a form of the count-time duality,

$$\begin{aligned} P(T_1 \leq s \mid N(t) = 1) &= \frac{P(T_1 \leq s, N(t) = 1)}{P(N(t) = 1)} \\ &= \frac{P(N(s) = 1, N(t) - N(s) = 0)}{P(N(t) = 1)} \\ &= \frac{(e^{-\lambda s} \lambda s) (e^{-\lambda(t-s)})}{e^{-\lambda t} \lambda t} \\ &= \frac{s}{t} \end{aligned}$$

Thus, the conditional CDF of T_1 given $N(t) = 1$ is the $\text{Unif}(0, t)$ CDF. $\square$

More generally, given that $N(t) = n$, the arrival times T_j look like the order statistics of n i.i.d. $\text{Unif}(0, t)$ r.v.s. We omit the proof since in introducing order statistics in Chapter 8, we focused on their marginal distributions.

Theorem 13.2.3 (Conditional times). In a Poisson process of rate λ , conditional on $N(t) = n$, the joint distribution of the arrival times $T_1, \dots, T_n$ is the same as the joint distribution of the order statistics of n i.i.d. $\text{Unif}(0, t)$ r.v.s.

From Chapter 8, we know that the order statistics of $\text{Unif}(0, 1)$ r.v.s are Betas, so the conditional distributions of the T_j are scaled Betas; to get Beta distributions, we can just divide the T_j by t so that their support is $(0, 1)$:

$$t^{-1}T_j \mid N(t) = n \sim \text{Beta}(j, n - j + 1).$$

We now have another way to generate arrivals from a Poisson process, this time for a specific interval instead of a specific number of arrivals.

Story 13.2.4 (Generative story for Poisson process, take 2). To generate arrivals from a Poisson process with rate λ in an interval $(0, t]$:

1. Generate the total number of events in the interval, $N(t) \sim \text{Pois}(\lambda t)$.
2. Given $N(t) = n$, generate n i.i.d. $\text{Unif}(0, t)$ r.v.s $U_1, \dots, U_n$.
3. For $j = 1, \dots, n$, set $T_j = U_{(j)}$.

This is actually the generative story we used to create Figure 13.1, since we knew we wanted to simulate in the interval $(0, 10]$.

Example 13.2.5 (Users on a website). Users visit a certain website according to a Poisson process with rate λ_1 users per minute, where an "arrival" at a certain time means that at that time someone starts browsing the site. After arriving at the site, each user browses the site for an $\text{Expo}(\lambda_2)$ amount of time (and then leaves), independent of other users.

Suppose that at time 0, no one is using the site. Let N_t be the number of users who arrive in the interval $(0, t]$, and let C_t be the number of users who are currently browsing the site at time t .

- (a) Let X be the time of arrival and Y be the time of departure for a user who arrives at a Uniform time point in $[0, t]$. Find the joint PDF of X and Y .
- (b) Let p_t be the probability that a user who arrives at a Uniform time point in $(0, t]$ is still browsing the site at time t . Find p_t .
- (c) Find the distribution of C_t in terms of λ_1, λ_2 , and t .
- (d) Little's law is a very general result, which says the following:

The long-run average number of customers in a stable system is the long-term average arrival rate multiplied by the average time a customer spends in the system.

Explain what happens to $E(C_t)$ for t large, and how this can be interpreted in terms of Little's law.

Solution:

(a) We have $X \sim \text{Unif}(0, t)$. Given $X = x$, Y is an $\text{Expo}(\lambda_2)$ shifted to start at x , i.e., $(Y - x) \mid (X = x) \sim \text{Expo}(\lambda_2)$. So the joint PDF of X and Y is

$$f(x, y) = f_X(x)f_{Y|X}(y \mid x) = \frac{\lambda_2}{t}e^{-\lambda_2(y-x)}, \text{ for } 0 < x < t \text{ and } x < y.$$

(b) With notation as in (a), we want to find $p_t = P(Y > t)$. This can be done by integrating the joint PDF over all (x, y) with $y > t$:

$$\begin{aligned} P(Y > t) &= \frac{1}{t} \int_0^t \int_t^\infty \lambda_2 e^{-\lambda_2(y-x)} dy dx \\ &= \frac{1}{t} \int_0^t e^{\lambda_2 x} \left(\int_t^\infty \lambda_2 e^{-\lambda_2 y} dy \right) dx \\ &= \frac{e^{-\lambda_2 t}}{t} \int_0^t e^{\lambda_2 x} dx \\ &= \frac{1 - e^{-\lambda_2 t}}{\lambda_2 t} \end{aligned}$$

(c) By Theorem 13.2.3, given $N_t = n$ we have that the n arrival times in $(0, t]$ are i.i.d. and uniform in that interval. Therefore, $C_t \mid N_t \sim \text{Bin}(N_t, p_t)$, with p_t as above, and $N_t \sim \text{Pois}(\lambda_1 t)$. So by the chicken-egg story (Story 7.1.9),

$$C_t \sim \text{Pois}(\lambda_1 p_t t)$$

That is,

$$C_t \sim \text{Pois}\left(\frac{\lambda_1 (1 - e^{-\lambda_2 t})}{\lambda_2}\right).$$

(d) As $t \rightarrow \infty$, $E(C_t) \rightarrow \lambda_1/\lambda_2$. This agrees with Little's law since it says the longrun average number of users in the system (browsing the site) is the rate at which users arrive (λ_1) times the average time a user browses in a session ($1/\lambda_2$).

13.2.2 Superposition

The next property we will examine is superposition: if we take two independent Poisson processes and overlay them, we get another Poisson process. (The combined process is a Poisson process in its own right, but we can imagine that the arrivals are tagged to indicate which of the underlying processes they came from.)

Theorem 13.2.6 (Superposition). Let $(N_1(t) : t > 0)$ and $(N_2(t) : t > 0)$ be independent Poisson processes with rates λ_1 and λ_2 , respectively. Then the combined process $N(t) = N_1(t) + N_2(t)$ is a Poisson process with rate $\lambda_1 + \lambda_2$.

Proof. Let's verify the two properties in the definition of Poisson process.

1. For all $t > 0$, $N_1(t) \sim \text{Pois}(\lambda_1 t)$ and $N_2(t) \sim \text{Pois}(\lambda_2 t)$, independently, so $N(t) \sim \text{Pois}((\lambda_1 + \lambda_2)t)$, by Theorem 4.8.1. The same argument applies for any interval of length t , not just intervals of the form $(0, t]$.

2. Arrivals in disjoint intervals are independent in the combined process because they are independent in the two individual processes, and the individual processes are independent of each other.

There is a very natural way to generate a superposition.

Story 13.2.7 (Generative story for superposition). To generate the superposition of two independent Poisson processes, $(N_1(t) : t > 0)$ with rate λ_1 and $(N_2(t) : t > 0)$ with rate λ_2 :

1. Generate arrivals from the Poisson process $(N_1(t) : t > 0)$.
2. Generate arrivals from the Poisson process $(N_2(t) : t > 0)$.
3. Superpose the results of steps 1 and 2.




FIGURE 13.3

Superposition. The superposition of independent Poisson processes is a Poisson process, and the rates add. The top two timelines are independent Poisson processes, each with rate 0.5 . The bottom timeline is the superposition of the top two Poisson processes and is itself a Poisson process with rate 1.

Figure 13.3 depicts a superposed Poisson process consisting of x's and ◊'s. Let's call the ×'s "type- 1 events" and the ◊'s "type- 2 events". A natural question to ask is: what is the probability of observing a type- 1 event before a type- 2 event?

Theorem 13.2.8 (Probability of type-1 event before type-2 event). Consider two independent Poisson processes: a Poisson process of type-1 arrivals, with rate λ_1 , and a Poisson process of type- 2 arrivals, with rate λ_2 . In the superposition of these two processes, the probability of the first arrival being type- 1 is $\lambda_1 / (\lambda_1 + \lambda_2)$.

Proof. Let T be the time until the first type-1 event and let V be the time until the first type-2 event. We seek $P(T \leq V)$. We could do this with 2D LOTUS, integrating the joint PDF of T and V over the region of interest in the 2D plane. But it turns out we can avoid calculus altogether.

We have $T \sim \text{Expo}(\lambda_1)$ and $V \sim \text{Expo}(\lambda_2)$, independently. Rescaling, let

$$\tilde{T} = \lambda_1 T, \tilde{V} = \lambda_2 V$$

Then $\tilde{T}, \tilde{V}$ are i.i.d. $\text{Expo}(1)$.

Letting $U = \tilde{T} / (\tilde{T} + \tilde{V})$, we have

$$\begin{aligned}
 P(T \leq V) &= P\left(\frac{\tilde{T}}{\lambda_1} \leq \frac{\tilde{V}}{\lambda_2}\right) \\
 &= P\left(\frac{\tilde{T}}{\tilde{T} + \tilde{V}} \leq \frac{\tilde{V}}{\tilde{T} + \tilde{V}} \cdot \frac{\lambda_1}{\lambda_2}\right) \\
 &= P\left(U \leq (1 - U) \cdot \frac{\lambda_1}{\lambda_2}\right) \\
 &= P\left(U \leq \frac{\lambda_1}{\lambda_1 + \lambda_2}\right)
 \end{aligned}$$

Since $\tilde{T}, \tilde{V} \sim \text{Expo}(1)$, the bank-post office story tells us that $U \sim \text{Beta}(1, 1)$. In other words, U is standard Uniform! Thus, $P(T \leq V) = \lambda_1 / (\lambda_1 + \lambda_2)$. Note that when $\lambda_1 = \lambda_2$ this reduces to $1/2$, as it should by symmetry. $\square$

The above result applies to the first arrival in the combined Poisson process. After the first arrival, however, the same reasoning applies to the second arrival: by the memoryless property, the time to the next type-1 event is $\text{Expo}(\lambda_1)$ and the time to the next type-2 event is $\text{Expo}(\lambda_2)$, independent of the past. Therefore the second arrival is a type-1 arrival with probability $\lambda_1 / (\lambda_1 + \lambda_2)$, independent of the first arrival. Similarly, all of the arrival types can be viewed as i.i.d. coin tosses with probability $\lambda_1 / (\lambda_1 + \lambda_2)$ of Heads, where Heads corresponds to type-1.

This yields an alternative generative story for the superposition of two independent Poisson processes: we can first generate an $\text{Expo}(\lambda_1 + \lambda_2)$ r.v. to decide when the next arrival occurs, and then independently flip a coin with probability $\lambda_1 / (\lambda_1 + \lambda_2)$ of Heads to decide what kind of arrival it is.

Story 13.2.9 (Generative story for superposition, take 2). To generate the superposition of two independent Poisson processes, with rates λ_1 and λ_2 :

1. Generate i.i.d. $\text{Expo}(\lambda_1 + \lambda_2)$ r.v.s $X_1, X_2, \dots$, and let the j th arrival be at time $T_j = X_1 + \dots + X_j$.
2. Generate i.i.d. r.v.s $I_1, I_2, \dots \sim \text{Bern}(\lambda_1 / (\lambda_1 + \lambda_2))$, independent of $X_1, X_2, \dots$. Let the j th arrival be type-1 if $I_j = 1$, and type-2 otherwise. $\square$

This story provides us with a quick proof of a result known as the competing risks theorem, which seems like a surprising independence result when stated on its own but becomes very intuitive when viewed in the context of Poisson processes.

Example 13.2.10 (Competing risks). The lifetime of Fred's refrigerator is $Y_1 \sim \text{Expo}(\lambda_1)$, and the lifetime of his dishwasher is $Y_2 \sim \text{Expo}(\lambda_2)$, independent of Y_1 . Show that $\min(Y_1, Y_2)$, the time of the first appliance failure, is independent of $I(Y_1 < Y_2)$, the indicator that the refrigerator failed first. This may seem surprising.

For example, if the average lifetime of his refrigerator is 15 years and the average lifetime of his dishwasher is 7 years, and then one of these appliances fails after 7 years, it would be natural to guess that it was the dishwasher that failed. But in fact, we will show that knowing the time of the first appliance failure provides no information about which appliance failed.

Solution: We will use an embedding strategy, using Poisson process ideas even though this problem doesn't mention Poisson processes anywhere! We will embed the r.v.s Y_1 and Y_2 into a Poisson process that we ourselves invent, in order to take advantage of the properties of Poisson processes. So let's imagine there is an entire Poisson process of refrigerator failures with rate λ_1 and a Poisson process of dishwasher failures with rate λ_2 .

We can interpret Y_1 as the waiting time for the first arrival in the refrigerator process and Y_2 as the waiting time for the first arrival in the dishwasher process. This approach is valid since $(\min(Y_1, Y_2), I(Y_1 < Y_2))$ is a function of (Y_1, Y_2) , so what matters is the joint distribution of (Y_1, Y_2) , and the way we construct (Y_1, Y_2) does have the correct joint distribution.

Furthermore, $\min(Y_1, Y_2)$ is the waiting time for the first arrival in the superposition of the two Poisson processes, and $I(Y_1 < Y_2)$ is the indicator of this arrival being a type-1 event. But in the above generative story, the arrival times and event types in the superposed Poisson process are generated completely independently! Hence, $\min(Y_1, Y_2)$ and $I(Y_1 < Y_2)$ are independent, with $\min(Y_1, Y_2) \sim \text{Expo}(\lambda_1 + \lambda_2)$ and $I(Y_1 < Y_2) \sim \text{Bern}(\lambda_1/(\lambda_1 + \lambda_2))$.

A direct consequence of Story 13.2.9 is that if we project a superposed Poisson process into discrete time, keeping the sequence of type-1 and type-2 arrivals but not the arrival times themselves, we are left with i.i.d. $\text{Bern}(\lambda_1/(\lambda_1 + \lambda_2))$ r.v.s $I_1, I_2, \dots$, where I_j is the indicator of the j th arrival being type-1. Figure 13.4 illustrates what it means to remove the continuous-time information from the Poisson process, and Theorem 13.2.11 states the result formally.




FIGURE 13.4

Projection into discrete time. Stripping out the continuous-time information from a superposed Poisson process produces a sequence of i.i.d. indicators of type-1 versus type-2 events. Here $\times$ represents type-1 events and $\diamond$ represents type-2 events.

Theorem 13.2.11 (Projection of superposition into discrete time). Consider the superposition $(N(t) : t > 0)$ of two independent Poisson processes with rates λ_1 and λ_2 . For $j = 1, 2, \dots$, let I_j be the indicator of the j th event being from the Poisson process with rate λ_1 . Then the I_j are i.i.d. $\text{Bern}(\lambda_1/(\lambda_1 + \lambda_2))$.

Using this result, we can prove with a story that a Gamma mixture of Poissons is Negative Binomial, which we first learned from Fred's adventures in Blotchville (Story 8.4.5). We'll consider a special case first.

Theorem 13.2.12 (Exponential mixture of Poissons is Geometric). Suppose that $X \sim \text{Expo}(\lambda)$ and $Y \mid X = x \sim \text{Pois}(x)$. Then $Y \sim \text{Geom}(\lambda/(\lambda + 1))$.

Proof. As with the competing risks theorem, we embed X and Y into Poisson processes. Consider two independent Poisson processes, a process of failures arriving at rate 1 and another of successes arriving at rate λ . Let X be the time of the first success; then $X \sim \text{Expo}(\lambda)$. Let Y be the number of failures before the time of the first success. By the definition of a Poisson process with rate 1, $Y \mid X = x \sim \text{Pois}(x)$. Therefore X and Y satisfy the conditions of the theorem.

To get the marginal distribution of Y , strip out the continuous-time information! In discrete time we have i.i.d. Bernoulli trials with success probability $\lambda/(\lambda + 1)$, and Y is defined as the number of failures before the first success, so by the story of the Geometric distribution, $Y \sim \text{Geom}(\lambda/(\lambda + 1))$. $\square$

The reasoning for the general case is analogous.

Theorem 13.2.13 (Gamma mixture of Poissons is Negative Binomial). Suppose that $X \sim \text{Gamma}(r, \lambda)$ and $Y \mid X = x \sim \text{Pois}(x)$. Then $Y \sim \text{NBin}(r, \lambda/(\lambda + 1))$.

Proof. Consider two independent Poisson processes, a process of failures arriving at rate 1 and another of successes arriving at rate λ . Let X be the time of the r th success, so $X \sim \text{Gamma}(r, \lambda)$. Let Y be the number of failures before the time of the r th success. Then $Y \mid X = x \sim \text{Pois}(x)$ by definition of Poisson process. We have that Y is the number of

failures before the r th success in a sequence of i.i.d. Bernoulli trials with success probability $\lambda/(\lambda + 1)$, so $Y \sim \text{NBin}(r, \lambda/(\lambda + 1))$. $\square$

13.2.3 Thinning

The last property of Poisson processes we will discuss is thinning: if we take a Poisson process and, for each arrival, independently flip a coin to decide whether it is a type-1 event or type-2 event, we end up with two independent Poisson processes.

Theorem 13.2.14 (Thinning). Let $(N(t) : t > 0)$ be a Poisson process with rate λ , and classify each arrival as a type-1 event with probability p and a type-2 event with probability $1 - p$, where these classifications are independent of each other and independent of the arrival times. Then the type-1 events form a Poisson process with rate λp , the type-2 events form a Poisson process with rate $\lambda(1 - p)$, and these two processes are independent.

Proof. Let $\lambda_1 = \lambda p$ and $\lambda_2 = \lambda(1 - p)$, so $\lambda = \lambda_1 + \lambda_2$ and $p = \lambda_1/(\lambda_1 + \lambda_2)$. Then we find ourselves in Story 13.2.9. But Story 13.2.7 and Story 13.2.9 are two equivalent generative stories for superposition! The probabilistic structures of the outputs of the two stories are identical, thinking of the arrivals in Story 13.2.7 from $(N_i(t) : t > 0)$ as being type- i . So we can assume instead that Story 13.2.7 was used to generate the process. Happily, in this story we know from the beginning that the type-1 process is independent of the type-2 process. $\square$

Thus we can superpose independent Poisson processes to get a combined Poisson process, or we can split a single Poisson process into independent Poisson processes. Figure 13.5 is an illustration of thinning. We simply flipped Figure 13.3 upside-down, which is appropriate because thinning is the flip side of superposition!




FIGURE 13.5

Thinning. Starting with a single Poisson process, if we let each arrival be type-1 with probability p and type-2 with probability $1 - p$, we obtain two independent Poisson processes with rates λp and $\lambda(1 - p)$. Here $\times$ represents type-1 events and $\diamond$ represents type-2 events, and we take $p = 0.5$.

Thinning is the Poisson process analog of the chicken-egg story. Suppose that we thin a Poisson process $(N(t) : t > 0)$ of rate λ as in the statement of Theorem 13.2.14, and let $(N_i(t) : t > 0)$ be the process of type- i arrivals. Thinking of type-1 arrivals as eggs that hatch, it follows immediately from the chicken-egg story that we have $N_1(t) \sim \text{Pois}(\lambda p t)$, $N_2(t) \sim \text{Pois}(\lambda(1 - p)t)$, with $N_1(t)$ independent of $N_2(t)$.

For practice and as an example of the elegance of thinning, let's find the distribution of a random sum in two different ways: using MGFs and conditional expectation, and using thinning.

Example 13.2.15 (First Success sum of Exponentials). Let $X_1, X_2, \dots$ be i.i.d. $\text{Expo}(\lambda)$ and $N \sim \text{FS}(p)$, independent of the X_j . Find the distribution of the random sum

$$Y = \sum_{j=1}^N X_j.$$

Solution:

We'll solve this problem twice, first with tools from Chapter 9 and then with a Poisson process story that uses thinning. For the first method, we recognize Y as the sum of a random number of random variables, so we can find the MGF of

Y using Adam's law, conditioning on N . Recalling that the MGF of the $\text{Expo}(\lambda)$ distribution is $\lambda/(\lambda - t)$ for $t < \lambda$, we have

$$\begin{aligned} E(e^{tY}) &= E\left(E\left(e^{t\sum_{j=1}^N X_j} \mid N\right)\right) \\ &= E\left(E(e^{tX_1}) E(e^{tX_2}) \dots E(e^{tX_N}) \mid N\right) \\ &= E\left(E(e^{tX_1})^N\right) \\ &= E\left(\left(\frac{\lambda}{\lambda - t}\right)^N\right). \end{aligned}$$

Now we can use LOTUS with the $\text{FS}(p)$ PMF,

$$P(N = k) = q^{k-1}p \text{ for } k = 1, 2, \dots$$

The LOTUS sum is

$$E\left(\left(\frac{\lambda}{\lambda - t}\right)^N\right) = \sum_{k=1}^{\infty} \left(\frac{\lambda}{\lambda - t}\right)^k q^{k-1}p$$

which simplifies to $\lambda p/(\lambda p - t)$ for $t < \lambda p$, after doing some algebra and summing a geometric series. This is the $\text{Expo}(\lambda p)$ MGF, so $Y \sim \text{Expo}(\lambda p)$.

Now let's see how Poisson processes can spare us from algebra while also providing insight into why Y is Exponential. Using the embedding strategy, since the X_j are i.i.d. $\text{Expo}(\lambda)$ we are free to interpret the X_j as interarrival times in a Poisson process with rate λ . So let's imagine such a Poisson process, and let's further imagine that each of the arrivals is a special arrival with probability p , independently. Then we can interpret N as the number of arrivals until the first special arrival and Y as the waiting time for the first special arrival. But by the thinning property, the special arrivals form a Poisson process with rate λp . The waiting time for the first special arrival is thus distributed $\text{Expo}(\lambda p)$.

Thinning works analogously with more than 2 types. In this setting, thinning is sometimes called the coloring theorem.

Theorem 13.2.16 (Coloring). Let $(N(t) : t > 0)$ be a Poisson process with rate λ , and C be a finite set of "colors", labeled from 1 through c . Suppose that each arrival gets randomly assigned a color from C , with color i having probability p_i . The color assignments are independent of each other and independent of the arrival times. Let $(N_i(t) : t > 0)$ be the color i process, i.e., $N_i(t)$ is the number of arrivals with color i in $(0, t]$. Then $(N_i(t) : t > 0)$ is a Poisson process with rate λp_i , for $i = 1, 2, \dots, c$, and these c monochromatic processes are independent.

Proof. We will induct on c . For $c = 1$ there is nothing to show. For $c = 2$ the result is the thinning theorem that we already proved. Now assume that the result holds for c colors, and show it when there are $c + 1$ colors. Let's call color 1 green. Use thinning to assign each arrival to be green with probability p_1 and non-green with probability $1 - p_1$. This splits the Poisson process into 2 independent processes: a green process of rate λp_1 and a non-green process of rate $\lambda(1 - p_1)$. By the inductive hypothesis, we can then split the non-green process into c independent processes, one for each of colors 2, 3, ..., $c + 1$, where color j now gets probability

$$\tilde{p}_j = \frac{p_j}{1 - p_1} = \frac{p_j}{p_2 + p_3 + \cdots + p_{c+1}},$$

for $j = 2, 3, \dots, c + 1$. We then have $c + 1$ independent Poisson processes, one for each color, such that the process for color j has rate λp_j for $j = 1, 2, \dots, c + 1$.

The next example shows how coloring can help us decompose a complicated Poisson process into more manageable components.

Example 13.2.17 (Cars on a highway). Suppose cars enter a one-way highway from a common entrance, following a Poisson process with rate λ . The i th car has velocity V_i and travels at this velocity forever; no time is lost when one car overtakes another car. Assume the V_i are i.i.d. discrete r.v.s whose support is a finite set of positive values. The process starts at time 0, and we'll consider the highway entrance to be at location 0.

For fixed locations a and b on the highway with $0 < a < b$, let Z_t be the number of cars located in the interval $[a, b]$ at time t . (For instance, on an interstate highway running west to east through the midwestern United States, a could be Kansas City and b could be St. Louis; then Z_t would be the number of cars on the highway that are in the state of Missouri at time t .) Figure 13.6 illustrates the setup of the problem and the definition of Z_t .

Assume that t is large enough that $t > b/V_i$ for all possible values of V_i . Show that Z_t has a Poisson distribution with mean $\lambda(b - a)E(V_i^{-1})$.

Solution:

Since the V_i are discrete with finite support, we can enumerate the set of possible velocities $v_1, \dots, v_m$ and their probabilities $p_1, \dots, p_m$. After doing so, we realize that the cars entering the highway represent m types of arrivals, each corresponding to a different velocity. This suggests breaking up our overall Poisson process into simpler ones. Let's color the Poisson process m ways according to the velocity of the cars, resulting in a Poisson process with rate λp_1 for cars of velocity v_1 , a Poisson process with rate λp_2 for cars of velocity v_2 , and so forth.

For each of these m separate Poisson processes, we can ask: within what time interval do cars from this process have to enter the highway in order to be in $[a, b]$ at time t ? This is a matter of physics, not statistics:

$$\text{distance} = \text{velocity} \cdot \text{time},$$

so a car that enters the highway at time s with velocity v will be at position $(t - s)v$




FIGURE 13.6

Cars enter a highway. Their entrance times form a Poisson process and are indicated by $\times$'s on the time axis. The i th car has velocity V_i , represented by the slope of the line emanating from the i th $\times$ symbol. We are interested in Z_t , the number of cars located in the interval $[a, b]$ at time t . Here we observe $Z_t = 3$, depicted by the three stars.

at time t . Thus, in order for the car's position to be between a and b , we require its arrival time to be between $t - b/v$ and $t - a/v$. (By our assumption that t is sufficiently large, we don't need to worry about $t - b/v$ being negative.) If the car arrives prior to time $t - b/v$, it will already have passed b by time t ; if the car arrives after time $t - a/v$, it won't have reached a by time t .

We now have the answer for each of the separate Poisson processes. Within the process where cars have velocity v_j , the number of cars arriving between $t - b/v_j$ and $t - a/v_j$, which we'll call Z_{tj} , is distributed $\text{Pois}(\lambda p_j(b - a)/v_j)$: the rate of the process is λp_j , and the length of the interval $[t - b/v_j, t - a/v_j]$ is $(b - a)/v_j$.

Since the separate processes are independent, Z_{t1} through Z_{tm} are independent Poisson r.v.s. Thus,

$$Z_t = Z_{t1} + \cdots + Z_{tm} \sim \text{Pois}\left(\lambda(b - a) \sum_{j=1}^m \frac{p_j}{v_j}\right),$$

and $\sum_{j=1}^m p_j/v_j$ is the expectation of V_i^{-1} by LOTUS. This is what we wanted. $\square$

To wrap up this section, here is a table describing the correspondences between properties of the Poisson process and properties of the Poisson distribution. In the second column, $Y_1 \sim \text{Pois}(\lambda_1)$ and $Y_2 \sim \text{Pois}(\lambda_2)$ are independent.

Poisson process	Poisson distribution
conditioning	$Y_1 \mid Y_1 + Y_2 = n \sim \text{Bin}(n, \lambda_1/(\lambda_1 + \lambda_2))$
superposition	$Y_1 + Y_2 \sim \text{Pois}(\lambda_1 + \lambda_2)$
thinning	chicken-egg story

4.13.3 13.3 Poisson Processes in Multiple Dimensions

Poisson processes in multiple dimensions are defined analogously to the 1D Poisson process: we just replace the notion of length with the notion of area or volume. For concreteness, we will now define 2D Poisson processes, after which it should also be clear by analogy how to define Poisson processes in higher dimensions.

Definition 13.3.1 (2D Poisson process). Events in the plane $\mathbb{R}^2$ are a 2D Poisson process with intensity λ if the following conditions hold:

1. The number of events in a region A is distributed $\text{Pois}(\lambda \cdot \text{area}(A))$.
2. The numbers of events in disjoint regions are independent of each other.

As one might guess, conditioning, superposition, and thinning properties apply to 2D Poisson processes. Let $N(A)$ be the number of events in a region A , and let $B \subseteq A$. Given $N(A) = n$, the conditional distribution of $N(B)$ is Binomial:

$$N(B) \mid N(A) = n \sim \text{Bin}\left(n, \frac{\text{area}(B)}{\text{area}(A)}\right).$$

Conditional on the total number of events in the larger region A , the probability of an event falling into a subregion is proportional to the area of the subregion; thus the locations of the events are conditionally Uniform, and we can generate a 2D Poisson process in A by

first generating the number of events $N(A) \sim \text{Pois}(\lambda \cdot \text{area}(A))$ and then placing the events uniformly at random in A . Figure 13.7 shows simulated 2D Poisson processes in the square $[0, 5] \times [0, 5]$ for intensities $\lambda = 1, 2, 5$.

As in the 1D case, the superposition of independent 2D Poisson processes is a 2D Poisson process, and the intensities add. We can also thin a 2D Poisson process to get independent 2D Poisson processes.

One property of 1D Poisson processes for which we haven't asserted a higherdimensional analog is the count-time duality. The next example, the last in this chapter, will lead us to a spatial analog: a count-distance duality.




FIGURE 13.7

Simulated 2D Poisson process in the square $[0, 5] \times [0, 5]$, for $\lambda = 1, 2, 5$.

Example 13.3.2 (Nearest star). Stars in a certain universe are distributed according to a 3D Poisson process with intensity λ . If you live in this universe, what is the distribution of the distance from you to the nearest star?

Solution:

In a 3D Poisson process with intensity λ , the number of events in a region of space V is Poisson with mean $\lambda \cdot \text{volume}(V)$. Let R be the distance from you to the nearest star. The key observation is that in order for the event $R > r$ to occur, there must be no stars within a sphere of radius r around you; in fact, these two events are equivalent. For any $r > 0$, let N_r be the number of events within radius r of you, so $N_r \sim \text{Pois}(\lambda \cdot \frac{4}{3}\pi r^3)$. Then we have a count-distance duality:

$$R > r \text{ is the same event as } N_r = 0,$$

Therefore,

$$P(R > r) = P(N_r = 0) = e^{-\frac{4}{3}\lambda\pi r^3},$$

which gives that the CDF of R is

$$P(R \leq r) = 1 - e^{-\frac{4}{3}\lambda\pi r^3},$$

for $r > 0$ (and 0 otherwise). This is a Weibull (see Example 6.5.5). Specifically,

$$R \sim \text{Wei}\left(\frac{4\pi\lambda}{3}, 3\right).$$

□

Poisson processes have numerous extensions, some of which are explored in the exercises. We can allow λ to vary as a function of time or space instead of remaining constant; this is called an inhomogeneous Poisson process. We can allow λ to be a random variable; this is called a Cox process. Finally, we can allow the rate to increase by λ after each successive arrival; this is called a Yule process.

4.13.4 13.4 Recap

A Poisson process in one dimension is a sequence of arrivals such that the number of arrivals in any interval is Poisson (with mean proportional to the length of the interval) and disjoint intervals have independent numbers of arrivals. Some operations that we can perform with Poisson processes are conditioning, superposition, and thinning. Conditioning on the total number of arrivals in an interval allows us to view the arrivals as independent Uniforms on the interval. Superposition and thinning are complementary, and they allow us to split and merge Poisson processes when convenient. All of these properties have analogs for higher dimensions.

Poisson processes tie together many of the named distributions we have studied in this book:

- Poisson for the arrival counts,
- Exponential and Gamma for the interarrival times and arrival times,
- Binomial for the conditional counts,
- Uniform and (scaled) Beta for the conditional arrival times,
- Geometric and Negative Binomial for the discrete waiting times for special arrivals.

Poisson processes are also especially amenable to story proofs. A problem-solving strategy we used several times in this chapter is to embed r.v.s into a Poisson process in the hopes of discovering a story proof, even when the original problem appears to be unrelated to Poisson processes.

Poisson processes unite two of the important themes of this book, named distributions and stories, in a natural way. We think it is fitting to end with a topic that weaves together the story threads from throughout the book.

4.13.5 13.5 R

1D Poisson Process

In Chapter 5, we discussed how to simulate a specified number of arrivals from a one-dimensional Poisson process by using the fact that the interarrival times are i.i.d. Exponentials. In this chapter, Story 13.2.4 tells us how to simulate a Poisson process within a specified interval $(0, L]$. We first generate the number of arrivals $N(L)$, which is distributed $\text{Pois}(\lambda L)$. Conditional on $N(L) = n$, the arrival times are distributed as the order statistics of n i.i.d. $\text{Unif}(0, L)$ r.v.s. The following code simulates arrivals from a Poisson process with rate 10 in the interval $(0, 5]$:

```
L <- 5
lambda <- 10
n <- rpois(1, lambda*L)
t <- sort(runif(n, 0, L))
```

To visualize the Poisson process we have generated, we can plot the cumulative number of arrivals $N(t)$ as a function of t :

```
plot(t, 1:n, type="s")
```

This produces a staircase plot as in Figure 13.8.




FIGURE 13.8

Number of arrivals in a Poisson process with rate 10 in the interval $(0, 5]$.

Thinning

It is straightforward to thin a Poisson process in R. The following code starts with a vector of arrival times $\mathbf{t}$ and the corresponding number of arrivals n from a Poisson process of rate λ on some interval, generated as above. For each arrival, we flip a coin with probability p of Heads; these coin tosses are stored in the vector $\mathbf{y}$. Finally, the arrivals for which the coin landed Heads are labeled as type-1; the rest are labeled as type-2. The resulting vectors of arrival times, $\mathbf{t}_1$ and $\mathbf{t}_2$, are realizations of independent Poisson processes of rates λp and $\lambda(1 - p)$, by Theorem 13.2.14.

We can carry out this procedure for $p = 0.3$ as follows:

```
p <- 0.3
y <- rbinom(n,1,p)
t1 <- t[y==1]
t2 <- t[y==0]
```

2D Poisson Process

Simulating a 2D Poisson process of rate λ on a square is nearly as easy as simulating a 1D Poisson process. In the square $(0, L] \times (0, L]$, the number of arrivals is distributed $\text{Pois}(\lambda L^2)$. Conditional on the number of arrivals, the locations of the arrivals are i.i.d. Uniform points in the square. By Example 7.1.23, the coordinates for each of these points are i.i.d. $\text{Unif}(0, L)$. So for $L = 5, \lambda = 10$, we can type:

```
L <- 5
lambda <- 10
n <- rpois(1,lambda*L^2)
x <- runif(n,0,L)
y <- runif(n,0,L)
```

4.13.6 13.6 Exercises

1. Passengers arrive at a bus stop according to a Poisson process with rate λ . The arrivals of buses are exactly t minutes apart. Show that on average, the sum of the waiting times of the riders on one of the buses is $\frac{1}{2}\lambda t^2$.

2. Earthquakes occur over time according to a Poisson process with rate λ . The j th earthquake has intensity Z_j , where the Z_j are i.i.d. with mean μ and variance σ^2 . Find the mean and variance of the cumulative intensity of all the earthquakes up to time t .
3. Alice receives phone calls according to a Poisson process with rate λ . Unfortunately she has lost her cell phone charger. The battery's remaining life is a random variable T with mean μ and variance σ^2 . Let $N(T)$ be the number of phone calls she receives before the battery dies; find $E(N(T))$, $\text{Var}(N(T))$, and $\text{Cov}(T, N(T))$.
4. Emails arrive in Bob's inbox according to a Poisson process with rate λ , measured in emails per hour; each email is work-related with probability p and personal with probability $1-p$. The amount of time it takes to answer a work-related email is a random variable with mean μ_W and variance σ_W^2 , the amount of time it takes to answer a personal email has mean μ_P and variance σ_P^2 , and the response times for different emails are independent. What is the average amount of time Bob has to spend answering all the emails that arrive in a t -hour interval? What about the variance?
5. In an endless soccer match, goals are scored according to a Poisson process with rate λ . Each goal is made by team A with probability p and team B with probability $1-p$. For $j > 1$, the j th goal is a turnaround if it is made by a different team than the $(j-1)$ st goal; for example, in the sequence AABBA..., the 3rd and 5th goals are turnarounds.
 - (a) In n goals, what is the expected number of turnarounds?
 - (b) If an A-to-B turnaround has just occurred, what is the expected time until the next B-to-A turnaround?
6. Let N_t be the number of arrivals up until time t in a Poisson process of rate λ , and let T_n be the time of the n th arrival. Consider statements of the form

$$P(N_t \lesseqgtr_1 n) = P(T_n \lesseqgtr_2 t)$$

where $\lesseqgtr_1$ and $\lesseqgtr_2$ are replaced by symbols from the list $<, \leq, \geq, >$. Which of these statements are true?

7. Claims against an insurance company follow a Poisson process with rate $\lambda > 0$. A total of N claims were received over two disjoint time periods of combined length $t = t_1 + t_2$, with t_i the length of period i .
 - (a) Given this information, derive the conditional probability distribution of N_1 , the number of claims made in the first period, given N .
 - (b) The amount paid for the i th claim is X_i , with $X_1, X_2, \dots$ i.i.d. and independent of the claims process. Let $E(X_i) = \mu$, $\text{Var}(X_i) = \sigma^2$, for $i = 1, \dots, N$. Given N , find the mean and variance of the total claims paid in period 1. That is, find these two conditional moments of the quantity

$$W_1 = \sum_{i=1}^{N_1} X_i$$

where (by convention) $W_1 = 0$ if $N_1 = 0$.

8. On a certain question-and-answer website, $N \sim \text{Pois}(\lambda_1)$ questions will be posted tomorrow, with λ_1 measured in questions/day. Given N , the post times are i.i.d. and uniformly distributed over the day (a day begins and ends at midnight). When a question is posted, it takes an $\text{Expo}(\lambda_2)$ amount of time (in days) for an answer to be posted, independent of what happens with other questions.
 - (a) Find the probability that a question posted at a uniformly random time tomorrow will not

yet have been answered by the end of that day.

(b) Find the joint distribution of how many answered and unanswered questions posted tomorrow there will be at the end of that day.

9. An inhomogeneous Poisson process in one dimension is a Poisson process whose rate, instead of being constant, is a nonnegative function $\lambda(t)$ of time. Formally, we require that the number of arrivals in the interval $[t_1, t_2)$ be Poisson-distributed with mean $\int_{t_1}^{t_2} \lambda(t) dt$ and that disjoint intervals be independent. When $\lambda(t)$ is constant, this reduces to the definition of the ordinary or homogeneous Poisson process.

(a) Show that we can generate arrivals from an inhomogeneous Poisson process in the interval $[t_1, t_2)$ using the following procedure.

1. Let $\lambda_{\max}$ be the maximum value of $\lambda(t)$ in the interval $[t_1, t_2)$. Create a 2D rectangle $[t_1, t_2) \times [0, \lambda_{\max}]$, and plot the function $\lambda(t)$ in the rectangle.
2. Generate $N \sim \text{Pois}(\lambda_{\max}(t_2 - t_1))$, and place N points uniformly at random in the rectangle.
3. For each of the N points: if the point falls below the curve $\lambda(t)$, accept it as an arrival in the process, and take its horizontal coordinate to be its arrival time. If the point falls above the curve $\lambda(t)$, reject it.

Hint: Verify that the two conditions in the definition are satisfied.

(b) Suppose we have an inhomogeneous Poisson process with rate function $\lambda(t)$. Let $N(t)$ be the number of arrivals up to time t and T_j be the time of the j th arrival. Explain why the hybrid joint PDF of $N(t)$ and $T_1, \dots, T_{N(t)}$, which constitute all the data observed up to time t , is given by

$$f(n, t_1, \dots, t_n) = \frac{e^{-\lambda_{\text{total}}} \lambda_{\text{total}}^n}{n!} \cdot n! \frac{\lambda(t_1) \dots \lambda(t_n)}{\lambda_{\text{total}}^n} = e^{-\lambda_{\text{total}}} \lambda(t_1) \dots \lambda(t_n)$$

for $0 < t_1 < t_2 < \dots < t_n$ and nonnegative integer n , where $\lambda_{\text{total}} = \int_0^t \lambda(u) du$.

10. A Cox process is a generalization of a Poisson process, where the rate λ is a random variable. That is, λ is generated according to some distribution on $(0, \infty)$ and then, given that value of λ , a Poisson process with that rate is generated.

(a) Explain intuitively why disjoint intervals in a 1D Cox process are not independent.

(b) In a 1D Cox process where $\lambda \sim \text{Gamma}(\alpha, \beta)$, find the covariance between the number of arrivals in $[0, t)$ and the number of arrivals in $[t, t + s)$.

Hint: Condition on λ .

11. In a Yule process with rate λ , the rate of arrivals increases after each new arrival, so that the time of the first arrival is $\text{Expo}(\lambda)$ and the time between the $(j - 1)$ st and j th arrivals is $\text{Expo}(j\lambda)$ for $j = 2, 3, \dots$. So interarrival times are independent but not i.i.d.

(a) Show that the superposition of two independent Yule processes with the same rate λ is a Yule process, except shifted so that the interarrival times are $\text{Expo}(2\lambda), \text{Expo}(3\lambda), \text{Expo}(4\lambda), \dots$ rather than $\text{Expo}(\lambda), \text{Expo}(2\lambda), \text{Expo}(3\lambda), \dots$.

(b) Show that if we project the process from Part (a) into discrete time, the resulting sequence of type-1 and type-2 arrivals is equivalent to the following discrete-time process:

1. Start with two balls in an urn, labeled 1 and 2.
2. Draw a ball out of the urn at random, note its number, and replace it along with another ball with the same number.
3. Repeat step 2 over and over.

4. Consider the coupon collector problem: there are n toy types, and toys are collected one by one, sampling with replacement from the set of toy types each time. We solved this problem in Chapter 4, assuming that all toy types are equally likely. Now suppose that at each stage, the j th toy type is collected with probability p_j , where the p_j are not necessarily equal. Let N be the number of toys needed until we have a full set; we wish to find $E(N)$. This problem outlines an embedding method for calculating $E(N)$.
- (a) Suppose that the toys arrive according to a Poisson process with rate 1, so that the interarrival times between toys are i.i.d. $X_j \sim \text{Expo}(1)$. For $j = 1, \dots, n$, let Y_j be the waiting time until the first toy of type j . What are the distributions of the Y_j ? Are the Y_j independent?
- (b) Explain why $T = \max(Y_1, \dots, Y_n)$, the waiting time until all toy types are collected, can also be written as $X_1 + \dots + X_N$, where the X_j are defined as in Part (a). Use this to show that $E(T) = E(N)$.
- (c) Show that $E(T)$, and hence $E(N)$, can be found by computing the integral

$$\int_0^\infty \left(1 - \prod_{j=1}^n (1 - e^{-p_j t}) \right) dt$$

Hint: Use the identity $E(T) = \int_0^\infty P(T > t) dt$, which was shown in Example 5.3.8.

5 A Comprehensive Course of Probability Theory by Klenke

5.1 1 Basic Measure Theory

In this chapter, we introduce the classes of sets that allow for a systematic treatment of events and random observations in the framework of probability theory. Furthermore, we construct measures, in particular probability measures, on such classes of sets. Finally, we define random variables as measurable maps.

5.1.1 1.1 Classes of Sets

In the following, let $\Omega \neq \emptyset$ be a nonempty set and let $\mathcal{A} \subset 2^\Omega$ (set of all subsets of Ω) be a class of subsets of Ω . Later, Ω will be interpreted as the space of elementary events and $\mathcal{A}$ will be the system of observable events. In this section, we introduce names for classes of subsets of Ω that are stable under certain set operations and we establish simple relations between such classes.

Definition 1.1 A class of sets $\mathcal{A}$ is called

- $\cap$ -closed (closed under intersections) or a π -system if $A \cap B \in \mathcal{A}$ whenever $A, B \in \mathcal{A}$,
- σ - $\cap$ -closed (closed under countable ¹ intersections) if $\bigcap_{n=1}^\infty A_n \in \mathcal{A}$ for any choice of countably many sets $A_1, A_2, \dots \in \mathcal{A}$,
- $\cup$ -closed (closed under unions) if $A \cup B \in \mathcal{A}$ whenever $A, B \in \mathcal{A}$,
- σ - $\cup$ -closed (closed under countable unions) if $\bigcup_{n=1}^\infty A_n \in \mathcal{A}$ for any choice of countably many sets $A_1, A_2, \dots \in \mathcal{A}$,
- $\setminus$ -closed (closed under differences) if $A \setminus B \in \mathcal{A}$ whenever $A, B \in \mathcal{A}$, and

- closed under complements if $A^c := \Omega \setminus A \in \mathcal{A}$ for any set $A \in \mathcal{A}$.

Definition 1.2 (σ -algebra) A class of sets $\mathcal{A} \subset 2^\Omega$ is called a σ -algebra if it fulfills the following three conditions:

- (i) $\Omega \in \mathcal{A}$.
- (ii) $\mathcal{A}$ is closed under complements.
- (iii) $\mathcal{A}$ is closed under countable unions.

Sometimes a σ -algebra is also named a σ -field. As we will see, we can define probabilities on σ -algebras in a consistent way. Hence these are the natural classes of sets to be considered as events in probability theory.

Theorem 1.3 If $\mathcal{A}$ is closed under complements, then we have the equivalences

$$\begin{aligned}\mathcal{A} \text{ is } \cap\text{-closed} &\iff \mathcal{A} \text{ is } \cup\text{-closed} , \\ \mathcal{A} \text{ is } \sigma\text{-}\cap\text{-closed} &\iff \mathcal{A} \text{ is } \sigma\text{-}\cup\text{-closed}.\end{aligned}$$

Proof The two statements are immediate consequences of de Morgan's rule (reminder: $(\bigcup A_i)^c = \bigcap A_i^c$). For example, let $\mathcal{A}$ be $\sigma\text{-}\cap\text{-closed}$ and let $A_1, A_2, \dots \in \mathcal{A}$. Hence

$$\bigcup_{n=1}^{\infty} A_n = \left(\bigcap_{n=1}^{\infty} A_n^c \right)^c \in \mathcal{A}.$$

Thus $\mathcal{A}$ is $\sigma\text{-}\cup\text{-closed}$. The other cases can be proved similarly.

Theorem 1.4 Assume that $\mathcal{A}$ is $\cap\text{-closed}$. Then the following statements hold:

- (i) $\mathcal{A}$ is $\cap\text{-closed}$.
- (ii) If in addition $\mathcal{A}$ is $\sigma\text{-}\cup\text{-closed}$, then $\mathcal{A}$ is $\sigma\text{-}\cap\text{-closed}$.
- (iii) Any countable (respectively finite) union of sets in $\mathcal{A}$ can be expressed as a countable (respectively finite) disjoint union of sets in $\mathcal{A}$.

Proof (i) Assume that $A, B \in \mathcal{A}$. Hence also $A \cap B = A \setminus (A \setminus B) \in \mathcal{A}$.

(ii) Assume that $A_1, A_2, \dots \in \mathcal{A}$. Hence

$$\bigcap_{n=1}^{\infty} A_n = \bigcap_{n=2}^{\infty} (A_1 \cap A_n) = \bigcap_{n=2}^{\infty} A_1 \setminus (A_1 \setminus A_n) = A_1 \setminus \bigcup_{n=2}^{\infty} (A_1 \setminus A_n) \in \mathcal{A}.$$

(iii) Assume that $A_1, A_2, \dots \in \mathcal{A}$. Hence a representation of $\bigcup_{n=1}^{\infty} A_n$ as a countable disjoint union of sets in $\mathcal{A}$ is

$$\bigcup_{n=1}^{\infty} A_n = A_1 \uplus (A_2 \setminus A_1) \uplus ((A_3 \setminus A_1) \setminus A_2) \uplus (((A_4 \setminus A_1) \setminus A_2) \setminus A_3) \uplus \dots$$

Remark 1.5 Sometimes the disjoint union of sets is denoted by the symbol $\uplus$. Note that this is not a new operation but only stresses the fact that the sets involved are mutually disjoint.

Definition 1.6 A class of sets $\mathcal{A} \subset 2^\Omega$ is called an algebra if the following three conditions are fulfilled:

- (i) $\Omega \in \mathcal{A}$.
- (ii) $\mathcal{A}$ is $\setminus\text{-closed}$.
- (iii) $\mathcal{A}$ is $\cup\text{-closed}$.

If $\mathcal{A}$ is an algebra, then obviously $\emptyset = \Omega \setminus \Omega$ is in $\mathcal{A}$. However, in general, this property is weaker than (i) in Definition 1.6.

Theorem 1.7 A class of sets $\mathcal{A} \subset 2^\Omega$ is an algebra if and only if the following three properties hold:

- (i) $\Omega \in \mathcal{A}$.

⁰¹ By "countable" we always mean either finite or countably infinite.

- (ii) $\mathcal{A}$ is closed under complements.
- (iii) $\mathcal{A}$ is closed under intersections.

Proof This is left as an exercise. $\square$

Definition 1.8 A class of sets $\mathcal{A} \subset 2^\Omega$ is called a ring if the following three conditions hold:

- (i) $\emptyset \in \mathcal{A}$.
- (ii) $\mathcal{A}$ is $\setminus$ -closed.
- (iii) $\mathcal{A}$ is $\cup$ -closed.

A ring is called a σ -ring if it is also σ - $\cup$ -closed.

Definition 1.9 A class of sets $\mathcal{A} \subset 2^\Omega$ is called a semiring if

- (i) $\emptyset \in \mathcal{A}$,
- (ii) for any two sets $A, B \in \mathcal{A}$ the difference set $B \setminus A$ is a finite union of mutually disjoint sets in $\mathcal{A}$,
- (iii) $\mathcal{A}$ is $\cap$ -closed.

Definition 1.10 A class of sets $\mathcal{A} \subset 2^\Omega$ is called a λ -system (or Dynkin's λ -system) if

- (i) $\Omega \in \mathcal{A}$,
- (ii) for any two sets $A, B \in \mathcal{A}$ with $A \subset B$, the difference set $B \setminus A$ is in $\mathcal{A}$, and
- (iii) $\biguplus_{n=1}^\infty A_n \in \mathcal{A}$ for any choice of countably many pairwise disjoint sets $A_1, A_2, \dots \in \mathcal{A}$.

Example 1.11

- (i) For any nonempty set Ω , the classes $\mathcal{A} = \{\emptyset, \Omega\}$ and $\mathcal{A} = 2^\Omega$ are the trivial examples of algebras, σ -algebras and λ -systems. On the other hand, $\mathcal{A} = \{\emptyset\}$ and $\mathcal{A} = 2^\Omega$ are the trivial examples of semirings, rings and σ -rings.
- (ii) Let $\Omega = \mathbb{R}$. Then $\mathcal{A} = \{A \subset \mathbb{R} : A \text{ is countable}\}$ is a σ -ring.
- (iii) $\mathcal{A} = \{(a, b] : a, b \in \mathbb{R}, a \leq b\}$ is a semiring on $\Omega = \mathbb{R}$ (but is not a ring).
- (iv) The class of finite unions of bounded intervals is a ring on $\Omega = \mathbb{R}$ (but is not an algebra).
- (v) The class of finite unions of arbitrary (also unbounded) intervals is an algebra on $\Omega = \mathbb{R}$ (but is not a σ -algebra).
- (vi) Let E be a finite nonempty set and let $\Omega := E^\mathbb{N}$ be the set of all E -valued sequences $\omega = (\omega_n)_{n \in \mathbb{N}}$. For any $\omega_1, \dots, \omega_n \in E$, let

$$[\omega_1, \dots, \omega_n] := \{\omega' \in \Omega : \omega'_i = \omega_i \text{ for all } i = 1, \dots, n\}$$

be the set of all sequences whose first n values are $\omega_1, \dots, \omega_n$. Let $\mathcal{A}_0 = \{\emptyset\}$. For $n \in \mathbb{N}$, define

$$\mathcal{A}_n := \{[\omega_1, \dots, \omega_n] : \omega_1, \dots, \omega_n \in E\}. \quad (1.1)$$

Hence $\mathcal{A} := \bigcup_{n=0}^\infty \mathcal{A}_n$ is a semiring but is not a ring (if $\#E > 1$).

- (vii) Let Ω be an arbitrary nonempty set. Then

$$\mathcal{A} := \{A \subset \Omega : A \text{ or } A^c \text{ is finite}\}$$

is an algebra. However, if $\#\Omega = \infty$, then $\mathcal{A}$ is not a σ -algebra.

- (viii) Let Ω be an arbitrary nonempty set. Then

$$\mathcal{A} := \{A \subset \Omega : A \text{ or } A^c \text{ is countable}\}$$

is a σ -algebra.

- (ix) Every σ -algebra is a λ -system.
- (x) Let $\Omega = \{1, 2, 3, 4\}$ and $\mathcal{A} = \{\emptyset, \{1, 2\}, \{1, 4\}, \{2, 3\}, \{3, 4\}, \{1, 2, 3, 4\}\}$. Hence $\mathcal{A}$ is a λ -system but is not an algebra. $\square$

Theorem 1.12 (Relations between classes of sets)

- (i) Every σ -algebra also is a λ -system, an algebra and a σ -ring.
- (ii) Every σ -ring is a ring, and every ring is a semiring.
- (iii) Every algebra is a ring. An algebra on a finite set Ω is a σ -algebra.

Proof (i) This is obvious.

- (ii) Let $\mathcal{A}$ be a ring. By Theorem 1.4, $\mathcal{A}$ is closed under intersections and is hence a semiring.
- (iii) Let $\mathcal{A}$ be an algebra. Then $\emptyset = \Omega \setminus \Omega \in \mathcal{A}$, and hence $\mathcal{A}$ is a ring. If in addition Ω is finite, then $\mathcal{A}$ is finite. Hence any countable union of sets in $\mathcal{A}$ is a finite union of sets. $\square$

Definition 1.13 (liminf and limsup) Let $A_1, A_2, \dots$ be subsets of Ω . The sets

$$\liminf_{n \rightarrow \infty} A_n := \bigcup_{n=1}^{\infty} \bigcap_{m=n}^{\infty} A_m \quad \text{and} \quad \limsup_{n \rightarrow \infty} A_n := \bigcap_{n=1}^{\infty} \bigcup_{m=n}^{\infty} A_m$$

are called limes inferior and limes superior, respectively, of the sequence $(A_n)_{n \in \mathbb{N}}$.

Remark 1.14

- (i) liminf and lim sup can be rewritten as

$$\begin{aligned} \liminf_{n \rightarrow \infty} A_n &= \{\omega \in \Omega : \#\{n \in \mathbb{N} : \omega \notin A_n\} < \infty\}, \\ \limsup_{n \rightarrow \infty} A_n &= \{\omega \in \Omega : \#\{n \in \mathbb{N} : \omega \in A_n\} = \infty\}. \end{aligned}$$

In other words, limes inferior is the event where eventually all of the A_n occur. On the other hand, limes superior is the event where infinitely many of the A_n occur. In particular, $A_* := \liminf_{n \rightarrow \infty} A_n \subset A^* := \limsup_{n \rightarrow \infty} A_n$.

- (ii) We define the indicator function on the set A by

$$\mathbb{1}_A(x) := \begin{cases} 1, & x \in A, \\ 0, & x \notin A. \end{cases} \quad (1.2)$$

With this notation,

$$\mathbb{1}_{A_*} = \liminf_{n \rightarrow \infty} \mathbb{1}_{A_n} \quad \text{and} \quad \mathbb{1}_{A^*} = \limsup_{n \rightarrow \infty} \mathbb{1}_{A_n}.$$

- (iii) If $\mathcal{A} \subset 2^\Omega$ is a σ -algebra and if $A_n \in \mathcal{A}$ for every $n \in \mathbb{N}$, then $A_* \in \mathcal{A}$ and $A^* \in \mathcal{A}$. $\square$

Proof This is left as an exercise. $\square$

Theorem 1.15 (Intersection of classes of sets) Let I be an arbitrary index set, and assume that $\mathcal{A}_i$ is a σ -algebra for every $i \in I$. Hence the intersection

$$\mathcal{A}_I := \{A \subset \Omega : A \in \mathcal{A}_i \text{ for every } i \in I\} = \bigcap_{i \in I} \mathcal{A}_i$$

is a σ -algebra. The analogous statement holds for rings, σ -rings, algebras and λ -systems. However, it fails for semirings.

Proof We give the proof for σ -algebras only. To this end, we check (i)-(iii) of Definition 1.2.

- (i) Clearly, $\Omega \in \mathcal{A}_i$ for every $i \in I$, and hence $\Omega \in \mathcal{A}_I$.
- (ii) Assume $A \in \mathcal{A}_I$. Hence $A \in \mathcal{A}_i$ for any $i \in I$. Thus also $A^c \in \mathcal{A}_i$ for any $i \in I$. We conclude that $A^c \in \mathcal{A}_I$.
- (iii) Assume $A_1, A_2, \dots \in \mathcal{A}_I$. Hence $A_n \in \mathcal{A}_i$ for every $n \in \mathbb{N}$ and $i \in I$. Thus $A := \bigcup_{n=1}^{\infty} A_n \in \mathcal{A}_i$ for every $i \in I$. We conclude $A \in \mathcal{A}_I$.

Counterexample for semirings: Let $\Omega = \{1, 2, 3, 4\}$, $\mathcal{A}_1 = \{\emptyset, \Omega, \{1\}, \{2, 3\}, \{4\}\}$ and $\mathcal{A}_2 = \{\emptyset, \Omega, \{1\}, \{2\}, \{3, 4\}\}$. Then $\mathcal{A}_1$ and $\mathcal{A}_2$ are semirings but $\mathcal{A}_1 \cap \mathcal{A}_2 = \{\emptyset, \Omega, \{1\}\}$ is not. $\square$

Theorem 1.16 (Generated σ -algebra) Let $\mathcal{E} \subset 2^\Omega$. Then there exists a smallest σ -algebra $\sigma(\mathcal{E})$ with $\mathcal{E} \subset \sigma(\mathcal{E})$:

$$\sigma(\mathcal{E}) := \bigcap_{\mathcal{A} \subset 2^\Omega \text{ is a } \sigma\text{-algebra}} \mathcal{A}.$$

$\sigma(\mathcal{E})$ is called the σ -algebra generated by $\mathcal{E}$. $\mathcal{E}$ is called a generator of $\sigma(\mathcal{E})$. Similarly, we define $\delta(\mathcal{E})$ as the λ -system generated by $\mathcal{E}$.

Proof $\mathcal{A} = 2^\Omega$ is a σ -algebra with $\mathcal{E} \subset \mathcal{A}$. Hence the intersection is nonempty. By Theorem 1.15, $\sigma(\mathcal{E})$ is a σ -algebra. Clearly, it is the smallest σ -algebra that contains $\mathcal{E}$. For λ -systems the proof is similar. $\square$

Remark 1.17 The following three statements hold:

- (i) $\mathcal{E} \subset \sigma(\mathcal{E})$.
- (ii) If $\mathcal{E}_1 \subset \mathcal{E}_2$, then $\sigma(\mathcal{E}_1) \subset \sigma(\mathcal{E}_2)$.
- (iii) $\mathcal{A}$ is a σ -algebra if and only if $\sigma(\mathcal{A}) = \mathcal{A}$.

The same statements hold for λ -systems. Furthermore, $\delta(\mathcal{E}) \subset \sigma(\mathcal{E})$. $\square$

Theorem 1.18 ($\cap$ -closed λ -system) Let $\mathcal{D} \subset 2^\Omega$ be a λ -system. Then

$$\mathcal{D} \text{ is a } \pi\text{-system} \iff \mathcal{D} \text{ is a } \sigma\text{-algebra}.$$

Proof " $\Leftarrow$ " This is obvious.

" $\Rightarrow$ " We check (i)-(iii) of Definition 1.2.

- (i) Clearly, $\Omega \in \mathcal{D}$.
- (ii) (Closedness under complements) Let $A \in \mathcal{D}$. Since $\Omega \in \mathcal{D}$ and by property (ii) of the λ -system, we get that $A^c = \Omega \setminus A \in \mathcal{D}$.
- (iii) (σ - $\cup$ -closedness) Let $A, B \in \mathcal{D}$. By assumption, $A \cap B \in \mathcal{D}$, and trivially $A \cap B \subset A$. Thus $A \setminus B = A \setminus (A \cap B) \in \mathcal{D}$. This implies that $\mathcal{D}$ is $\setminus$ -closed. Now let $A_1, A_2, \dots \in \mathcal{D}$. By Theorem 1.4(iii), there exist mutually disjoint sets $B_1, B_2, \dots \in \mathcal{D}$ with $\bigcup_{n=1}^\infty A_n = \biguplus_{n=1}^\infty B_n \in \mathcal{D}$. $\square$

Theorem 1.19 (Dynkin's $\pi - \lambda$ theorem) If $\mathcal{E} \subset 2^\Omega$ is a π -system, then

$$\sigma(\mathcal{E}) = \delta(\mathcal{E}).$$

Proof " $\supset$ " This follows from Remark 1.17.

" $\subset$ " We have to show that $\delta(\mathcal{E})$ is a σ -algebra. By Theorem 1.18, it is enough to show that $\delta(\mathcal{E})$ is a π -system. For any $B \in \delta(\mathcal{E})$ define

$$\mathcal{D}_B := \{A \in \delta(\mathcal{E}) : A \cap B \in \delta(\mathcal{E})\}.$$




Fig. 1.1 Inclusions between classes of sets $\mathcal{A} \subset 2^\Omega$

In order to show that $\delta(\mathcal{E})$ is a π -system, it is enough to show that

$$\delta(\mathcal{E}) \subset \mathcal{D}_B \quad \text{for any } B \in \delta(\mathcal{E}). \quad (1.3)$$

In order to show that $\mathcal{D}_E$ is a λ -system for any $E \in \delta(\mathcal{E})$, we check (i)-(iii) of Definition 1.10:

- (i) Clearly, $\Omega \cap E = E \in \delta(\mathcal{E})$; hence $\Omega \in \mathcal{D}_E$.
- (ii) For any $A, B \in \mathcal{D}_E$ with $A \subset B$, we have $(B \setminus A) \cap E = (B \cap E) \setminus (A \cap E) \in \delta(\mathcal{E})$.
- (iii) Assume that $A_1, A_2, \dots \in \mathcal{D}_E$ are mutually disjoint. Hence

$$\left(\bigcup_{n=1}^{\infty} A_n \right) \cap E = \biguplus_{n=1}^{\infty} (A_n \cap E) \in \delta(\mathcal{E}).$$

By assumption, $A \cap E \in \mathcal{E}$ if $A, E \in \mathcal{E}$; thus $\mathcal{E} \subset \mathcal{D}_E$ if $E \in \mathcal{E}$. By Remark 1.17(ii), we conclude that $\delta(\mathcal{E}) \subset \mathcal{D}_E$ for any $E \in \mathcal{E}$. Hence we get that $B \cap E \in \delta(\mathcal{E})$ for any $B \in \delta(\mathcal{E})$ and $E \in \mathcal{E}$. This implies that $E \in \mathcal{D}_B$ for any $B \in \delta(\mathcal{E})$. Thus $\mathcal{E} \subset \mathcal{D}_B$ for any $B \in \delta(\mathcal{E})$, and hence (1.3) follows. $\square$

For an illustration of the inclusions between the classes of sets, see Fig. 1.1.

We are particularly interested in σ -algebras that are generated by topologies. The most prominent role is played by the Euclidean space $\mathbb{R}^n$; however, we will also consider the (infinite-dimensional) space $C([0, 1])$ of continuous functions $[0, 1] \rightarrow \mathbb{R}$. On $C([0, 1])$ the norm $\|f\|_{\infty} = \sup_{x \in [0, 1]} |f(x)|$ induces a topology. For the convenience of the reader, we recall the definition of a topology.

Definition 1.20 (Topology) Let $\Omega \neq \emptyset$ be an arbitrary set. A class of sets $\tau \subset 2^{\Omega}$ is called a topology on Ω if it has the following three properties:

- (i) $\emptyset, \Omega \in \tau$.
- (ii) $A \cap B \in \tau$ for any $A, B \in \tau$.
- (iii) $(\bigcup_{A \in \mathcal{F}} A) \in \tau$ for any $\mathcal{F} \subset \tau$.

The pair (Ω, τ) is called a topological space. The sets $A \in \tau$ are called open, and the sets $A \subset \Omega$ with $A^c \in \tau$ are called closed.

In contrast with σ -algebras, topologies are closed under finite intersections only, but they are also closed under arbitrary unions.

Let d be a metric on Ω , and denote the open ball with radius $r > 0$ centered at $x \in \Omega$ by

$$B_r(x) = \{y \in \Omega : d(x, y) < r\}.$$

Then the usual class of open sets is the topology

$$\tau = \left\{ \bigcup_{(x,r) \in F} B_r(x) : F \subset \Omega \times (0, \infty) \right\}$$

Definition 1.21 (Borel σ -algebra) Let (Ω, τ) be a topological space. The σ -algebra

$$\mathcal{B}(\Omega) := \mathcal{B}(\Omega, \tau) := \sigma(\tau)$$

that is generated by the open sets is called the Borel σ -algebra Ω . The elements $A \in \mathcal{B}(\Omega, \tau)$ are called Borel sets or Borel measurable sets.

Remark 1.22 In many cases, we are interested in $\mathcal{B}(\mathbb{R}^n)$, where $\mathbb{R}^n$ is equipped with the Euclidean distance

$$d(x, y) = \|x - y\|_2 = \sqrt{\sum_{i=1}^n (x_i - y_i)^2}.$$

(i) There are subsets of $\mathbb{R}^n$ that are not Borel sets. These sets are not easy to construct like, for example, Vitali sets that can be found in calculus books (see also [37, Theorem 3.4.4]). Here we do not want to stress this point but state that, vaguely speaking, all sets that can be constructed explicitly are Borel sets.

(ii) If $C \subset \mathbb{R}^n$ is a closed set, then $C^c \in \tau$ is in $\mathcal{B}(\mathbb{R}^n)$ and hence C is a Borel set. In particular, $\{x\} \in \mathcal{B}(\mathbb{R}^n)$ for every $x \in \mathbb{R}^n$.

(iii) $\mathcal{B}(\mathbb{R}^n)$ is not a topology. To show this, let $V \subset \mathbb{R}^n$ such that $V \notin \mathcal{B}(\mathbb{R}^n)$. If $\mathcal{B}(\mathbb{R}^n)$ were a topology, then it would be closed under arbitrary unions. As $\{x\} \in \mathcal{B}(\mathbb{R}^n)$ for all $x \in \mathbb{R}^n$, we would get the contradiction $V = \bigcup_{x \in V} \{x\} \in \mathcal{B}(\mathbb{R}^n)$.

In most cases the class of open sets that generates the Borel σ -algebra is too big to work with efficiently. Hence we aim at finding smaller (in particular, countable) classes of sets that generate the Borel σ -algebra and that are more amenable. In

some of the examples, the elements of the generating class are simpler sets such as rectangles or compact sets.

We introduce the following notation. We denote by $\mathbb{Q}$ the set of rational numbers and by $\mathbb{Q}^+$ the set of strictly positive rational numbers. For $a, b \in \mathbb{R}^n$, we write

$$a < b \quad \text{if } a_i < b_i \quad \text{for all } i = 1, \dots, n. \quad (1.4)$$

For $a < b$, we define the open rectangle as the Cartesian product

$$(a, b) := X_{i=1}^n (a_i, b_i) := (a_1, b_1) \times (a_2, b_2) \times \dots \times (a_n, b_n). \quad (1.5)$$

Analogously, we define $[a, b]$, $(a, b]$ and $[a, b)$. Furthermore, we define $(-\infty, b) := \times_{i=1}^n (-\infty, b_i)$, and use an analogous definition for $(-\infty, b]$ and so on. We introduce the following classes of sets:

$$\begin{aligned} \mathcal{E}_1 &:= \{A \subset \mathbb{R}^n : A \text{ is open}\}, & \mathcal{E}_2 &:= \{A \subset \mathbb{R}^n : A \text{ is closed}\}, \\ \mathcal{E}_3 &:= \{A \subset \mathbb{R}^n : A \text{ is compact}\}, & \mathcal{E}_4 &:= \{B_r(x) : x \in \mathbb{Q}^n, r \in \mathbb{Q}^+\}, \\ \mathcal{E}_5 &:= \{(a, b) : a, b \in \mathbb{Q}^n, a < b\}, & \mathcal{E}_6 &:= \{[a, b) : a, b \in \mathbb{Q}^n, a < b\}, \\ \mathcal{E}_7 &:= \{(a, b] : a, b \in \mathbb{Q}^n, a < b\}, & \mathcal{E}_8 &:= \{[a, b] : a, b \in \mathbb{Q}^n, a < b\}, \\ \mathcal{E}_9 &:= \{(-\infty, b) : b \in \mathbb{Q}^n\}, & \mathcal{E}_{10} &:= \{(-\infty, b] : b \in \mathbb{Q}^n\}, \\ \mathcal{E}_{11} &:= \{(a, \infty) : a \in \mathbb{Q}^n\}, & \mathcal{E}_{12} &:= \{[a, \infty) : a \in \mathbb{Q}^n\}. \end{aligned}$$

Theorem 1.23 The Borel σ -algebra $\mathcal{B}(\mathbb{R}^n)$ is generated by any of the classes of sets $\mathcal{E}_1, \dots, \mathcal{E}_{12}$, that is, $\mathcal{B}(\mathbb{R}^n) = \sigma(\mathcal{E}_i)$ for any $i = 1, \dots, 12$.

Proof We show only some of the identities.

(1) By definition, $\mathcal{B}(\mathbb{R}^n) = \sigma(\mathcal{E}_1)$.

(2) Let $A \in \mathcal{E}_1$. Then $A^c \in \mathcal{E}_2$, and hence $A = (A^c)^c \in \sigma(\mathcal{E}_2)$. It follows that $\mathcal{E}_1 \subset \sigma(\mathcal{E}_2)$. By Remark 1.17, this implies $\sigma(\mathcal{E}_1) \subset \sigma(\mathcal{E}_2)$. Similarly, we obtain $\sigma(\mathcal{E}_2) \subset \sigma(\mathcal{E}_1)$ and hence equality.

(3) Any compact set is closed; hence $\sigma(\mathcal{E}_3) \subset \sigma(\mathcal{E}_2)$. Now let $A \in \mathcal{E}_2$. The sets $A_K := A \cap [-K, K]^n$, $K \in \mathbb{N}$, are compact; hence the countable union $A = \bigcup_{K=1}^{\infty} A_K$ is in $\sigma(\mathcal{E}_3)$. It follows that $\mathcal{E}_2 \subset \sigma(\mathcal{E}_3)$ and thus $\sigma(\mathcal{E}_2) = \sigma(\mathcal{E}_3)$.

(4) Clearly, $\mathcal{E}_4 \subset \mathcal{E}_1$; hence $\sigma(\mathcal{E}_4) \subset \sigma(\mathcal{E}_1)$. Now let $A \subset \mathbb{R}^n$ be an open set. For any $x \in A$, define $R(x) = \min(1, \sup\{r > 0 : B_r(x) \subset A\})$. Note that $R(x) > 0$, as A is open. Let $r(x) \in (R(x)/2, R(x)) \cap \mathbb{Q}$. For any $y \in A$ and $x \in (B_{R(y)/3}(y)) \cap \mathbb{Q}^n$, we have $R(x) \geq R(y) - \|x - y\|_2 > \frac{2}{3}R(y)$, and hence $r(x) > \frac{1}{3}R(y)$ and thus $y \in B_{r(x)}(x)$. It follows that $A = \bigcup_{x \in A \cap \mathbb{Q}^n} B_{r(x)}(x)$ is a countable union of sets from $\mathcal{E}_4$ and is hence in $\sigma(\mathcal{E}_4)$. We have shown that $\mathcal{E}_1 \subset \sigma(\mathcal{E}_4)$. By Remark 1.17, this implies $\sigma(\mathcal{E}_1) \subset \sigma(\mathcal{E}_4)$.

(5-12) Exhaustion arguments similar to that in (4) also work for rectangles. If in (4) we take open rectangles instead of open balls $B_r(x)$, we get $\mathcal{B}(\mathbb{R}^n) = \sigma(\mathcal{E}_{\nabla})$. For example, we have

$$\bigcap_{i=1}^n [a_i, b_i] = \bigcap_{k=1}^{\infty} \bigcap_{i=1}^n \left(a_i - \frac{1}{k}, b_i \right) \in \sigma(\mathcal{E}_5).$$

The other inclusions $\mathcal{E}_i \subset \sigma(\mathcal{E}_j)$ can be shown similarly. $\square$

Remark 1.24 Any of the classes $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3, \mathcal{E}_5, \dots, \mathcal{E}_{12}$ (but not $\mathcal{E}_4$) is a π -system. Hence, the Borel σ -algebra equals the generated λ -system: $\mathcal{B}(\mathbb{R}^n) = \delta(\mathcal{E}_i)$ for $i = 1, 2, 3, 5, \dots, 12$. In addition, the classes $\mathcal{E}_4, \dots, \mathcal{E}_{12}$ are countable. This is a crucial property that will be needed later. $\square$

Definition 1.25 (Trace of a class of sets) Let $\mathcal{A} \subset 2^\Omega$ be an arbitrary class of subsets of Ω and let $A \in 2^\Omega \setminus \{\emptyset\}$. The class

$$\mathcal{A}|_A := \{A \cap B : B \in \mathcal{A}\} \subset 2^A \quad (1.6)$$

is called the trace of $\mathcal{A}$ on A or the restriction of $\mathcal{A}$ to A .

Theorem 1.26 Let $A \subset \Omega$ be a nonempty set and let $\mathcal{A}$ be a σ -algebra on Ω or any of the classes of Definitions 1.6-1.9. Then $\mathcal{A}|_A$ is a class of sets of the same type as $\mathcal{A}$; however, on A instead of Ω . For λ -systems this is not true in general.

Proof This is left as an exercise. $\square$

Exercise 1.1.1 Let $\mathcal{A}$ be a semiring. Show that any countable (respectively finite) union of sets in $\mathcal{A}$ can be written as a countable (respectively finite) disjoint union of sets in $\mathcal{A}$.

Exercise 1.1.2 Give a counterexample that shows that, in general, the union $\mathcal{A} \cup \mathcal{A}'$ of two σ -algebras need not be a σ -algebra.

Exercise 1.1.3 Let (Ω_1, d_1) and (Ω_2, d_2) be metric spaces and let $f : \Omega_1 \rightarrow \Omega_2$ be an arbitrary map. Denote by $U_f = \{x \in \Omega_1 : f \text{ is discontinuous at } x\}$ the set of points of discontinuity of f . Show that $U_f \in \mathcal{B}(\Omega_1)$.

Hint: First show that for any $\varepsilon > 0$ and $\delta > 0$ the set

$$U_f^{\delta, \varepsilon} := \{x \in \Omega_1 : \text{there are } y, z \in B_\varepsilon(x) \text{ with } d_2(f(y), f(z)) > \delta\}$$

is open (where $B_\varepsilon(x) = \{y \in \Omega_1 : d_1(x, y) < \varepsilon\}$). Then construct U_f from such $U_f^{\delta, \varepsilon}$.

Exercise 1.1.4 Let Ω be an uncountably infinite set and $\mathcal{A} = \sigma(\{\omega\} : \omega \in \Omega)$. Show that

$$\mathcal{A} = \{A \subset \Omega : A \text{ is countable or } A^c \text{ is countable}\}.$$

Exercise 1.1.5 Let $\mathcal{A}$ be a ring on the set Ω . Show that $\mathcal{A}$ is an Abelian algebraic ring with multiplication " $\cap$ " and addition " Δ ".

5.1.2 1.2 Set Functions

Definition 1.27 Let $\mathcal{A} \subset 2^\Omega$ and let $\mu : \mathcal{A} \rightarrow [0, \infty]$ be a set function. We say that μ is

- (i) monotone if $\mu(A) \leq \mu(B)$ for any two sets $A, B \in \mathcal{A}$ with $A \subset B$,
- (ii) additive if $\mu(\biguplus_{i=1}^n A_i) = \sum_{i=1}^n \mu(A_i)$ for any choice of finitely many mutually disjoint sets $A_1, \dots, A_n \in \mathcal{A}$ with $\bigcup_{i=1}^n A_i \in \mathcal{A}$,
- (iii) σ -additive if $\mu(\biguplus_{i=1}^\infty A_i) = \sum_{i=1}^\infty \mu(A_i)$ for any choice of countably many mutually disjoint sets $A_1, A_2, \dots \in \mathcal{A}$ with $\bigcup_{i=1}^\infty A_i \in \mathcal{A}$,
- (iv) subadditive if for any choice of finitely many sets $A, A_1, \dots, A_n \in \mathcal{A}$ with $A \subset \bigcup_{i=1}^n A_i$, we have $\mu(A) \leq \sum_{i=1}^n \mu(A_i)$, and
- (v) σ -subadditive if for any choice of countably many sets $A, A_1, A_2, \dots \in \mathcal{A}$ with $A \subset \bigcup_{i=1}^\infty A_i$, we have $\mu(A) \leq \sum_{i=1}^\infty \mu(A_i)$.

Definition 1.28 Let $\mathcal{A}$ be a semiring and let $\mu : \mathcal{A} \rightarrow [0, \infty]$ be a set function with $\mu(\emptyset) = 0$. μ is called a

- content if μ is additive,
- premeasure if μ is σ -additive,
- measure if μ is a premeasure and $\mathcal{A}$ is a σ -algebra, and
- probability measure if μ is a measure and $\mu(\Omega) = 1$.

Definition 1.29 Let $\mathcal{A}$ be a semiring. A content μ on $\mathcal{A}$ is called

- (i) finite if $\mu(A) < \infty$ for every $A \in \mathcal{A}$ and
- (ii) σ -finite if there exists a sequence of sets $\Omega_1, \Omega_2, \dots \in \mathcal{A}$ such that $\Omega = \bigcup_{n=1}^{\infty} \Omega_n$ and such that $\mu(\Omega_n) < \infty$ for all $n \in \mathbb{N}$.

Example 1.30 (Contents, measures)

- (i) Let $\omega \in \Omega$ and $\delta_\omega(A) = \mathbb{1}_A(\omega)$ (see(1.2)). Then δ_ω is a probability measure on any σ -algebra $\mathcal{A} \subset 2^\Omega$. δ_ω is called the Dirac measure for the point ω .
- (ii) Let Ω be a finite nonempty set. By

$$\mu(A) := \frac{\#A}{\#\Omega} \quad \text{for } A \subset \Omega,$$

we define a probability measure on $\mathcal{A} = 2^\Omega$. This μ is called the uniform distribution on Ω . For this distribution, we introduce the symbol $\mathcal{U}_\Omega := \mu$. The resulting triple $(\Omega, \mathcal{A}, \mathcal{U}_\Omega)$ is called a Laplace space.

- (iii) Let Ω be countably infinite and let

$$\mathcal{A} := \{A \subset \Omega : \#A < \infty \text{ or } \#A^c < \infty\}.$$

Then $\mathcal{A}$ is an algebra. The set function μ on $\mathcal{A}$ defined by

$$\mu(A) = \begin{cases} 0, & A \text{ is finite,} \\ \infty, & A^c \text{ is finite,} \end{cases}$$

is a content but is not a premeasure. Indeed, $\mu(\bigcup_{\omega \in \Omega} \{\omega\}) = \mu(\Omega) = \infty$, but $\sum_{\omega \in \Omega} \mu(\{\omega\}) = 0$.

- (iv) Let $(\mu_n)_{n \in \mathbb{N}}$ be a sequence of measures (premeasures, contents) and let $(\alpha_n)_{n \in \mathbb{N}}$ be a sequence of nonnegative numbers. Then also $\mu := \sum_{n=1}^{\infty} \alpha_n \mu_n$ is a measure (premeasure, content).

- (v) Let Ω be an (at most) countable nonempty set and let $\mathcal{A} = 2^\Omega$. Further, let $(p_\omega)_{\omega \in \Omega}$ be nonnegative numbers. Then $A \mapsto \mu(A) := \sum_{\omega \in A} p_\omega$ defines a σ -finite measure on 2^Ω . We call $p = (p_\omega)_{\omega \in \Omega}$ the weight function of μ . The number p_ω is called the weight of μ at point ω .

- (vi) If in (v) the sum $\sum_{\omega \in \Omega} p_\omega$ equals one, then μ is a probability measure. In this case, we interpret p_ω as the probability of the elementary event ω . The vector $p = (p_\omega)_{\omega \in \Omega}$ is called a probability vector.

- (vii) If in (v) $p_\omega = 1$ for every $\omega \in \Omega$, then μ is called counting measure on Ω . If Ω is finite, then so is μ .

- (viii) Let $\mathcal{A}$ be the ring of finite unions of intervals $(a, b] \subset \mathbb{R}$. For $a_1 < b_1 < a_2 < b_2 < \dots < b_n$ and $A = \biguplus_{i=1}^n (a_i, b_i]$, define

$$\mu(A) = \sum_{i=1}^n (b_i - a_i).$$

Then μ is a σ -finite content on $\mathcal{A}$ (even a premeasure) since $\bigcup_{n=1}^{\infty} (-n, n] = \mathbb{R}$ and $\mu((-n, n]) = 2n < \infty$ for all $n \in \mathbb{N}$.

- (ix) Let $f : \mathbb{R} \rightarrow [0, \infty)$ be continuous. In a similar way to (viii), we define

$$\mu_f(A) = \sum_{i=1}^n \int_{a_i}^{b_i} f(x) dx$$

Then μ_f is a σ -finite content on $\mathcal{A}$ (even a premeasure). The function f is called the density of μ and plays a role similar to the weight function p in (v).

Lemma 1.31 (Properties of contents) Let $\mathcal{A}$ be a semiring and let μ be a content on $\mathcal{A}$. Then the following statements hold.

- (i) If $\mathcal{A}$ is a ring, then $\mu(A \cup B) + \mu(A \cap B) = \mu(A) + \mu(B)$ for any two sets $A, B \in \mathcal{A}$.
- (ii) μ is monotone. If $\mathcal{A}$ is a ring, then $\mu(B) = \mu(A) + \mu(B \setminus A)$ for any two sets $A, B \in \mathcal{A}$ with $A \subset B$.
- (iii) μ is subadditive. If μ is σ -additive, then μ is also σ -subadditive.
- (iv) If $\mathcal{A}$ is a ring, then $\sum_{n=1}^{\infty} \mu(A_n) \leq \mu(\bigcup_{n=1}^{\infty} A_n)$ for any choice of countably many mutually disjoint sets $A_1, A_2, \dots \in \mathcal{A}$ with $\bigcup_{n=1}^{\infty} A_n \in \mathcal{A}$.

Proof (i) Note that $A \cup B = A \uplus (B \setminus A)$ and $B = (A \cap B) \uplus (B \setminus A)$. As μ is additive, we obtain

$$\mu(A \cup B) = \mu(A) + \mu(B \setminus A) \quad \text{and} \quad \mu(B) = \mu(A \cap B) + \mu(B \setminus A).$$

This implies (i).

- (ii) Let $A \subset B$. Since $A \cap B = A$, we obtain $\mu(B) = \mu(A \uplus (B \setminus A)) = \mu(A) + \mu(B \setminus A)$ if $B \setminus A \in \mathcal{A}$. In particular, this is true if $\mathcal{A}$ is a ring. If $\mathcal{A}$ is only a semiring, then there exists an $n \in \mathbb{N}$ and mutually disjoint sets $C_1, \dots, C_n \in \mathcal{A}$ such that $B \setminus A = \biguplus_{i=1}^n C_i$. Hence $\mu(B) = \mu(A) + \sum_{i=1}^n \mu(C_i) \geq \mu(A)$ and thus μ is monotone.
- (iii) Let $n \in \mathbb{N}$ and $A, A_1, \dots, A_n \in \mathcal{A}$ with $A \subset \bigcup_{i=1}^n A_i$. Define $B_1 = A_1$ and

$$B_k = A_k \setminus \bigcup_{i=1}^{k-1} A_i = \bigcap_{i=1}^{k-1} (A_k \setminus (A_k \cap A_i)) \quad \text{for } k = 2, \dots, n.$$

By the definition of a semiring, any $A_k \setminus (A_k \cap A_i)$ is a finite disjoint union of sets in $\mathcal{A}$. Hence there exists a $c_k \in \mathbb{N}$ and sets $C_{k,1}, \dots, C_{k,c_k} \in \mathcal{A}$ such that $\biguplus_{i=1}^{c_k} C_{k,i} = B_k \subset A_k$. Similarly, there exist $d_k \in \mathbb{N}$ and $D_{k,1}, \dots, D_{k,d_k} \in \mathcal{A}$ such that $A_k \setminus B_k = \biguplus_{i=1}^{d_k} D_{k,i}$. Since μ is additive, we have

$$\mu(A_k) = \sum_{i=1}^{c_k} \mu(C_{k,i}) + \sum_{i=1}^{d_k} \mu(D_{k,i}) \geq \sum_{i=1}^{c_k} \mu(C_{k,i}).$$

Again due to additivity and monotonicity, we get

$$\begin{aligned} \mu(A) &= \mu\left(\biguplus_{k=1}^n \biguplus_{i=1}^{c_k} (C_{k,i} \cap A)\right) = \sum_{k=1}^n \sum_{i=1}^{c_k} \mu(C_{k,i} \cap A) \\ &\leq \sum_{k=1}^n \sum_{i=1}^{c_k} \mu(C_{k,i}) \leq \sum_{k=1}^n \mu(A_k). \end{aligned}$$

Hence μ is subadditive. By a similar argument, σ -subadditivity follows from σ additivity.

- (iv) Let $\mathcal{A}$ be a ring and let $A = \bigcup_{n=1}^{\infty} A_n \in \mathcal{A}$. Since μ is additive (and thus monotone), we have by (ii)

$$\sum_{n=1}^m \mu(A_n) = \mu\left(\biguplus_{n=1}^m A_n\right) \leq \mu(A) \quad \text{for any } m \in \mathbb{N}.$$

It follows that $\sum_{n=1}^{\infty} \mu(A_n) \leq \mu(A)$.

Remark 1.32 The inequality in (iv) can be strict (see Example 1.30(iii)). In other words, there are contents that are not premeasures.

Theorem 1.33 (Inclusion-exclusion formula) Let $\mathcal{A}$ be a ring and let μ be a content on $\mathcal{A}$. Let $n \in \mathbb{N}$ and $A_1, \dots, A_n \in \mathcal{A}$ such that $\mu(A_1 \cup \dots \cup A_n) < \infty$. Then the following inclusion and exclusion formulas hold:

$$\begin{aligned} \mu(A_1 \cup \dots \cup A_n) &= \sum_{k=1}^n (-1)^{k-1} \sum_{\{i_1, \dots, i_k\} \subset \{1, \dots, n\}} \mu(A_{i_1} \cap \dots \cap A_{i_k}), \\ \mu(A_1 \cap \dots \cap A_n) &= \sum_{k=1}^n (-1)^{k-1} \sum_{\{i_1, \dots, i_k\} \subset \{1, \dots, n\}} \mu(A_{i_1} \cup \dots \cup A_{i_k}). \end{aligned}$$

Here summation is over all subsets of $\{1, \dots, n\}$ with k elements.

Proof This is left as an exercise. Hint: Use induction on n .

The next goal is to characterize σ -subadditivity by a certain continuity property (Theorem 1.36). To this end, we agree on the following conventions.

Definition 1.34 Let $A, A_1, A_2, \dots$ be sets. We write

- $A_n \uparrow A$ and say that $(A_n)_{n \in \mathbb{N}}$ increases to A if $A_1 \subset A_2 \subset \dots$ and $\bigcup_{n=1}^{\infty} A_n = A$, and
- $A_n \downarrow A$ and say that $(A_n)_{n \in \mathbb{N}}$ decreases to A if $A_1 \supset A_2 \supset A_3 \supset \dots$ and $\bigcap_{n=1}^{\infty} A_n = A$.

Definition 1.35 (Continuity of contents) Let μ be a content on the ring $\mathcal{A}$.

- (i) μ is called lower semicontinuous if $\mu(A_n) \xrightarrow{n \rightarrow \infty} \mu(A)$ for any $A \in \mathcal{A}$ and any sequence $(A_n)_{n \in \mathbb{N}}$ in $\mathcal{A}$ with $A_n \uparrow A$.
- (ii) μ is called upper semicontinuous if $\mu(A_n) \xrightarrow{n \rightarrow \infty} \mu(A)$ for any $A \in \mathcal{A}$ and any sequence $(A_n)_{n \in \mathbb{N}}$ in $\mathcal{A}$ with $\mu(A_n) < \infty$ for some (and then eventually all) $n \in \mathbb{N}$ and $A_n \downarrow A$.
- (iii) μ is called $\emptyset$ -continuous if (ii) holds for $A = \emptyset$.

In the definition of upper semicontinuity, we needed the assumption $\mu(A_n) < \infty$ since otherwise we would not even have $\emptyset$ -continuity for an example as simple as the counting measure μ on $(\mathbb{N}, 2^{\mathbb{N}})$. Indeed, $A_n := \{n, n+1, \dots\} \downarrow \emptyset$ but $\mu(A_n) = \infty$ for all $n \in \mathbb{N}$.

Theorem 1.36 (Continuity and premeasure) Let μ be a content on the ring $\mathcal{A}$. Consider the following five properties.

- (i) μ is σ -additive (and hence a premeasure).
- (ii) μ is σ -subadditive.
- (iii) μ is lower semicontinuous.
- (iv) μ is $\emptyset$ -continuous.
- (v) μ is upper semicontinuous.

Then the following implications hold:

$$(i) \iff (ii) \iff (iii) \implies (iv) \iff (v).$$

If μ is finite, then we also have $(iv) \implies (iii)$.

Proof "(i) $\implies$ (ii)" Let $A, A_1, A_2, \dots \in \mathcal{A}$ with $A \subset \bigcup_{i=1}^{\infty} A_i$. Define $B_1 = A_1$ and $B_n = A_n \setminus \bigcup_{i=1}^{n-1} A_i \in \mathcal{A}$ for $n = 2, 3, \dots$. Then $A = \bigsqcup_{n=1}^{\infty} (A \cap B_n)$. Since μ is monotone and σ -additive, we infer

$$\mu(A) = \sum_{n=1}^{\infty} \mu(A \cap B_n) \leq \sum_{n=1}^{\infty} \mu(A_n).$$

Hence μ is σ -subadditive.

"(ii) $\implies$ (i)" This follows from Lemma 1.31(iv).

"(i) $\implies$ (iii)" Let μ be a premeasure and $A \in \mathcal{A}$. Let $(A_n)_{n \in \mathbb{N}}$ be a sequence in $\mathcal{A}$ such that $A_n \uparrow A$ and let $A_0 = \emptyset$. Then

$$\mu(A) = \sum_{i=1}^{\infty} \mu(A_i \setminus A_{i-1}) = \lim_{n \rightarrow \infty} \sum_{i=1}^n \mu(A_i \setminus A_{i-1}) = \lim_{n \rightarrow \infty} \mu(A_n).$$

"(iii) $\implies$ (i)" Assume now that (iii) holds. Let $B_1, B_2, \dots \in \mathcal{A}$ be mutually disjoint, and assume that $B = \bigsqcup_{n=1}^{\infty} B_n \in \mathcal{A}$. Define $A_n = \bigcup_{i=1}^n B_i$ for all $n \in \mathbb{N}$. Then it follows from (iii) that

$$\mu(B) = \lim_{n \rightarrow \infty} \mu(A_n) = \sum_{i=1}^{\infty} \mu(B_i).$$

Hence μ is σ -additive and therefore a premeasure.

"(iv) $\implies$ (v)" Let $A, A_1, A_2, \dots \in \mathcal{A}$ with $A_n \downarrow A$ and $\mu(A_1) < \infty$. Define $B_n = A_n \setminus A \in \mathcal{A}$ for all $n \in \mathbb{N}$. Then $B_n \downarrow \emptyset$. This implies $\mu(A_n) - \mu(A) = \mu(B_n) \xrightarrow{n \rightarrow \infty} 0$.

"(v) $\implies$ (iv)" This is evident.

"(iii) $\implies$ (iv)" Let $A_1, A_2, \dots \in \mathcal{A}$ with $A_n \downarrow \emptyset$ and $\mu(A_1) < \infty$. Then $A_1 \setminus A_n \in \mathcal{A}$ for any $n \in \mathbb{N}$ and $A_1 \setminus A_n \uparrow A_1$. Hence

$$\mu(A_1) = \lim_{n \rightarrow \infty} \mu(A_1 \setminus A_n) = \mu(A_1) - \lim_{n \rightarrow \infty} \mu(A_n)$$

Since $\mu(A_1) < \infty$, we have $\lim_{n \rightarrow \infty} \mu(A_n) = 0$.

"(iv) $\implies$ (iii)" (for finite μ) Assume that $\mu(A) < \infty$ for every $A \in \mathcal{A}$ and that μ is $\emptyset$ -continuous. Let $A, A_1, A_2, \dots \in \mathcal{A}$ with $A_n \uparrow A$. Then we have $A \setminus A_n \downarrow \emptyset$ and

$$\mu(A) - \mu(A_n) = \mu(A \setminus A_n) \xrightarrow{n \rightarrow \infty} 0.$$

Hence (iii) follows.

Example 1.37 (Compare Example 1.30(iii)) Let Ω be a countable set, and define

$$\mathcal{A} = \{A \subset \Omega : \#A < \infty \text{ or } \#A^c < \infty\},$$

$$\mu(A) = \begin{cases} 0, & A \text{ is finite,} \\ \infty, & A \text{ is infinite.} \end{cases}$$

Then μ is an $\emptyset$ -continuous content but not a premeasure.

Definition 1.38

(i) A pair $(\Omega, \mathcal{A})$ consisting of a nonempty set Ω and a σ -algebra $\mathcal{A} \subset 2^\Omega$ is called a measurable space. The sets $A \in \mathcal{A}$ are called measurable sets. If Ω is at most countably infinite and if $\mathcal{A} = 2^\Omega$, then the measurable space $(\Omega, 2^\Omega)$ is called discrete.

(ii) A triple $(\Omega, \mathcal{A}, \mu)$ is called a measure space if $(\Omega, \mathcal{A})$ is a measurable space and if μ is a measure on $\mathcal{A}$.

(iii) If in addition $\mu(\Omega) = 1$, then $(\Omega, \mathcal{A}, \mu)$ is called a probability space. In this case, the sets $A \in \mathcal{A}$ are called events.

(iv) The set of all finite measures on $(\Omega, \mathcal{A})$ is denoted by $\mathcal{M}_f(\Omega) := \mathcal{M}_f(\Omega, \mathcal{A})$. The subset of probability measures is denoted by $\mathcal{M}_1(\Omega) := \mathcal{M}_1(\Omega, \mathcal{A})$. Finally, the set of σ -finite measures on $(\Omega, \mathcal{A})$ is denoted by $\mathcal{M}_\sigma(\Omega, \mathcal{A})$.

Exercise 1.2.1 Let $\mathcal{A} = \{(a, b] \cap \mathbb{Q} : a, b \in \mathbb{R}, a \leq b\}$. Define $\mu : \mathcal{A} \rightarrow [0, \infty)$ by $\mu((a, b] \cap \mathbb{Q}) = b - a$. Show that $\mathcal{A}$ is a semiring and μ is a content on $\mathcal{A}$ that is lower and upper semicontinuous but is not σ -additive.

5.1.3 1.3 The Measure Extension Theorem

In this section, we construct measures μ on σ -algebras. The starting point will be to define the values of μ on a smaller class of sets; that is, on a semiring. Under a mild consistency condition, the resulting set function can be extended to the whole σ -algebra.

Before we develop the complete theory, we begin with two examples.

Example 1.39 (Lebesgue measure) Let $n \in \mathbb{N}$ and let

$$\mathcal{A} = \{(a, b] : a, b \in \mathbb{R}^n, a < b\}$$

be the semiring of half open rectangles $(a, b] \subset \mathbb{R}^n$ (see (1.5)). The n -dimensional volume of such a rectangle is

$$\mu((a, b]) = \prod_{i=1}^n (b_i - a_i)$$

Can we extend the set function μ to a (uniquely determined) measure on the Borel σ -algebra $\mathcal{B}(\mathbb{R}^n) = \sigma(\mathcal{A})$? We will see that this is indeed possible. The resulting measure is called Lebesgue measure (or sometimes Lebesgue-Borel measure) λ on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$.

Example 1.40 (Product measure, Bernoulli measure) We construct a measure for an infinitely often repeated random experiment with finitely many possible outcomes. Let E be the set of possible outcomes. For $e \in E$, let $p_e \geq 0$ be the probability that e occurs. Hence $\sum_{e \in E} p_e = 1$. For a fixed realization of the repeated experiment, let $\omega_1, \omega_2, \dots \in E$ be the observed outcomes. Hence the space of all possible outcomes of the repeated experiment is $\Omega = E^{\mathbb{N}}$. As in Example 1.11(vi), we define the set of all sequences whose first n values are $\omega_1, \dots, \omega_n$:

$$[\omega_1, \dots, \omega_n] := \{\omega' \in \Omega : \omega'_i = \omega_i \text{ for any } i = 1, \dots, n\}. \quad (1.7)$$

Let $\mathcal{A}_0 = \{\emptyset\}$. For $n \in \mathbb{N}$, define the class of cylinder sets that depend only on the first n coordinates

$$\mathcal{A}_n := \{[\omega_1, \dots, \omega_n] : \omega_1, \dots, \omega_n \in E\} \quad (1.8)$$

and let $\mathcal{A} := \bigcup_{n=0}^{\infty} \mathcal{A}_n$.

We interpret $[\omega_1, \dots, \omega_n]$ as the event where the outcome of the first experiment is ω_1 , the outcome of the second experiment is ω_2 and finally the outcome of the n th experiment is ω_n . The outcomes of the other experiments do not play a role for the occurrence of this event. As the individual experiments ought to be independent, we should have for any choice $\omega_1, \dots, \omega_n \in E$ that the probability of the event $[\omega_1, \dots, \omega_n]$ is the product of the probabilities of the individual events; that is,

$$\mu([\omega_1, \dots, \omega_n]) = \prod_{i=1}^n p_{\omega_i}$$

This formula defines a content μ on the semiring $\mathcal{A}$, and our aim is to extend μ in a unique way to a probability measure on the σ -algebra $\sigma(\mathcal{A})$ that is generated by $\mathcal{A}$.

Before we do so, we make the following definition. Define the (ultra-)metric d on Ω by

$$d(\omega, \omega') = \begin{cases} 2^{-\inf\{n \in \mathbb{N} : \omega_n \neq \omega'_n\}}, & \omega \neq \omega' \\ 0, & \omega = \omega' \end{cases} \quad (1.9)$$

Hence (Ω, d) is a compact metric space. Clearly,

$$[\omega_1, \dots, \omega_n] = B_{2^{-n}}(\omega) = \{\omega' \in \Omega : d(\omega, \omega') < 2^{-n}\}.$$

The complement of $[\omega_1, \dots, \omega_n]$ is an open set, as it is the union of $(\#E)^n - 1$ open balls

$$[\omega_1, \dots, \omega_n]^c = \bigcup_{(\omega'_1, \dots, \omega'_n) \neq (\omega_1, \dots, \omega_n)} [\omega'_1, \dots, \omega'_n]$$

Since Ω is compact, the closed subset $[\omega_1, \dots, \omega_n]$ is compact. As in Theorem 1.23, it can be shown that $\sigma(\mathcal{A}) = \mathcal{B}(\Omega, d)$.

Exercise: Prove the statements made above.

The main result of this chapter is Carathéodory's measure extension theorem.

Theorem 1.41 (Carathéodory) Let $\mathcal{A} \subset 2^\Omega$ be a ring and let μ be a σ -finite premeasure on $\mathcal{A}$. There exists a unique measure $\tilde{\mu}$ on $\sigma(\mathcal{A})$ such that $\tilde{\mu}(A) = \mu(A)$ for all $A \in \mathcal{A}$. Furthermore, $\tilde{\mu}$ is σ -finite.

We prepare for the proof of this theorem with a couple of lemmas. In fact, we will show a slightly stronger statement in Theorem 1.53.

Lemma 1.42 (Uniqueness by an $\cap$ -closed generator) Let $(\Omega, \mathcal{A}, \mu)$ be a σ -finite measure space and let $\mathcal{E} \subset \mathcal{A}$ be a π -system that generates $\mathcal{A}$. Assume that there exist sets $\Omega_1, \Omega_2, \dots \in \mathcal{E}$ such that $\bigcup_{n=1}^\infty \Omega_n = \Omega$ and $\mu(\Omega_n) < \infty$ for all $n \in \mathbb{N}$. Then μ is uniquely determined by the values $\mu(E), E \in \mathcal{E}$.

If μ is a probability measure, the existence of the sequence $(\Omega_n)_{n \in \mathbb{N}}$ is not needed.

Proof Let ν be a (possibly different) σ -finite measure on $(\Omega, \mathcal{A})$ such that

$$\mu(E) = \nu(E) \quad \text{for every } E \in \mathcal{E}.$$

Let $E \in \mathcal{E}$ with $\mu(E) < \infty$. Consider the class of sets

$$\mathcal{D}_E = \{A \in \mathcal{A} : \mu(A \cap E) = \nu(A \cap E)\}.$$

In order to show that $\mathcal{D}_E$ is a λ -system, we check the properties of Definition 1.10:

- (i) Clearly, $\Omega \in \mathcal{D}_E$.
- (ii) Let $A, B \in \mathcal{D}_E$ with $A \supset B$. Then

$$\begin{aligned} \mu((A \setminus B) \cap E) &= \mu(A \cap E) - \mu(B \cap E) \\ &= \nu(A \cap E) - \nu(B \cap E) = \nu((A \setminus B) \cap E) \end{aligned}$$

Hence $A \setminus B \in \mathcal{D}_E$.

- (iii) Let $A_1, A_2, \dots \in \mathcal{D}_E$ be mutually disjoint and $A = \bigcup_{n=1}^\infty A_n$. Then

$$\mu(A \cap E) = \sum_{n=1}^\infty \mu(A_n \cap E) = \sum_{n=1}^\infty \nu(A_n \cap E) = \nu(A \cap E).$$

Hence $A \in \mathcal{D}_E$.

Clearly, $\mathcal{E} \subset \mathcal{D}_E$; hence $\delta(\mathcal{E}) \subset \mathcal{D}_E$. Since $\mathcal{E}$ is a π -system, Theorem 1.19 yields

$$\mathcal{A} \supset \mathcal{D}_E \supset \delta(\mathcal{E}) = \sigma(\mathcal{E}) = \mathcal{A}$$

Hence $\mathcal{D}_E = \mathcal{A}$.

This implies $\mu(A \cap E) = \nu(A \cap E)$ for any $A \in \mathcal{A}$ and $E \in \mathcal{E}$ with $\mu(E) < \infty$. Now let $\Omega_1, \Omega_2, \dots \in \mathcal{E}$ be a sequence such that $\bigcup_{n=1}^\infty \Omega_n = \Omega$ and $\mu(\Omega_n) < \infty$ for all $n \in \mathbb{N}$. Let $E_n := \bigcup_{i=1}^n \Omega_i, n \in \mathbb{N}$, and $E_0 = \emptyset$. Hence $E_n = \biguplus_{i=1}^n (E_{i-1}^c \cap \Omega_i)$. For any $A \in \mathcal{A}$ and $n \in \mathbb{N}$, we thus get

$$\mu(A \cap E_n) = \sum_{i=1}^n \mu((A \cap E_{i-1}^c) \cap \Omega_i) = \sum_{i=1}^n \nu((A \cap E_{i-1}^c) \cap \Omega_i) = \nu(A \cap E_n).$$

Since $E_n \uparrow \Omega$ and since μ and ν are lower semicontinuous, we infer

$$\mu(A) = \lim_{n \rightarrow \infty} \mu(A \cap E_n) = \lim_{n \rightarrow \infty} \nu(A \cap E_n) = \nu(A).$$

The additional statement is trivial as $\tilde{\mathcal{E}} := \mathcal{E} \cup \{\Omega\}$ is a π -system that generates $\mathcal{A}$, and the value $\mu(\Omega) = 1$ is given. Hence one can choose the constant sequence $E_n = \Omega, n \in \mathbb{N}$. However, note that it is not enough to assume that μ is finite. In this case, in general, the total mass $\mu(\Omega)$ is not uniquely determined by the values $\mu(E), E \in \mathcal{E}$; see Example 1.45(ii).

Example 1.43 Let $\Omega = \mathbb{Z}$ and $\mathcal{E} = \{E_n : n \in \mathbb{Z}\}$ where $E_n = (-\infty, n] \cap \mathbb{Z}$. Then $\mathcal{E}$ is a π -system and $\sigma(\mathcal{E}) = 2^\Omega$. Hence a finite measure μ on $(\Omega, 2^\Omega)$ is uniquely determined by the values $\mu(E_n), n \in \mathbb{Z}$.

However, a σ -finite measure on $\mathbb{Z}$ is not uniquely determined by the values on $\mathcal{E}$: Let μ be the counting measure on $\mathbb{Z}$ and let $\nu = 2\mu$. Hence $\mu(E) = \infty = \nu(E)$ for all $E \in \mathcal{E}$. In order to distinguish μ and ν one needs a generator that contains sets of finite measure (of μ). Do the sets $\tilde{F}_n = [-n, n] \cap \mathbb{Z}, n \in \mathbb{N}$ do the trick? Indeed, for any σ -finite measure μ , we have $\mu(\tilde{F}_n) < \infty$ for all $n \in \mathbb{N}$. However, the sets $\tilde{F}_n$ do not generate 2^Ω (but which σ -algebra?). We get things to work out better if we

modify the definition: $F_n = [-n/2, (n+1)/2] \cap \mathbb{Z}$. Now $\sigma(\{F_n, n \in \mathbb{N}\}) = 2^\Omega$, and hence $\mathcal{E} = \{F_n, n \in \mathbb{N}\}$ is a π -system that generates 2^Ω and such that $\mu(F_n) < \infty$ for all $n \in \mathbb{N}$. The conditions of the theorem are fulfilled as $F_n \uparrow \Omega$.

Example 1.44 (Distribution function) A probability measure μ on the space $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ is uniquely determined by the values $\mu((-\infty, b])$ (where $(-\infty, b] = \times_{i=1}^n (-\infty, b_i], b \in \mathbb{R}^n$). In fact, these sets form a π -system that generates $\mathcal{B}(\mathbb{R}^n)$ (see Theorem 1.23). In particular, a probability measure μ on $\mathbb{R}$ is uniquely determined by its distribution function $F : \mathbb{R} \rightarrow [0, 1], x \mapsto \mu((-\infty, x])$.

Example 1.45

(i) Let $\Omega = \{1, 2, 3, 4\}$ and $\mathcal{E} = \{\{1, 2\}, \{2, 3\}\}$. Clearly, $\sigma(\mathcal{E}) = 2^\Omega$ but $\mathcal{E}$ is not a π -system. In fact, here a probability measure μ is not uniquely determined by the values, say $\mu(\{1, 2\}) = \mu(\{2, 3\}) = \frac{1}{2}$. We give just two different possibilities: $\mu = \frac{1}{2}\delta_1 + \frac{1}{2}\delta_3$ and $\mu' = \frac{1}{2}\delta_2 + \frac{1}{2}\delta_4$.

(ii) Let $\Omega = \{1, 2\}$ and $\mathcal{E} = \{\{1\}\}$. Then $\mathcal{E}$ is a π -system that generates 2^Ω . Hence a probability measure μ is uniquely determined by the value $\mu(\{1\})$. However, a finite measure is not determined by its value on $\{1\}$, as $\mu = 0$ and $\nu = \delta_2$ are different finite measures that agree on $\mathcal{E}$.

Definition 1.46 (Outer measure) A set function $\mu^* : 2^\Omega \rightarrow [0, \infty]$ is called an outer measure if

- (i) $\mu^*(\emptyset) = 0$, and
- (ii) μ^* is monotone,
- (iii) μ^* is σ -subadditive.

Lemma 1.47 Let $\mathcal{A} \subset 2^\Omega$ be an arbitrary class of sets with $\emptyset \in \mathcal{A}$ and let μ be a nonnegative set function on $\mathcal{A}$ with $\mu(\emptyset) = 0$. For $A \subset \Omega$, define the set of countable coverings $\mathcal{F}$ with sets $F \in \mathcal{A}$:

$$\mathcal{U}(A) = \left\{ \mathcal{F} \subset \mathcal{A} : \mathcal{F} \text{ is at most countable and } A \subset \bigcup_{F \in \mathcal{F}} F \right\}$$

Define

$$\mu^*(A) := \inf \left\{ \sum_{F \in \mathcal{F}} \mu(F) : \mathcal{F} \in \mathcal{U}(A) \right\},$$

where $\inf \emptyset = \infty$. Then μ^* is an outer measure. If in addition μ is σ -subadditive, then $\mu^*(A) = \mu(A)$ for all $A \in \mathcal{A}$.

Proof We check properties (i)-(iii) of an outer measure.

- (i) Since $\emptyset \in \mathcal{A}$, we have $\{\emptyset\} \in \mathcal{U}(\emptyset)$; hence $\mu^*(\emptyset) = 0$.
- (ii) If $A \subset B$, then $\mathcal{U}(A) \supset \mathcal{U}(B)$; hence $\mu^*(A) \leq \mu^*(B)$.
- (iii) Let $A_n \subset \Omega$ for any $n \in \mathbb{N}$ and let $A \subset \bigcup_{n=1}^{\infty} A_n$. We show that $\mu^*(A) \leq \sum_{n=1}^{\infty} \mu^*(A_n)$. Without loss of generality, assume $\mu^*(A_n) < \infty$ and hence $\mathcal{U}(A_n) \neq \emptyset$ for all $n \in \mathbb{N}$. Fix $\varepsilon > 0$. For every $n \in \mathbb{N}$, choose a covering $\mathcal{F}_n \in \mathcal{U}(A_n)$ such that

$$\sum_{F \in \mathcal{F}_n} \mu(F) \leq \mu^*(A_n) + \varepsilon 2^{-n}.$$

Then $\mathcal{F} := \bigcup_{n=1}^{\infty} \mathcal{F}_n \in \mathcal{U}(A)$ and

$$\mu^*(A) \leq \sum_{F \in \mathcal{F}} \mu(F) \leq \sum_{n=1}^{\infty} \sum_{F \in \mathcal{F}_n} \mu(F) \leq \sum_{n=1}^{\infty} \mu^*(A_n) + \varepsilon.$$

Let $A \in \mathcal{A}$. Since $\{A\} \in \mathcal{U}(A)$, we have $\mu^*(A) \leq \mu(A)$. If μ is σ -subadditive, then for any $\mathcal{F} \in \mathcal{U}(A)$, we have $\sum_{F \in \mathcal{F}} \mu(F) \geq \mu(A)$; hence $\mu^*(A) \geq \mu(A)$. $\square$

Definition 1.48 (μ^* -measurable sets) Let μ^* be an outer measure. A set $A \in 2^\Omega$ is called μ^* -measurable if

$$\mu^*(A \cap E) + \mu^*(A^c \cap E) = \mu^*(E) \quad \text{for any } E \in 2^\Omega. \quad (1.10)$$

We write $\mathcal{M}(\mu^*) = \{A \in 2^\Omega : A \text{ is } \mu^*\text{-measurable}\}$.

Lemma 1.49 $A \in \mathcal{M}(\mu^*)$ if and only if

$$\mu^*(A \cap E) + \mu^*(A^c \cap E) \leq \mu^*(E) \quad \text{for any } E \in 2^\Omega.$$

Proof As μ^* is subadditive, the other inequality is trivial. $\square$

Lemma 1.50 $\mathcal{M}(\mu^*)$ is an algebra.

Proof We check properties (i)-(iii) of an algebra from Theorem 1.7.

- (i) $\Omega \in \mathcal{M}(\mu^*)$ is evident.
- (ii) (Closedness under complements) By definition, $A \in \mathcal{M}(\mu^*) \iff A^c \in \mathcal{M}(\mu^*)$.
- (iii) (π -system) Let $A, B \in \mathcal{M}(\mu^*)$ and $E \in 2^\Omega$. Then

$$\begin{aligned} & \mu^*((A \cap B) \cap E) + \mu^*((A \cap B)^c \cap E) \\ &= \mu^*(A \cap B \cap E) + \mu^*((A^c \cap B \cap E) \cup (A^c \cap B^c \cap E) \cup (A \cap B^c \cap E)) \\ &\leq \mu^*(A \cap B \cap E) + \mu^*(A^c \cap B \cap E) \\ &\quad + \mu^*(A^c \cap B^c \cap E) + \mu^*(A \cap B^c \cap E) \\ &= \mu^*(B \cap E) + \mu^*(B^c \cap E) \\ &= \mu^*(E) \end{aligned}$$

Here we used $A \in \mathcal{M}(\mu^*)$ in the last but one equality and $B \in \mathcal{M}(\mu^*)$ in the last equality.

Lemma 1.51 An outer measure μ^* is σ -additive on $\mathcal{M}(\mu^*)$.

Proof Let $A, B \in \mathcal{M}(\mu^*)$ with $A \cap B = \emptyset$. Then

$$\mu^*(A \cup B) = \mu^*(A \cap (A \cup B)) + \mu^*(A^c \cap (A \cup B)) = \mu^*(A) + \mu^*(B).$$

Inductively, we get (finite) additivity. By definition, μ^* is σ -subadditive; hence we conclude by Theorem 1.36 that μ^* is also σ -additive.

Lemma 1.52 If μ^* is an outer measure, then $\mathcal{M}(\mu^*)$ is a σ -algebra. In particular, μ^* is a measure on $\mathcal{M}(\mu^*)$.

Proof By Lemma 1.50, $\mathcal{M}(\mu^*)$ is an algebra and hence a π -system. By Theorem 1.18, it is sufficient to show that $\mathcal{M}(\mu^*)$ is a λ -system.

Hence, let $A_1, A_2, \dots \in \mathcal{M}(\mu^*)$ be mutually disjoint, and define $A := \biguplus_{n=1}^{\infty} A_n$. We have to show $A \in \mathcal{M}(\mu^*)$; that is,

$$\mu^*(A \cap E) + \mu^*(A^c \cap E) \leq \mu^*(E) \quad \text{for any } E \in 2^{\Omega}. \quad (1.11)$$

Let $B_n = \bigcup_{i=1}^n A_i$ for all $n \in \mathbb{N}$. For all $n \in \mathbb{N}$, we have

$$\begin{aligned} \mu^*(E \cap B_{n+1}) &= \mu^*((E \cap B_{n+1}) \cap B_n) + \mu^*((E \cap B_{n+1}) \cap B_n^c) \\ &= \mu^*(E \cap B_n) + \mu^*(E \cap A_{n+1}) \end{aligned}$$

Inductively, we get $\mu^*(E \cap B_n) = \sum_{i=1}^n \mu^*(E \cap A_i)$. The monotonicity of μ^* now implies that

$$\begin{aligned} \mu^*(E) &= \mu^*(E \cap B_n) + \mu^*(E \cap B_n^c) \\ &\geq \mu^*(E \cap B_n) + \mu^*(E \cap A^c) \\ &= \sum_{i=1}^n \mu^*(E \cap A_i) + \mu^*(E \cap A^c) \end{aligned}$$

Letting $n \rightarrow \infty$ and using the σ -subadditivity of μ^* , we conclude

$$\mu^*(E) \geq \sum_{i=1}^{\infty} \mu^*(E \cap A_i) + \mu^*(E \cap A^c) \geq \mu^*(E \cap A) + \mu^*(E \cap A^c)$$

Hence (1.11) holds and the proof is complete.

We come to an extension theorem for measures that makes slightly weaker assumptions than Carathéodory's theorem (Theorem 1.41).

Theorem 1.53 (Extension theorem for measures) Let $\mathcal{A}$ be a semiring and let $\mu : \mathcal{A} \rightarrow [0, \infty]$ be an additive, σ -subadditive and σ -finite set function with $\mu(\emptyset) = 0$.

Then there is a unique σ -finite measure $\tilde{\mu} : \sigma(\mathcal{A}) \rightarrow [0, \infty]$ such that $\tilde{\mu}(A) = \mu(A)$ for all $A \in \mathcal{A}$.

Proof As $\mathcal{A}$ is a π -system, uniqueness follows by Lemma 1.42.

In order to establish the existence of $\tilde{\mu}$, we define as in Lemma 1.47

$$\mu^*(A) := \inf \left\{ \sum_{F \in \mathcal{F}} \mu(F) : \mathcal{F} \in \mathcal{U}(A) \right\} \quad \text{for any } A \in 2^{\Omega}.$$

By Lemma 1.47, μ^* is an outer measure and $\mu^*(A) = \mu(A)$ for any $A \in \mathcal{A}$. We have to show that $\mathcal{M}(\mu^*) \supset \sigma(\mathcal{A})$. Since $\mathcal{M}(\mu^*)$ is a σ -algebra (Lemma 1.52), it is enough to show $\mathcal{A} \subset \mathcal{M}(\mu^*)$.

To this end, let $A \in \mathcal{A}$ and $E \in 2^{\Omega}$ with $\mu^*(E) < \infty$. Fix $\varepsilon > 0$. Then there is a sequence $E_1, E_2, \dots \in \mathcal{A}$ such that

$$E \subset \bigcup_{n=1}^{\infty} E_n \quad \text{and} \quad \sum_{n=1}^{\infty} \mu(E_n) \leq \mu^*(E) + \varepsilon.$$

Define $B_n := E_n \cap A \in \mathcal{A}$. Since $\mathcal{A}$ is a semiring, for every $n \in \mathbb{N}$ there is an $m_n \in \mathbb{N}$ and sets $C_n^1, \dots, C_n^{m_n} \in \mathcal{A}$ such that $E_n \setminus A = E_n \setminus B_n = \biguplus_{k=1}^{m_n} C_n^k$. Hence

$$E \cap A \subset \bigcup_{n=1}^{\infty} B_n, \quad E \cap A^c \subset \bigcup_{n=1}^{\infty} \biguplus_{k=1}^{m_n} C_n^k \quad \text{and} \quad E_n = B_n \uplus \biguplus_{k=1}^{m_n} C_n^k.$$

By the definition of the outer measure and since μ is assumed to be (finitely) additive, we get

$$\begin{aligned}\mu^*(E \cap A) + \mu^*(E \cap A^c) &\leq \sum_{n=1}^{\infty} \mu(B_n) + \sum_{n=1}^{\infty} \sum_{k=1}^{m_n} \mu(C_n^k) \\ &= \sum_{n=1}^{\infty} \left(\mu(B_n) + \sum_{k=1}^{m_n} \mu(C_n^k) \right) \\ &= \sum_{n=1}^{\infty} \mu(E_n) \\ &\leq \mu^*(E) + \varepsilon.\end{aligned}$$

Hence $\mu^*(E \cap A) + \mu^*(E \cap A^c) \leq \mu^*(E)$ and thus $A \in \mathcal{M}(\mu^*)$, which implies $\mathcal{A} \subset \mathcal{M}(\mu^*)$. Now define $\tilde{\mu} : \sigma(\mathcal{A}) \rightarrow [0, \infty]$, $A \mapsto \mu^*(A)$. By Lemma 1.51, $\tilde{\mu}$ is a measure and $\tilde{\mu}$ is σ -finite since μ is σ -finite.

Example 1.54 (Lebesgue measure, continuation of Example 1.39) We aim at extending the volume $\mu((a, b]) = \prod_{i=1}^n (b_i - a_i)$ that was defined on the class of rectangles $\mathcal{A} = \{(a, b] : a, b \in \mathbb{R}^n, a < b\}$ to the Borel σ -algebra $\mathcal{B}(\mathbb{R}^n)$. In order to check the assumptions of Theorem 1.53, we only have to check that μ is σ -subadditive. To this end, let $(a, b], (a(1), b(1)], (a(2), b(2)], \dots \in \mathcal{A}$ with

$$(a, b] \subset \bigcup_{k=1}^{\infty} (a(k), b(k)]$$

We show that

$$\mu((a, b]) \leq \sum_{k=1}^{\infty} \mu((a(k), b(k)]). \quad (1.12)$$

For this purpose we use a compactness argument to reduce (1.12) to finite additivity. Fix $\varepsilon > 0$. For any $k \in \mathbb{N}$, choose $b_\varepsilon(k) > b(k)$ such that

$$\mu((a(k), b_\varepsilon(k)]) \leq \mu((a(k), b(k)]) + \varepsilon 2^{-k-1}.$$

Further choose $a_\varepsilon \in (a, b)$ such that $\mu((a_\varepsilon, b]) \geq \mu((a, b]) - \frac{\varepsilon}{2}$. Now $[a_\varepsilon, b]$ is compact and

$$\bigcup_{k=1}^{\infty} (a(k), b_\varepsilon(k)) \supset \bigcup_{k=1}^{\infty} (a(k), b(k)] \supset (a, b] \supset [a_\varepsilon, b],$$

whence there exists a K_0 such that $\bigcup_{k=1}^{K_0} (a(k), b_\varepsilon(k)) \supset (a_\varepsilon, b]$. As μ is (finitely) subadditive (see Lemma 1.31(iii)), we obtain

$$\begin{aligned}\mu((a, b]) &\leq \frac{\varepsilon}{2} + \mu((a_\varepsilon, b]) \leq \frac{\varepsilon}{2} + \sum_{k=1}^{K_0} \mu((a(k), b_\varepsilon(k)]) \\ &\leq \frac{\varepsilon}{2} + \sum_{k=1}^{K_0} (\varepsilon 2^{-k-1} + \mu((a(k), b(k)])) \leq \varepsilon + \sum_{k=1}^{\infty} \mu((a(k), b(k)])\end{aligned}$$

Letting $\varepsilon \downarrow 0$ yields (1.12); hence μ is σ -subadditive.

Combining the last example with Theorem 1.53, we have shown the following theorem.

Theorem 1.55 (Lebesgue measure) There exists a uniquely determined measure λ^n on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ with the property that

$$\lambda^n((a, b]) = \prod_{i=1}^n (b_i - a_i) \quad \text{for all } a, b \in \mathbb{R}^n \text{ with } a < b.$$

λ^n is called the Lebesgue measure on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ or Lebesgue-Borel measure.

Example 1.56 (Lebesgue-Stieltjes measure) Let $\Omega = \mathbb{R}$ and $\mathcal{A} = \{(a, b] : a, b \in \mathbb{R}, a \leq b\}$. $\mathcal{A}$ is a semiring and $\sigma(\mathcal{A}) = \mathcal{B}(\mathbb{R})$, where $\mathcal{B}(\mathbb{R})$ is the Borel σ -algebra on $\mathbb{R}$. Furthermore, let $F : \mathbb{R} \rightarrow \mathbb{R}$ be monotone increasing and right continuous. We define a set function

$$\tilde{\mu}_F : \mathcal{A} \rightarrow [0, \infty), \quad (a, b] \mapsto F(b) - F(a).$$

Clearly, $\tilde{\mu}_F(\emptyset) = 0$ and $\tilde{\mu}_F$ is additive.

Let $(a, b], (a(1), b(1)], (a(2), b(2)], \dots \in \mathcal{A}$ such that $(a, b] \subset \bigcup_{n=1}^{\infty} (a(n), b(n)]$. Fix $\varepsilon > 0$ and choose $a_\varepsilon \in (a, b)$ such that $F(a_\varepsilon) - F(a) < \varepsilon/2$. This is possible, as F is right continuous. For any $k \in \mathbb{N}$, choose $b_\varepsilon(k) > b(k)$ such that

$$F(b_\varepsilon(k)) - F(b(k)) < \varepsilon 2^{-k-1}$$

As in Example 1.54, it can be shown that $\tilde{\mu}_F((a, b]) \leq \varepsilon + \sum_{k=1}^{\infty} \tilde{\mu}_F((a(k), b(k)])$. This implies that $\tilde{\mu}_F$ is σ -subadditive. By Theorem 1.53, we can extend $\tilde{\mu}_F$ uniquely to a σ -finite measure μ_F on $\mathcal{B}(\mathbb{R})$.

Definition 1.57 (Lebesgue-Stieltjes measure) The measure μ_F on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ defined by

$$\mu_F((a, b]) = F(b) - F(a) \quad \text{for all } a, b \in \mathbb{R} \text{ with } a < b$$

is called the Lebesgue-Stieltjes measure with distribution function F .

Example 1.58 Important special cases for the Lebesgue-Stieltjes measure are the following:

(i) If $F(x) = x$, then $\mu_F = \lambda^1$ is the Lebesgue measure on $\mathbb{R}$.

(ii) Let $f : \mathbb{R} \rightarrow [0, \infty)$ be continuous and let $F(x) = \int_0^x f(t)dt$ for all $x \in \mathbb{R}$. Then μ_F is the extension of the premeasure with density f that was defined in Example 1.30(ix).

(iii) Let $x_1, x_2, \dots \in \mathbb{R}$ and $\alpha_n \geq 0$ for all $n \in \mathbb{N}$ such that $\sum_{n=1}^{\infty} \alpha_n < \infty$. Then $F = \sum_{n=1}^{\infty} \alpha_n \mathbb{1}_{[x_n, \infty)}$ is the distribution function of the finite measure $\mu_F = \sum_{n=1}^{\infty} \alpha_n \delta_{x_n}$.

(iv) Let $x_1, x_2, \dots \in \mathbb{R}$ such that $\mu = \sum_{n=1}^{\infty} \delta_{x_n}$ is a σ -finite measure. Then μ is a Lebesgue-Stieltjes measure if and only if the sequence $(x_n)_{n \in \mathbb{N}}$ does not have a limit point. Indeed, if $(x_n)_{n \in \mathbb{N}}$ does not have a limit point, then by the Bolzano-Weierstraß theorem, $\#\{n \in \mathbb{N} : x_n \in [-K, K]\} < \infty$ for every $K > 0$. If we let $F(x) = \#\{n \in \mathbb{N} : x_n \in [0, x]\}$ for $x \geq 0$ and $F(x) = -\#\{n \in \mathbb{N} : x_n \in [x, 0]\}$, then $\mu = \mu_F$. On the other hand, if μ is a Lebesgue-Stieltjes measure, this is $\mu = \mu_F$ for some F , then $\#\{n \in \mathbb{N} : x_n \in (-K, K]\} = F(K) - F(-K) < \infty$ for all $K > 0$; hence $(x_n)_{n \in \mathbb{N}}$ does not have a limit point.

(v) If $\lim_{x \rightarrow \infty} (F(x) - F(-x)) = 1$, then μ_F is a probability measure.

We will now have a closer look at the case where μ_F is a probability measure.

Definition 1.59 (Distribution function) A right continuous monotone increasing function $F : \mathbb{R} \rightarrow [0, 1]$ with $F(-\infty) := \lim_{x \rightarrow -\infty} F(x) = 0$ and $F(\infty) := \lim_{x \rightarrow \infty} F(x) = 1$ is called a (proper) probability distribution function (p.d.f.). If we only have $F(\infty) \leq 1$ instead of $F(\infty) = 1$, then F is called a (possibly) defective p.d.f. If μ is a (sub-)probability measure on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$, then $F_\mu : x \mapsto \mu((-\infty, x])$ is called the distribution function of μ .

Clearly, F_μ is right continuous and $F(-\infty) = 0$, since μ is upper semicontinuous and finite (Theorem 1.36). Since μ is lower semicontinuous, we have $F(\infty) = \mu(\mathbb{R})$; hence F_μ is indeed a (possibly defective) distribution function if μ is a (sub-)probability measure.

The argument of Example 1.56 yields the following theorem.

Theorem 1.60 The map $\mu \mapsto F_\mu$ is a bijection from the set of probability measures on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ to the set of probability distribution functions, respectively from the set of sub-probability measures to the set of defective distribution functions.

We have established that every finite measure on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ is a Lebesgue-Stieltjes measure for some function F . For σ -finite measures, the corresponding statement does not hold in this generality as we saw in Example 1.58(iv).

We come now to a theorem that combines Theorem 1.55 with the idea of Lebesgue-Stieltjes measures. Later we will see that the following theorem is valid in greater generality. In particular, the assumption that the factors are of Lebesgue-Stieltjes type can be dropped.

Theorem 1.61 (Finite products of measures) Let $n \in \mathbb{N}$ and let $\mu_1, \dots, \mu_n$ be finite measures or, more generally, Lebesgue-Stieltjes measures on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$. Then there exists a unique σ -finite measure μ on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ such that

$$\mu((a, b]) = \prod_{i=1}^n \mu_i((a_i, b_i]) \quad \text{for all } a, b \in \mathbb{R}^n \text{ with } a < b.$$

We call $\mu =: \bigotimes_{i=1}^n \mu_i$ the product measure of the measures $\mu_1, \dots, \mu_n$.

Proof The proof is the same as for Theorem 1.55. One has to check that the intervals $(a, b_\varepsilon]$ and so on can be chosen such that $\mu((a, b_\varepsilon]) < \mu((a, b]) + \varepsilon$. Here we employ the right continuity of the increasing function F_i that belongs to μ_i . The details are left as an exercise.

Remark 1.62 Later we will see in Theorem 14.14 that the statement holds even for arbitrary σ -finite measures $\mu_1, \dots, \mu_n$ on arbitrary (even different) measurable spaces. One can even construct infinite products if all factors are probability spaces (Theorem 14.36).

Example 1.63 (Infinite product measure, continuation of Example 1.40) Let E be a finite set and let $\Omega = E^{\mathbb{N}}$ be the space of E -valued sequences. Further, let $(p_e)_{e \in E}$ be a probability vector. Define a content μ on $\mathcal{A} = \{[\omega_1, \dots, \omega_n] : \omega_1, \dots, \omega_n \in E, n \in \mathbb{N}\}$ by

$$\mu([\omega_1, \dots, \omega_n]) = \prod_{i=1}^n p_{\omega_i}$$

We aim at extending μ to a measure on $\sigma(\mathcal{A})$. In order to check the assumptions of Theorem 1.53, we have to show that μ is σ -subadditive. As in the preceding example, we use a compactness argument.

Let $A, A_1, A_2, \dots \in \mathcal{A}$ and $A \subset \bigcup_{n=1}^{\infty} A_n$. We are done if we can show that there exists an $N \in \mathbb{N}$ such that

$$A \subset \bigcup_{n=1}^N A_n. \quad (1.13)$$

Indeed, due to the (finite) subadditivity of μ (see Lemma 1.31(iii)), this implies $\mu(A) \leq \sum_{n=1}^N \mu(A_n) \leq \sum_{n=1}^{\infty} \mu(A_n)$; hence μ is σ -subadditive.

We now give two different proofs for (1.13).

1st Proof. The metric d from (1.9) induces the product topology on Ω ; hence, as remarked in Example 1.40, (Ω, d) is a compact metric space. Every $A \in \mathcal{A}$ is closed and thus compact. Since every A_n is also open, A can be covered by finitely many A_n ; hence (1.13) holds.

2nd Proof. We now show by elementary means the validity of (1.13). The procedure imitates the proof that Ω is compact. Let $B_n := A \setminus \bigcup_{i=1}^n A_i$. We assume $B_n \neq \emptyset$ for all $n \in \mathbb{N}$ in order to get a contradiction. By Dirichlet's pigeonhole principle (recall that E is finite), we can choose $\omega_1 \in E$ such that $[\omega_1] \cap B_n \neq \emptyset$ for infinitely many $n \in \mathbb{N}$. Since $B_1 \supset B_2 \supset \dots$, we obtain

$$[\omega_1] \cap B_n \neq \emptyset \quad \text{for all } n \in \mathbb{N}.$$

Successively choose $\omega_2, \omega_3, \dots \in E$ in such a way that

$$[\omega_1, \dots, \omega_k] \cap B_n \neq \emptyset \quad \text{for all } k, n \in \mathbb{N}.$$

B_n is a disjoint union of certain sets $C_{n,1}, \dots, C_{n,m_n} \in \mathcal{A}$. Hence, for every $n \in \mathbb{N}$ there is an $i_n \in \{1, \dots, m_n\}$ such that $[\omega_1, \dots, \omega_k] \cap C_{n,i_n} \neq \emptyset$ for infinitely many $k \in \mathbb{N}$. Since $[\omega_1] \supset [\omega_1, \omega_2] \supset \dots$, we obtain

$$[\omega_1, \dots, \omega_k] \cap C_{n,i_n} \neq \emptyset \quad \text{for all } k, n \in \mathbb{N}.$$

For fixed $n \in \mathbb{N}$ and large k , we have $[\omega_1, \dots, \omega_k] \subset C_{n,i_n}$. Hence $\omega = (\omega_1, \omega_2, \dots) \in C_{n,i_n} \subset B_n$. This implies $\bigcap_{n=1}^{\infty} B_n \neq \emptyset$, contradicting the assumption.

Combining the last example with Theorem 1.53, we have shown the following theorem.

Theorem 1.64 (Product measure, Bernoulli measure) Let E be a finite nonempty set and $\Omega = E^{\mathbb{N}}$. Let $(p_e)_{e \in E}$ be a probability vector. Then there exists a unique probability measure μ on $\sigma(\mathcal{A}) = \mathcal{B}(\Omega)$ such that

$$\mu([\omega_1, \dots, \omega_n]) = \prod_{i=1}^n p_{\omega_i} \quad \text{for all } \omega_1, \dots, \omega_n \in E \text{ and } n \in \mathbb{N}$$

μ is called the product measure or Bernoulli measure on Ω with weights $(p_e)_{e \in E}$. We write $(\sum_{e \in E} p_e \delta_e)^{\otimes \mathbb{N}} := \mu$. The σ -algebra $(2^E)^{\otimes \mathbb{N}} := \sigma(\mathcal{A})$ is called the product σ -algebra on Ω .

We will study product measures in a systematic way in Chapter 14.

The measure extension theorem yields an abstract statement of existence and uniqueness for measures on $\sigma(\mathcal{A})$ that were first defined on a semiring $\mathcal{A}$ only. The following theorem, however, shows that the measure of a set from $\sigma(\mathcal{A})$ can be well approximated by finite and countable operations with sets from $\mathcal{A}$.

Denote by

$$A \Delta B := (A \setminus B) \cup (B \setminus A) \quad \text{for } A, B \subset \Omega \quad (1.14)$$

the symmetric difference of the two sets A and B .

Theorem 1.65 (Approximation theorem for measures) Let $\mathcal{A} \subset 2^{\Omega}$ be a semiring and let μ be a measure on $\sigma(\mathcal{A})$ that is σ -finite on $\mathcal{A}$.

- (i) For any $A \in \sigma(\mathcal{A})$ and $\varepsilon > 0$, there exist mutually disjoint sets $A_1, A_2, \dots \in \mathcal{A}$ such that $A \subset \bigcup_{n=1}^{\infty} A_n$ and $\mu(\bigcup_{n=1}^{\infty} A_n \setminus A) < \varepsilon$.
- (ii) For any $A \in \sigma(\mathcal{A})$ with $\mu(A) < \infty$ and any $\varepsilon > 0$, there exists an $n \in \mathbb{N}$ and mutually disjoint sets $A_1, \dots, A_n \in \mathcal{A}$ such that $\mu(A \Delta \bigcup_{k=1}^n A_k) < \varepsilon$.
- (iii) For any $A \in \mathcal{M}(\mu^*)$, there are sets $A_-, A_+ \in \sigma(\mathcal{A})$ with $A_- \subset A \subset A_+$ and $\mu(A_+ \setminus A_-) = 0$.

Remark 1.66 (iii) implies that (i) and (ii) also hold for $A \in \mathcal{M}(\mu^*)$ (with μ^* instead of μ). If $\mathcal{A}$ is an algebra, then in (ii) for any $A \in \sigma(\mathcal{A})$, we even have $\inf_{B \in \mathcal{A}} \mu(A \Delta B) = 0$.

Proof (ii) As μ and the outer measure μ^* coincide on $\sigma(\mathcal{A})$ and since $\mu(A)$ is finite, by the very definition of μ^* (see Lemma 1.47) there exists a covering $B_1, B_2, \dots \in \mathcal{A}$ of A such that

$$\mu(A) \geq \sum_{i=1}^{\infty} \mu(B_i) - \varepsilon/2$$

Let $n \in \mathbb{N}$ with $\sum_{i=n+1}^{\infty} \mu(B_i) < \frac{\varepsilon}{2}$ (such an n exists since $\mu(A) < \infty$). For any three sets C, D, E , we have

$$C \Delta D = (D \setminus C) \cup (C \setminus D) \subset (D \setminus C) \cup (C \setminus (D \cup E)) \cup E \subset (C \Delta (D \cup E)) \cup E.$$

Choosing $C = A, D = \bigcup_{i=1}^n B_i$ and $E = \bigcup_{i=n+1}^{\infty} B_i$, this yields

$$\begin{aligned} \mu\left(A \Delta \bigcup_{i=1}^n B_i\right) &\leq \mu\left(A \Delta \bigcup_{i=1}^{\infty} B_i\right) + \mu\left(\bigcup_{i=n+1}^{\infty} B_i\right) \\ &\leq \mu\left(\bigcup_{i=1}^{\infty} B_i\right) - \mu(A) + \frac{\varepsilon}{2} \leq \varepsilon. \end{aligned}$$

As $\mathcal{A}$ is a semiring, there exist a $k \in \mathbb{N}$ and $A_1, \dots, A_k \in \mathcal{A}$ such that

$$\bigcup_{i=1}^n B_i = B_1 \uplus \biguplus_{i=2}^n \bigcap_{j=1}^{i-1} (B_i \setminus B_j) =: \biguplus_{i=1}^k A_i.$$

(i) Let $A \in \sigma(\mathcal{A})$ and $E_n \uparrow \Omega$, $E_n \in \sigma(\mathcal{A})$ with $\mu(E_n) < \infty$ for any $n \in \mathbb{N}$. For every $n \in \mathbb{N}$, choose a covering $(B_{n,m})_{m \in \mathbb{N}}$ of $A \cap E_n$ with

$$\mu(A \cap E_n) \geq \sum_{m=1}^{\infty} \mu(B_{n,m}) - 2^{-n} \varepsilon$$

(This is possible due to the definition of the outer measure μ^* , which coincides with μ on $\mathcal{A}$.) Let $\bigcup_{n=1}^{\infty} B_{n,m} = \biguplus_{n=1}^{\infty} A_n$ for certain $A_n \in \mathcal{A}$, $n \in \mathbb{N}$ (Exercise 1.1.1). Then

$$\begin{aligned} \mu\left(\biguplus_{n=1}^{\infty} A_n \setminus A\right) &= \mu\left(\bigcup_{n=1}^{\infty} \bigcup_{m=1}^{\infty} B_{n,m} \setminus A\right) \\ &\leq \mu\left(\bigcup_{n=1}^{\infty} \bigcup_{m=1}^{\infty} (B_{n,m} \setminus (A \cap E_n))\right) \\ &\leq \sum_{n=1}^{\infty} \left(\left(\sum_{m=1}^{\infty} \mu(B_{n,m}) \right) - \mu(A \cap E_n) \right) \leq \varepsilon \end{aligned}$$

(iii) Let $A \in \mathcal{M}(\mu^*)$ and $(E_n)_{n \in \mathbb{N}}$ as above. For any $m, n \in \mathbb{N}$, choose $A_{n,m} \in \sigma(\mathcal{A})$ such that $A_{n,m} \supset A \cap E_n$ and $\mu^*(A_{n,m}) \leq \mu^*(A \cap E_n) + \frac{2^{-n}}{m}$.

Define $A_m := \bigcup_{n=1}^{\infty} A_{n,m} \in \sigma(\mathcal{A})$. Then $A_m \supset A$ and $\mu^*(A_m \setminus A) \leq \frac{1}{m}$. Define $A_+ := \bigcap_{m=1}^{\infty} A_m$. Then $\sigma(\mathcal{A}) \ni A_+ \supset A$ and $\mu^*(A_+ \setminus A) = 0$. Similarly, choose $(A_-)^c \in \sigma(\mathcal{A})$ with $(A_-)^c \supset A^c$ and $\mu^*((A_-)^c \setminus A^c) = 0$. Then $A_+ \supset A \supset A_-$ and $\mu(A_+ \setminus A_-) = \mu^*(A_+ \setminus A_-) = \mu^*(A_+ \setminus A) + \mu^*(A \setminus A_-) = 0$.

Remark 1.67 (Regularity of measures) (Compare with Theorem 13.6.) Let λ^n be the Lebesgue measure on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$. Let $\mathcal{A}$ be the semiring of rectangles of the form $(a, b] \subset \mathbb{R}^n$; hence $\mathcal{B}(\mathbb{R}^n) = \sigma(\mathcal{A})$ by Theorem 1.23. By the approximation theorem, for any $A \in \mathcal{B}(\mathbb{R}^n)$ and $\varepsilon > 0$, there exist countably many $A_1, A_2, \dots \in \mathcal{A}$ with $A \subset \bigcup_{i=1}^{\infty} A_i$ and

$$\lambda^n\left(\bigcup_{i=1}^{\infty} A_i \setminus A\right) < \varepsilon/2$$

For any A_i , there exists an open rectangle $B_i \supset A_i$ with $\lambda^n(B_i \setminus A_i) < \varepsilon 2^{-i-1}$ (upper semicontinuity of λ^n). Hence $U = \bigcup_{i=1}^{\infty} B_i$ is an open set $U \supset A$ with

$$\lambda^n(U \setminus A) < \varepsilon.$$

This property of λ^n is called outer regularity.

If $\lambda^n(A)$ is finite, then for any $\varepsilon > 0$ there exists a compact $K \subset A$ such that

$$\lambda^n(A \setminus K) < \varepsilon.$$

This property of λ^n is called inner regularity. Indeed, let $N > 0$ be such that $\lambda^n(A) - \lambda^n(A \cap [-N, N]^n) < \varepsilon/2$. Choose an open set $U \supset (A \cap [-N, N]^n)^c$ such that $\lambda^n(U \setminus (A \cap [-N, N]^n)^c) < \varepsilon/2$, and let $K := [-N, N]^n \setminus U \subset A$.

Definition 1.68 (Null set) Let $(\Omega, \mathcal{A}, \mu)$ be a measure space.

(i) A set $A \in \mathcal{A}$ is called a μ -null set, or briefly a null set, if $\mu(A) = 0$. By $\mathcal{N}_{\mu}$ we denote the class of all subsets of μ -null sets.

(ii) Let $E(\omega)$ be a property that a point $\omega \in \Omega$ can have or not have. We say that E holds μ -almost everywhere (a.e.) or for almost all (a.a.) ω if there exists a null set N such that $E(\omega)$ holds for every $\omega \in \Omega \setminus N$. If $A \in \mathcal{A}$ and if there exists a null set N such that $E(\omega)$ holds for every $\omega \in A \setminus N$, then we say that E holds almost everywhere on A .

If $\mu = P$ is a probability measure, then we say that E holds P -almost surely (a.s.), respectively almost surely on A .

(iii) Let $A, B \in \mathcal{A}$ be such that $\mu(A \Delta B) = 0$. Then we write $A = B \pmod{\mu}$.

Definition 1.69 A measure space $(\Omega, \mathcal{A}, \mu)$ is called complete if $\mathcal{N}_\mu \subset \mathcal{A}$.

Remark 1.70 (Completion of a measure space) Let $(\Omega, \mathcal{A}, \mu)$ be a σ -finite measure space. There exists a unique smallest σ -algebra $\mathcal{A}^* \supset \mathcal{A}$ and an extension μ^* of μ to $\mathcal{A}^*$ such that $(\Omega, \mathcal{A}^*, \mu^*)$ is complete. $(\Omega, \mathcal{A}^*, \mu^*)$ is called the completion of $(\Omega, \mathcal{A}, \mu)$. With the notation of Theorem 1.53, this completion is

$$(\Omega, \mathcal{M}(\mu^*), \mu^*|_{\mathcal{M}(\mu^*)}).$$

Furthermore,

$$\mathcal{M}(\mu^*) = \sigma(\mathcal{A} \cup \mathcal{N}_\mu) = \{A \cup N : A \in \mathcal{A}, N \in \mathcal{N}_\mu\}$$

and $\mu^*(A \cup N) = \mu(A)$ for any $A \in \mathcal{A}$ and $N \in \mathcal{N}_\mu$.

In the following, we will not need these statements. Hence, instead of giving a proof, we refer to the textbooks on measure theory (e.g., [37]).

Example 1.71 Let λ be the Lebesgue measure (more accurately, the Lebesgue-Borel measure) on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$. Then λ can be extended uniquely to a measure λ^* on

$$\mathcal{B}^*(\mathbb{R}^n) = \sigma(\mathcal{B}(\mathbb{R}^n) \cup \mathcal{N}),$$

where $\mathcal{N}$ is the class of subsets of Lebesgue-Borel null sets. $\mathcal{B}^*(\mathbb{R}^n)$ is called the σ algebra of Lebesgue measurable sets. For the sake of distinction, we sometimes call λ the Lebesgue-Borel measure and λ^* the Lebesgue measure. However, in practice, this distinction will not be needed in this book.

Example 1.72 Let $\mu = \delta_\omega$ be the Dirac measure for the point $\omega \in \Omega$ on some measurable space $(\Omega, \mathcal{A})$. If $\{\omega\} \in \mathcal{A}$, then the completion is $\mathcal{A}^* = 2^\Omega$, $\mu^* = \delta_\omega$. In the extreme case of a trivial σ -algebra $\mathcal{A} = \{\emptyset, \Omega\}$, however, the empty set is the only null set, $\mathcal{N}_\mu = \{\emptyset\}$; hence $\mathcal{A}^* = \{\emptyset, \Omega\}$, $\mu^* = \delta_\omega$. Note that, on the trivial σ -algebra, Dirac measures for different points $\omega \in \Omega$ cannot be distinguished.

Definition 1.73 Let $(\Omega, \mathcal{A}, \mu)$ be a measure space and $\Omega' \in \mathcal{A}$. On the trace σ algebra $\mathcal{A}|_{\Omega'}$, we define a measure by

$$\mu|_{\Omega'}(A) := \mu(A) \quad \text{for } A \in \mathcal{A} \text{ with } A \subset \Omega'.$$

This measure is called the restriction of μ to Ω' .

Example 1.74 The restriction of the Lebesgue-Borel measure λ on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ to $[0, 1]$ is a probability measure on $([0, 1], \mathcal{B}(\mathbb{R})|_{[0, 1]})$. More generally, for a measurable $A \in \mathcal{B}(\mathbb{R})$, we call the restriction $\lambda|_A$ the Lebesgue measure on A . Often this measure will be denoted by the same symbol λ when there is no danger of ambiguity.

Later we will see (Corollary 1.84) that $\mathcal{B}(\mathbb{R})|_A = \mathcal{B}(A)$, where $\mathcal{B}(A)$ is the Borel σ -algebra on A that is generated by the (relatively) open subsets of A .

Example 1.75 (Uniform distribution) Let $A \in \mathcal{B}(\mathbb{R}^n)$ be a measurable set with n dimensional Lebesgue measure $\lambda^n(A) \in (0, \infty)$. Then we can define a probability measure on $\mathcal{B}(\mathbb{R}^n)|_A$ by

$$\mu(B) := \frac{\lambda^n(B)}{\lambda^n(A)} \quad \text{for } B \in \mathcal{B}(\mathbb{R}^n) \text{ with } B \subset A.$$

This measure μ is called the uniform distribution on A and will be denoted by $\mathcal{U}_A := \mu$.

Exercise 1.3.1 Show the following generalization of Example 1.58(iv): A measure $\sum_{n=1}^{\infty} \alpha_n \delta_{x_n}$ is a Lebesgue-Stieltjes measure for a suitable function F if and only if $\sum_{n:|x_n| \leq K} \alpha_n < \infty$ for all $K > 0$.

Exercise 1.3.2 Let Ω be an uncountably infinite set and let $\omega_0 \in \Omega$ be an arbitrary element. Let $\mathcal{A} = \sigma(\{\omega\} : \omega \in \Omega \setminus \{\omega_0\})$.

(i) Give a characterization of $\mathcal{A}$ as in Exercise 1.1.4 (p. 11).

(ii) Show that $(\Omega, \mathcal{A}, \delta_{\omega_0})$ is complete.

Exercise 1.3.3 Let $(\mu_n)_{n \in \mathbb{N}}$ be a sequence of finite measures on the measurable space $(\Omega, \mathcal{A})$. Assume that for any $A \in \mathcal{A}$ there exists the limit $\mu(A) := \lim_{n \rightarrow \infty} \mu_n(A)$.

Show that μ is a measure on $(\Omega, \mathcal{A})$.

Hint: In particular, one has to show that μ is $\emptyset$ -continuous.

5.1.4 1.4 Measurable Maps

A major task of mathematics is to study homomorphisms between objects; that is, structure-preserving maps. For topological spaces, these are the continuous maps, and for measurable spaces, these are the measurable maps.

In the rest of this chapter, we let $(\Omega, \mathcal{A})$ and $(\Omega', \mathcal{A}')$ be measurable spaces.

Definition 1.76 (Measurable maps)

(i) A map $X : \Omega \rightarrow \Omega'$ is called $\mathcal{A} - \mathcal{A}'$ -measurable (or, briefly, measurable) if $X^{-1}(\mathcal{A}') := \{X^{-1}(A') : A' \in \mathcal{A}'\} \subset \mathcal{A}$; that is, if

$$X^{-1}(A') \in \mathcal{A} \quad \text{for any } A' \in \mathcal{A}'.$$

If X is measurable, we write $X : (\Omega, \mathcal{A}) \rightarrow (\Omega', \mathcal{A}')$.

(ii) If $\Omega' = \mathbb{R}$ and $\mathcal{A}' = \mathcal{B}(\mathbb{R})$ is the Borel σ -algebra on $\mathbb{R}$, then $X : (\Omega, \mathcal{A}) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$ is called an $\mathcal{A}$ -measurable real map.

Example 1.77

- (i) The identity map $\text{id} : \Omega \rightarrow \Omega$ is $\mathcal{A} - \mathcal{A}$ -measurable.
- (ii) If $\mathcal{A} = 2^\Omega$ or $\mathcal{A}' = \{\emptyset, \Omega'\}$, then any map $X : \Omega \rightarrow \Omega'$ is $\mathcal{A} - \mathcal{A}'$ -measurable.
- (iii) Let $A \subset \Omega$. The indicator function $\mathbb{1}_A : \Omega \rightarrow \{0, 1\}$ is $\mathcal{A} - 2^{\{0,1\}}$ -measurable if and only if $A \in \mathcal{A}$.

Theorem 1.78 (Generated σ -algebra) Let $(\Omega', \mathcal{A}')$ be a measurable space and let Ω be a nonempty set. Let $X : \Omega \rightarrow \Omega'$ be a map. The preimage

$$X^{-1}(\mathcal{A}') := \{X^{-1}(A') : A' \in \mathcal{A}'\} \tag{1.15}$$

is the smallest σ -algebra with respect to which X is measurable. We say that $\sigma(X) := X^{-1}(\mathcal{A}')$ is the σ -algebra on Ω that is generated by X .

Proof This is left as an exercise. $\square$

We now consider σ -algebras that are generated by more than one map.

Definition 1.79 (Generated σ -algebra) Let Ω be a nonempty set. Let I be an arbitrary index set. For any $i \in I$, let $(\Omega_i, \mathcal{A}_i)$ be a measurable space and let $X_i : \Omega \rightarrow \Omega_i$ be an arbitrary map. Then

$$\sigma(X_i, i \in I) := \sigma\left(\bigcup_{i \in I} \sigma(X_i)\right) = \sigma\left(\bigcup_{i \in I} X_i^{-1}(\mathcal{A}_i)\right)$$

is called the σ -algebra on Ω that is generated by $(X_i, i \in I)$. This is the smallest σ -algebra with respect to which all X_i are measurable.

As with continuous maps, the composition of measurable maps is again measurable.

Theorem 1.80 (Composition of maps) Let $(\Omega, \mathcal{A})$, $(\Omega', \mathcal{A}')$ and $(\Omega'', \mathcal{A}'')$ be measurable spaces and let $X : \Omega \rightarrow \Omega'$ and $X' : \Omega' \rightarrow \Omega''$ be measurable maps. Then the map $Y := X' \circ X : \Omega \rightarrow \Omega''$, $\omega \mapsto X'(X(\omega))$ is $\mathcal{A} - \mathcal{A}''$ -measurable.

Proof Obvious, since $Y^{-1}(\mathcal{A}'') = X^{-1}((X')^{-1}(\mathcal{A}'')) \subset X^{-1}(\mathcal{A}') \subset \mathcal{A}$. $\square$

In practice, it is often not possible to check if a map X is measurable by checking if all preimages $X^{-1}(A')$, $A' \in \mathcal{A}'$ are measurable. Most σ -algebras $\mathcal{A}'$ are simply too large. Thus it comes in very handy that it is sufficient to check measurability on a generator $\mathcal{E}'$ by the following theorem.

Theorem 1.81 (Measurability on a generator) Let $\mathcal{E}' \subset \mathcal{A}'$ be a class of $\mathcal{A}'$ measurable sets. Then $\sigma(X^{-1}(\mathcal{E}')) = X^{-1}(\sigma(\mathcal{E}'))$ and hence

$$X \text{ is } \mathcal{A} - \sigma(\mathcal{E}')\text{-measurable} \iff X^{-1}(E') \in \mathcal{A} \text{ for all } E' \in \mathcal{E}'.$$

If in particular $\sigma(\mathcal{E}') = \mathcal{A}'$, then

$$X \text{ is } \mathcal{A} - \mathcal{A}'\text{-measurable} \iff X^{-1}(\mathcal{E}') \subset \mathcal{A}.$$

Proof Clearly, $X^{-1}(\mathcal{E}') \subset X^{-1}(\sigma(\mathcal{E}')) = \sigma(X^{-1}(\sigma(\mathcal{E}')))$. Hence also

$$\sigma(X^{-1}(\mathcal{E}')) \subset X^{-1}(\sigma(\mathcal{E}')).$$

For the other inclusion, consider the class of sets

$$\mathcal{A}'_0 := \{A' \in \sigma(\mathcal{E}') : X^{-1}(A') \in \sigma(X^{-1}(\mathcal{E}'))\}.$$

We first show that $\mathcal{A}'_0$ is a σ -algebra by checking (i)-(iii) of Definition 1.2:

- (i) Clearly, $\Omega' \in \mathcal{A}'_0$.
- (ii) (Stability under complements) If $A' \in \mathcal{A}'_0$, then

$$X^{-1}((A')^c) = (X^{-1}(A'))^c \in \sigma(X^{-1}(\mathcal{E}'));$$

hence $(A')^c \in \mathcal{A}'_0$.

- (iii) ($\sigma - \cup$ -stability) Let $A'_1, A'_2, \dots \in \mathcal{A}'_0$. Then

$$X^{-1}\left(\bigcup_{n=1}^{\infty} A'_n\right) = \bigcup_{n=1}^{\infty} X^{-1}(A'_n) \in \sigma(X^{-1}(\mathcal{E}'));$$

hence $\bigcup_{n=1}^{\infty} A'_n \in \mathcal{A}'_0$.

Now $\mathcal{A}'_0 = \sigma(\mathcal{E}')$ since $\mathcal{E}' \subset \mathcal{A}'_0$. Hence $X^{-1}(A') \in \sigma(X^{-1}(\mathcal{E}'))$ for any $A' \in \sigma(\mathcal{E}')$ and thus $X^{-1}(\sigma(\mathcal{E}')) \subset \sigma(X^{-1}(\mathcal{E}'))$.

Corollary 1.82 (Measurability of composed maps) Let I be a nonempty index set and let $(\Omega, \mathcal{A})$, $(\Omega', \mathcal{A}')$ and $(\Omega_i, \mathcal{A}_i)$ be measurable spaces for any $i \in I$. Further, let $(X_i : i \in I)$ be a family of measurable maps $X_i : \Omega' \rightarrow \Omega_i$ with $\mathcal{A}' = \sigma(X_i : i \in I)$. Then the following holds: A map $Y : \Omega \rightarrow \Omega'$ is $\mathcal{A} - \mathcal{A}'$ -measurable if and only if $X_i \circ Y$ is $\mathcal{A} - \mathcal{A}_i$ -measurable for all $i \in I$.

Proof If Y is measurable, then by Theorem 1.80 every $X_i \circ Y$ is measurable. Now assume that all of the composed maps $X_i \circ Y$ are $\mathcal{A} - \mathcal{A}_i$ -measurable. By assumption, the set $\mathcal{E}' := \{X_i^{-1}(A'') : A'' \in \mathcal{A}_i, i \in I\}$ is a generator of $\mathcal{A}'$. Since all $X_i \circ Y$ are measurable, we have $Y^{-1}(A') \in \mathcal{A}$ for any $A' \in \mathcal{E}'$. Hence Theorem 1.81 yields that Y is measurable.

Recall the definition of the trace of a class of sets from Definition 1.25.

Corollary 1.83 (Trace of a generated σ -algebra) Let $\mathcal{E} \subset 2^\Omega$ and assume that $A \subset \Omega$ is nonempty. Then $\sigma(\mathcal{E}|_A) = \sigma(\mathcal{E})|_A$.

Proof Let $X : A \hookrightarrow \Omega, \omega \mapsto \omega$ be the canonical inclusion; hence $X^{-1}(B) = A \cap B$ for all $B \subset \Omega$. By Theorem 1.81, we have

$$\begin{aligned}\sigma(\mathcal{E}|_A) &= \sigma(\{E \cap A : E \in \mathcal{E}\}) \\ &= \sigma(\{X^{-1}(E) : E \in \mathcal{E}\}) = \sigma(X^{-1}(\mathcal{E})) \\ &= X^{-1}(\sigma(\mathcal{E})) = \{A \cap B : B \in \sigma(\mathcal{E})\} = \sigma(\mathcal{E})|_A.\end{aligned}$$

Recall that, for any subset $A \subset \Omega$ of a topological space (Ω, τ) , the class $\tau|_A$ is the topology of relatively open sets (in A). We denote by $\mathcal{B}(\Omega, \tau) = \sigma(\tau)$ the Borel σ -algebra on (Ω, τ) .

Corollary 1.84 (Trace of the Borel σ -algebra) Let (Ω, τ) be a topological space and let $A \subset \Omega$ be a nonempty subset of Ω . Then

$$\mathcal{B}(\Omega, \tau)|_A = \mathcal{B}(A, \tau|_A).$$

Example 1.85

(i) Let Ω' be countable. Then $X : \Omega \rightarrow \Omega'$ is $\mathcal{A}$ - $2^{\Omega'}$ -measurable if and only if $X^{-1}(\{\omega'\}) \in \mathcal{A}$ for all $\omega' \in \Omega'$. If Ω' is uncountably infinite, this is wrong in general. (For example, consider $\Omega = \Omega' = \mathbb{R}, \mathcal{A} = \mathcal{B}(\mathbb{R})$, and $X(\omega) = \omega$ for all $\omega \in \Omega$. Clearly, $X^{-1}(\{\omega\}) = \{\omega\} \in \mathcal{B}(\mathbb{R})$. If, on the other hand, $A \subset \mathbb{R}$ is not in $\mathcal{B}(\mathbb{R})$, then $A \in 2^{\mathbb{R}}$, but $X^{-1}(A) \notin \mathcal{B}(\mathbb{R})$.)

(ii) For $x \in \mathbb{R}$, we agree on the following notation for rounding:

$$\lfloor x \rfloor := \max\{k \in \mathbb{Z} : k \leq x\} \quad \text{and} \quad \lceil x \rceil := \min\{k \in \mathbb{Z} : k \geq x\}. \quad (1.16)$$

The maps $\mathbb{R} \rightarrow \mathbb{Z}, x \mapsto \lfloor x \rfloor$ and $x \mapsto \lceil x \rceil$ are $\mathcal{B}(\mathbb{R}) - 2^{\mathbb{Z}}$ -measurable since for all $k \in \mathbb{Z}$ the preimages $\{x \in \mathbb{R} : \lfloor x \rfloor = k\} = [k, k+1)$ and $\{x \in \mathbb{R} : \lceil x \rceil = k\} = (k-1, k]$ are in $\mathcal{B}(\mathbb{R})$. By the composition theorem (Theorem 1.80), for any measurable map $f : (\Omega, \mathcal{A}) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$ the maps $\lfloor f \rfloor$ and $\lceil f \rceil$ are also $\mathcal{A} - 2^{\mathbb{Z}}$ -measurable.

(iii) A map $X : \Omega \rightarrow \mathbb{R}^d$ is $\mathcal{A} - \mathcal{B}(\mathbb{R}^d)$ -measurable if and only if

$$X^{-1}((-\infty, a]) \in \mathcal{A} \quad \text{for any } a \in \mathbb{R}^d$$

In fact $\sigma((-\infty, a], a \in \mathbb{R}^d) = \mathcal{B}(\mathbb{R}^d)$ by Theorem 1.23. The analogous statement holds for any of the classes $\mathcal{E}_1, \dots, \mathcal{E}_{12}$ from Theorem 1.23.

Example 1.86 Let $d(x, y) = \|x - y\|_2$ be the usual Euclidean distance on $\mathbb{R}^n$ and let $\mathcal{B}(\mathbb{R}^n, d) = \mathcal{B}(\mathbb{R}^n)$ be the Borel σ -algebra with respect to the topology generated by d . For any subset A of $\mathbb{R}^n$, we have $\mathcal{B}(A, d) = \mathcal{B}(\mathbb{R}^n, d)|_A$.

We want to extend the real line by the points $-\infty$ and $+\infty$. Thus we define

$$\overline{\mathbb{R}} := \mathbb{R} \cup \{-\infty, +\infty\}$$

From a topological point of view, $\overline{\mathbb{R}}$ will be considered as the so-called two point compactification by considering $\overline{\mathbb{R}}$ as topologically isomorphic to $[-1, 1]$ via the map

$$\varphi : [-1, 1] \rightarrow \overline{\mathbb{R}}, \quad x \mapsto \begin{cases} \tan(\pi x/2), & x \in (-1, 1), \\ -\infty, & x = -1, \\ \infty, & x = +1. \end{cases}$$

In fact, $\bar{d}(x, y) = |\varphi^{-1}(x) - \varphi^{-1}(y)|$ for $x, y \in \overline{\mathbb{R}}$ defines a metric on $\overline{\mathbb{R}}$ such that φ and φ^{-1} are continuous. Hence φ is a topological isomorphism. We denote by $\bar{\tau}$ the corresponding topology induced on $\overline{\mathbb{R}}$ and by τ the usual topology on $\mathbb{R}$.

Corollary 1.87 With the above notation, $\bar{\tau}|_{\mathbb{R}} = \tau$ and hence $\mathcal{B}(\bar{\mathbb{R}})|_{\mathbb{R}} = \mathcal{B}(\mathbb{R})$.

In particular, if $X : (\Omega, \mathcal{A}) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$ is measurable, then in a canonical way X is also an $\bar{\mathbb{R}}$ -valued measurable map.

Thus $\bar{\mathbb{R}}$ is really an extension of the real line, and the inclusion $\mathbb{R} \hookrightarrow \bar{\mathbb{R}}$ is measurable.

Theorem 1.88 (Measurability of continuous maps) Let (Ω, τ) and (Ω', τ') be topological spaces and let $f : \Omega \rightarrow \Omega'$ be a continuous map. Then f is $\mathcal{B}(\Omega)$ - $\mathcal{B}(\Omega')$ -measurable.

Proof As $\mathcal{B}(\Omega') = \sigma(\tau')$ and by Theorem 1.81, it is sufficient to show that $f^{-1}(A') \in \sigma(\tau)$ for all $A' \in \tau'$. However, since f is continuous, we even have $f^{-1}(A') \in \tau$ for all $A' \in \tau'$.

For $x, y \in \bar{\mathbb{R}}$, we agree on the following notation.

$$\begin{aligned} x \vee y &= \max(x, y) && (\text{maximum}), \\ x \wedge y &= \min(x, y) && (\text{minimum}), \\ x^+ &= \max(x, 0) && (\text{positive part}), \\ x^- &= \max(-x, 0) && (\text{negative part}), \\ |x| &= \max(x, -x) = x^- + x^+ && (\text{modulus}), \\ \text{sign}(x) &= \mathbb{1}_{\{x>0\}} - \mathbb{1}_{\{x<0\}} && (\text{sign function}). \end{aligned}$$

Analogously, for measurable real maps we write, for example, $X^+ = \max(X, 0)$. The maps $x \mapsto x^+$, $x \mapsto x^-$ and $x \mapsto |x|$ are continuous (and hence measurable by the preceding theorem). Clearly, the map $x \mapsto \text{sign}(x)$ also is measurable. Using Corollary 1.82, we thus get the following corollary.

Corollary 1.89 If X is a real or $\bar{\mathbb{R}}$ -valued measurable map, then the maps X^- , X^+ , $|X|$ and $\text{sign}(X)$ also are measurable.

Theorem 1.90 (Coordinate maps are measurable) Let $(\Omega, \mathcal{A})$ be a measurable space and let $f_1, \dots, f_n : \Omega \rightarrow \mathbb{R}$ be maps. Define $f := (f_1, \dots, f_n) : \Omega \rightarrow \mathbb{R}^n$. Then

$$f \text{ is } \mathcal{A} - \mathcal{B}(\mathbb{R}^n)\text{-measurable} \iff \text{each } f_i \text{ is } \mathcal{A} - \mathcal{B}(\mathbb{R})\text{-measurable.}$$

The analogous statement holds for $f_i : \Omega \rightarrow \bar{\mathbb{R}} := \mathbb{R} \cup \{\pm\infty\}$.

Proof For $b \in \mathbb{R}^n$, we have $f^{-1}((-\infty, b)) = \bigcap_{i=1}^n f_i^{-1}((-\infty, b_i))$. If each f_i is measurable, then $f^{-1}((-\infty, b)) \in \mathcal{A}$. However, the rectangles $(-\infty, b)$, $b \in \mathbb{R}^n$, generate $\mathcal{B}(\mathbb{R}^n)$, and hence f is measurable. Now assume that f is measurable. For $i = 1, \dots, n$, let $\pi_i : \mathbb{R}^n \rightarrow \mathbb{R}$, $x \mapsto x_i$ be the projection on the i th coordinate. Clearly, π_i is continuous and thus $\mathcal{B}(\mathbb{R}^n) - \mathcal{B}(\mathbb{R})$ -measurable. Hence $f_i = \pi_i \circ f$ is measurable by Theorem 1.80.

In the following theorem, we agree that $\frac{x}{0} := 0$ for all $x \in \mathbb{R}$.

Theorem 1.91 Let $(\Omega, \mathcal{A})$ be a measurable space. Let $h : (\Omega, \mathcal{A}) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$ and $f, g : (\Omega, \mathcal{A}) \rightarrow (\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ be measurable maps. Then also the maps $f + g$, $f - g$, $f \cdot h$ and f/h are measurable.

Proof The map $\pi : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}^n$, $(x, \alpha) \mapsto \alpha \cdot x$ is continuous and thus measurable. By Theorem 1.90, $(f, h) : \Omega \rightarrow \mathbb{R}^n \times \mathbb{R}$ is measurable. Hence also the composed map $f \cdot h = \pi \circ (f, h)$ is measurable. Similarly, we obtain the measurability of $f + g$ and $f - g$.

In order to show measurability of f/h , we define the map $H : \mathbb{R} \rightarrow \mathbb{R}$, $x \mapsto 1/x$. Note that by our convention $H(0) = 0$. Hence $f/h = f \cdot H \circ h$. Thus it is enough to show that H is measurable. Clearly, $H|_{\mathbb{R} \setminus \{0\}}$ is continuous. For any open set $U \subset \mathbb{R}$, $U \setminus \{0\}$ is also open and hence $H^{-1}(U \setminus \{0\}) \in \mathcal{B}(\mathbb{R})$. Furthermore, $H^{-1}(\{0\}) = \{0\}$. Concluding, we get $H^{-1}(U) = H^{-1}(U \setminus \{0\}) \cup (U \cap \{0\}) \in \mathcal{B}(\mathbb{R})$.

Theorem 1.92 Let $X_1, X_2, \dots$ be measurable maps $(\Omega, \mathcal{A}) \rightarrow (\bar{\mathbb{R}}, \mathcal{B}(\bar{\mathbb{R}}))$. Then the following maps are also measurable:

$$\inf_{n \in \mathbb{N}} X_n, \quad \sup_{n \in \mathbb{N}} X_n, \quad \liminf_{n \rightarrow \infty} X_n, \quad \limsup_{n \rightarrow \infty} X_n.$$

Proof For any $a \in \overline{\mathbb{R}}$, we have

$$\left(\inf_{n \in \mathbb{N}} X_n \right)^{-1}([-\infty, a)) = \bigcup_{n=1}^{\infty} X_n^{-1}([-\infty, a)) \in \mathcal{A}$$

By Theorem 1.81, this implies that $\inf_{n \in \mathbb{N}} X_n$ is measurable. The proof for $\sup_{n \in \mathbb{N}} X_n$ is similar.

For any $n \in \mathbb{N}$, we define $Y_n := \inf_{m \geq n} X_m$. Note that Y_n is measurable and hence $\liminf_{n \rightarrow \infty} X_n := \sup_{n \in \mathbb{N}} Y_n$ also is measurable. The proof for the limes superior is similar.

We come to an important example of measurable maps $(\Omega, \mathcal{A}) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$, the so-called simple functions.

Definition 1.93 (Simple function) Let $(\Omega, \mathcal{A})$ be a measurable space. A map $f : \Omega \rightarrow \mathbb{R}$ is called a simple function if there is an $n \in \mathbb{N}$ and mutually disjoint measurable sets $A_1, \dots, A_n \in \mathcal{A}$, as well as numbers $\alpha_1, \dots, \alpha_n \in \mathbb{R}$, such that

$$f = \sum_{i=1}^n \alpha_i \mathbb{1}_{A_i}.$$

Remark 1.94 A measurable map that assumes only finitely many values is a simple function. (Exercise: Show this!)

Definition 1.95 Assume that $f, f_1, f_2, \dots$ are maps $\Omega \rightarrow \overline{\mathbb{R}}$ such that

$$f_1(\omega) \leq f_2(\omega) \leq \dots \quad \text{and} \quad \lim_{n \rightarrow \infty} f_n(\omega) = f(\omega) \quad \text{for any } \omega \in \Omega.$$

Then we write $f_n \uparrow f$ and say that $(f_n)_{n \in \mathbb{N}}$ increases (pointwise) to f . Analogously, we write $f_n \downarrow f$ if $(-f_n) \uparrow (-f)$.

Theorem 1.96 Let $(\Omega, \mathcal{A})$ be a measurable space and let $f : \Omega \rightarrow [0, \infty]$ be measurable. Then the following statements hold.

- (i) There exists a sequence $(f_n)_{n \in \mathbb{N}}$ of nonnegative simple functions such that $f_n \uparrow f$.
- (ii) There are measurable sets $A_1, A_2, \dots \in \mathcal{A}$ and numbers $\alpha_1, \alpha_2, \dots \geq 0$ such that $f = \sum_{n=1}^{\infty} \alpha_n \mathbb{1}_{A_n}$.

Proof (i) For $n \in \mathbb{N}_0$, define $f_n = (2^{-n} \lfloor 2^n f \rfloor) \wedge n$. Then f_n is measurable (by Theorem 1.92 and Example 1.85(ii)) and assumes at most $n2^n + 1$ different values. Hence it is a simple function. Clearly, $f_n \uparrow f$.

(ii) Let f_n be as above. Let $B_{n,i} := \{\omega : f_n(\omega) - f_{n-1}(\omega) = i2^{-n}\}$ and $\beta_{n,i} = i2^{-n}$ for $n \in \mathbb{N}$ and $i = 1, \dots, 2^n$. Hence $f_n - f_{n-1} = \sum_{i=1}^{2^n} \beta_{n,i} \mathbb{1}_{B_{n,i}}$. By changing the numeration $(n, i) \mapsto m$, we get $(\alpha_m)_{m \in \mathbb{N}}$ and $(A_m)_{m \in \mathbb{N}}$ such that

$$f = f_0 + \sum_{n=1}^{\infty} (f_n - f_{n-1}) = \sum_{m=1}^{\infty} \alpha_m \mathbb{1}_{A_m}$$

As a corollary to this statement on the structure of $[0, \infty]$ -valued measurable maps, we show the following factorization lemma.

Corollary 1.97 (Factorization lemma) Let $(\Omega', \mathcal{A}')$ be a measurable space and let Ω be a nonempty set. Let $f : \Omega \rightarrow \Omega'$ be a map. A map $g : \Omega \rightarrow \overline{\mathbb{R}}$ is $\sigma(f) - \mathcal{B}(\overline{\mathbb{R}})$ measurable if and only if there is a measurable map $\varphi : (\Omega', \mathcal{A}') \rightarrow (\overline{\mathbb{R}}, \mathcal{B}(\overline{\mathbb{R}}))$ such that $g = \varphi \circ f$.

Proof " $\Leftarrow$ " If φ is measurable and $g = \varphi \circ f$, then g is measurable by Theorem 1.80.

" $\Rightarrow$ " Now assume that g is $\sigma(f) - \mathcal{B}(\overline{\mathbb{R}})$ -measurable. First consider the case $g \geq 0$. Then there exist measurable sets $A_1, A_2, \dots \in \sigma(f)$ as well as numbers $\alpha_1, \alpha_2, \dots \in [0, \infty)$ such that $g = \sum_{n=1}^{\infty} \alpha_n \mathbb{1}_{A_n}$. By the definition of $\sigma(f)$, for any $n \in \mathbb{N}$ there is a set $B_n \in \mathcal{A}'$ such that $f^{-1}(B_n) = A_n$; that is, such that $\mathbb{1}_{A_n} = \mathbb{1}_{B_n} \circ f$. Define $\varphi : \Omega' \rightarrow \overline{\mathbb{R}}$ by

$$\varphi = \sum_{n=1}^{\infty} \alpha_n \mathbb{1}_{B_n}$$

Clearly, φ is $\mathcal{A}' - \mathcal{B}(\overline{\mathbb{R}})$ -measurable and $g = \varphi \circ f$.

Now drop the assumption that g is nonnegative. Then there exist measurable maps φ^- and φ^+ such that $g^- = \varphi^- \circ f$ and $g^+ = \varphi^+ \circ f$. Hence $\varphi := \varphi^+ - \varphi^-$ does the trick.

A measurable map transports a measure from one space to another.

Definition 1.98 (Image measure) Let $(\Omega, \mathcal{A})$ and $(\Omega', \mathcal{A}')$ be measurable spaces and let μ be a measure on $(\Omega, \mathcal{A})$. Further, let $X : (\Omega, \mathcal{A}) \rightarrow (\Omega', \mathcal{A}')$ be measurable. The image measure of μ under the map X is the measure $\mu \circ X^{-1}$ on $(\Omega', \mathcal{A}')$ that is defined by

$$\mu \circ X^{-1} : \mathcal{A}' \rightarrow [0, \infty], \quad A' \mapsto \mu(X^{-1}(A')).$$

Example 1.99 Let μ be a measure on $\mathbb{Z}^2$ and let $X : \mathbb{Z}^2 \rightarrow \mathbb{Z}, (x, y) \mapsto x + y$. Then

$$\mu \circ X^{-1}(\{x\}) = \sum_{y \in \mathbb{Z}} \mu(\{(x - y, y)\}).$$

Example 1.100 Let $L : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear bijection and let λ be the Lebesgue measure on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$. Then $\lambda \circ L^{-1} = |\det(L)|^{-1} \lambda$. This is clear since for any $a, b \in \mathbb{R}^n$ with $a < b$, the parallelepiped $L^{-1}((a, b])$ has volume $|\det(L^{-1})| \times \prod_{i=1}^n (b_i - a_i)$.

As a generalization of the last example, we state without proof the transformation formula for measures with continuous densities under differentiable maps. The proof can be found in textbooks on calculus.

Theorem 1.101 (Transformation formula in $\mathbb{R}^n$) Let μ be a measure on $\mathbb{R}^n$ that has a continuous (or piecewise continuous) density $f : \mathbb{R}^n \rightarrow [0, \infty)$. That is,

$$\mu((-\infty, x]) = \int_{-\infty}^{x_1} dt_1 \dots \int_{-\infty}^{x_n} dt_n f(t_1, \dots, t_n) \quad \text{for all } x \in \mathbb{R}^n$$

Let $A \subset \mathbb{R}^n$ be an open or a closed subset of $\mathbb{R}^n$ with $\mu(\mathbb{R}^n \setminus A) = 0$. Further, let $B \subset \mathbb{R}^n$ be open or closed. Finally, assume that $\varphi : A \rightarrow B$ is a continuously differentiable bijection with derivative φ' . Then the image measure $\mu \circ \varphi^{-1}$ has the density

$$f_{\varphi}(x) = \begin{cases} \frac{f(\varphi^{-1}(x))}{|\det(\varphi'(\varphi^{-1}(x)))|}, & \text{if } x \in B, \\ 0, & \text{if } x \in \mathbb{R}^n \setminus B. \end{cases}$$

Exercise 1.4.1 Let $f : \mathbb{R} \rightarrow \mathbb{R}, x \mapsto |x|$. Show that a Borel measurable map $g : \mathbb{R} \rightarrow \mathbb{R}$ is $\sigma(f) = f^{-1}(\mathcal{B}(\mathbb{R}))$ -measurable if and only if g is even.

Exercise 1.4.2 Let $(\Omega, \mathcal{A}, \mu)$ be a measure space and let $f : \Omega \rightarrow \mathbb{R}$ be measurable. Assume that $g : \Omega \rightarrow \mathbb{R}$ fulfills $g = f\mu$ -almost everywhere. Show that g need not be measurable.

Exercise 1.4.3 Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be differentiable with derivative f' . Show that f' is $\mathcal{B}(\mathbb{R}) - \mathcal{B}(\mathbb{R})$ -measurable.

Exercise 1.4.4 (Compare Examples 1.40 and 1.63.) Let $\Omega = \{0, 1\}^{\mathbb{N}}$ and let $\mathcal{A} = (2^{\{0, 1\}})^{\otimes \mathbb{N}}$ be the σ -algebra generated by the cylinder sets

$$\{[\omega_1, \dots, \omega_n] : n \in \mathbb{N}, \omega_1, \dots, \omega_n \in \{0, 1\}\}.$$

Further, let $\mu = (\frac{1}{2}\delta_0 + \frac{1}{2}\delta_1)^{\otimes \mathbb{N}}$ be the Bernoulli measure on Ω with equal weights on 0 and 1. For all $n \in \mathbb{N}$, let $X_n : \Omega \rightarrow \{0, 1\}, \omega \mapsto \omega_n$ be the n th coordinate map. Finally, let

$$U(\omega) = \sum_{n=1}^{\infty} X_n(\omega) 2^{-n} \quad \text{for } \omega \in \Omega$$

- (i) Show that $\mathcal{A} = \sigma(X_n : n \in \mathbb{N})$.
- (ii) Show that U is $\mathcal{A} - \mathcal{B}([0, 1])$ -measurable.
- (iii) Determine the image measure $\mu \circ U^{-1}$ on $([0, 1], \mathcal{B}([0, 1]))$.
- (iv) Determine an $\Omega_0 \in \mathcal{A}$ such that $\tilde{U} := U|_{\Omega_0}$ is bijective.
- (v) Show that $\tilde{U}^{-1}$ is $\mathcal{B}([0, 1]) - \mathcal{A}|_{\Omega_0}$ -measurable.
- (vi) Give an interpretation of the map $X_n \circ \tilde{U}^{-1}$.

Exercise 1.4.5 (Lusin's theorem) Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a Borel measurable map. Show that for any $\varepsilon > 0$, there exists a closed set $C \subset \mathbb{R}$ with $\lambda(\mathbb{R} \setminus C) < \varepsilon$ such that the restriction $f|_C$ of f to C is continuous. (Note: Clearly, this does not mean that f would be continuous in every point $x \in C$.)

Hint: Use the inner regularity of Lebesgue measure λ (Remark 1.67) to show the assertion first for indicator functions. Next construct a sequence of maps that approximates f uniformly on a suitable set C .

5.1.5 1.5 Random Variables

The fundamental idea of modern probability theory is to model one or more random experiments as a probability space $(\Omega, \mathcal{A}, \mathbf{P})$. The sets $A \in \mathcal{A}$ are called events. In most cases, the events of Ω are not observed directly. Rather, the observations are aspects of the single experiments that are coded as measurable maps from Ω to a space of possible observations. In short, every random observation (the technical term is random variable) is a measurable map. The probabilities of the possible random observations will be described in terms of the distribution of the corresponding random variable, which is the image measure of $\mathbf{P}$ under X . Hence we have to develop a calculus to determine the distributions of, for example, sums of random variables.

Definition 1.102 (Random variables) Let $(\Omega', \mathcal{A}')$ be a measurable space and let $X : \Omega \rightarrow \Omega'$ be measurable.

- (i) X is called a random variable with values in $(\Omega', \mathcal{A}')$. If $(\Omega', \mathcal{A}') = (\mathbb{R}, \mathcal{B}(\mathbb{R}))$, then X is called a real random variable or simply a random variable.
- (ii) For $A' \in \mathcal{A}'$, we denote $\{X \in A'\} := X^{-1}(A')$ and $\mathbf{P}[X \in A'] := \mathbf{P}[X^{-1}(A')]$. In particular, we let $\{X \geq 0\} := X^{-1}([0, \infty))$ and define $\{X \leq b\}$ similarly and so on.

Definition 1.103 (Distributions) Let X be a random variable.

- (i) The probability measure $\mathbf{P}_X := \mathbf{P} \circ X^{-1}$ is called the distribution of X .
- (ii) For a real random variable X , the map $F_X : x \mapsto \mathbf{P}[X \leq x]$ is called the distribution function of X (or, more accurately, of $\mathbf{P}_X$). We write $X \sim \mu$ if $\mu = \mathbf{P}_X$ and say that X has distribution μ .
- (iii) A family $(X_i)_{i \in I}$ of random variables is called identically distributed if $\mathbf{P}_{X_i} = \mathbf{P}_{X_j}$ for all $i, j \in I$. We write $X \stackrel{\mathcal{D}}{=} Y$ if $\mathbf{P}_X = \mathbf{P}_Y$ ($\mathcal{D}$ for distribution).

Theorem 1.104 For any distribution function F , there exists a real random variable X with $F_X = F$.

Proof We explicitly construct a probability space $(\Omega, \mathcal{A}, \mathbf{P})$ and a random variable $X : \Omega \rightarrow \mathbb{R}$ such that $F_X = F$.

The simplest choice would be $(\Omega, \mathcal{A}) = (\mathbb{R}, \mathcal{B}(\mathbb{R}))$, $X : \mathbb{R} \rightarrow \mathbb{R}$ the identity map and $\mathbf{P}$ the Lebesgue-Stieltjes measure with distribution function F (see Example 1.56).

A more instructive approach is based on first constructing, independently of F , a sort of standard probability space on which we define a random variable with uniform distribution on $(0, 1)$. In a second step, this random variable will be transformed by applying the inverse map F^{-1} : Let $\Omega := (0, 1)$, $\mathcal{A} := \mathcal{B}(\mathbb{R})|_{\Omega}$ and let $\mathbf{P}$ be the Lebesgue measure on $(\Omega, \mathcal{A})$ (see Example 1.74). Define the left continuous inverse of F :

$$F^{-1}(t) := \inf\{x \in \mathbb{R} : F(x) \geq t\} \quad \text{for } t \in (0, 1).$$

Then

$$F^{-1}(t) \leq x \iff t \leq F(x).$$

In particular, $\{t : F^{-1}(t) \leq x\} = (0, F(x)] \cap (0, 1)$; hence $F^{-1} : (\Omega, \mathcal{A}) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$ is measurable and

$$\mathbf{P}[\{t : F^{-1}(t) \leq x\}] = F(x).$$

Concluding, $X := F^{-1}$ is the random variable that we wanted to construct.

Example 1.105 We present some prominent distributions of real random variables X . These are some of the most important distributions in probability theory, and we will come back to these examples in many places.

(i) Let $p \in [0, 1]$ and $\mathbf{P}[X = 1] = p, \mathbf{P}[X = 0] = 1 - p$. Then $\mathbf{P}_X =: \text{Ber}_p$ is called the Bernoulli distribution with parameter p ; formally

$$\text{Ber}_p = (1 - p)\delta_0 + p\delta_1$$

Its distribution function is

$$F_X(x) = \begin{cases} 0, & x < 0 \\ 1 - p, & x \in [0, 1) \\ 1, & x \geq 1 \end{cases}$$

The distribution $\mathbf{P}_Y$ of $Y := 2X - 1$ is sometimes called the Rademacher distribution with parameter p ; formally $\text{Rad}_p = (1 - p)\delta_{-1} + p\delta_{+1}$. In particular, $\text{Rad}_{1/2}$ is called the Rademacher distribution.

(ii) Let $p \in [0, 1]$ and $n \in \mathbb{N}$, and let $X : \Omega \rightarrow \{0, \dots, n\}$ be such that

$$\mathbf{P}[X = k] = \binom{n}{k} p^k (1 - p)^{n-k}.$$

Then $\mathbf{P}_X =: b_{n,p}$ is called the binomial distribution with parameters n and p ; formally

$$b_{n,p} = \sum_{k=0}^n \binom{n}{k} p^k (1 - p)^{n-k} \delta_k.$$

(iii) Let $p \in (0, 1]$ and $X : \Omega \rightarrow \mathbb{N}_0$ with

$$\mathbf{P}[X = n] = p(1 - p)^n \quad \text{for any } n \in \mathbb{N}_0.$$

Then $\gamma_p := b_{1,p}^- := \mathbf{P}_X$ is called the geometric distribution ² with parameter p ; formally

$$\gamma_p = \sum_{n=0}^{\infty} p(1 - p)^n \delta_n.$$

Its distribution function is $F(x) = 1 - (1 - p)^{\lfloor x+1 \rfloor \vee 0}$ for $x \in \mathbb{R}$.

We can interpret $X+1$ as the waiting time for the first success in a series of independent random experiments, any of which yields a success with probability p . Indeed, let $\Omega = \{0, 1\}^{\mathbb{N}}$ and let $\mathbf{P}$ be the product measure $((1 - p)\delta_0 + p\delta_1)^{\otimes \mathbb{N}}$ (Theorem 1.64), as well as $\mathcal{A} = \sigma([\omega_1, \dots, \omega_n] : \omega_1, \dots, \omega_n \in \{0, 1\}, n \in \mathbb{N})$. Define

$$X(\omega) := \inf \{n \in \mathbb{N} : \omega_n = 1\} - 1,$$

where $\inf \emptyset = \infty$. Clearly, any map

$$X_n : \Omega \rightarrow \mathbb{R}, \quad \omega \mapsto \begin{cases} n-1, & \omega_n = 1, \\ \infty, & \omega_n = 0, \end{cases}$$

is $\mathcal{A} - \mathcal{B}(\overline{\mathbb{R}})$ -measurable. Thus also $X = \inf_{n \in \mathbb{N}} X_n$ is $\mathcal{A} - \mathcal{B}(\overline{\mathbb{R}})$ -measurable and is hence a random variable. Let $\omega^0 := (0, 0, \dots) \in \Omega$. Then $\mathbf{P}[X \geq n] = \mathbf{P}[[\omega_1^0, \dots, \omega_n^0]] = (1-p)^n$. Hence $\mathbf{P}[X = n] = \mathbf{P}[X \geq n] - \mathbf{P}[X \geq n+1] = (1-p)^n - (1-p)^{n+1} = p(1-p)^n$.

(iv) Let $r > 0$ (note that r need not be an integer) and let $p \in (0, 1]$. We denote by

$$b_{r,p}^- := \sum_{k=0}^{\infty} \binom{-r}{k} (-1)^k p^r (1-p)^k \delta_k \quad (1.17)$$

the negative binomial distribution or Pascal distribution with parameters r and p . (Here $\binom{x}{k} = \frac{x(x-1)\dots(x-k+1)}{k!}$ for $x \in \mathbb{R}$ and $k \in \mathbb{N}$ is the generalized binomial coefficient.) If $r \in \mathbb{N}$, then one can show as in the preceding example that $b_{r,p}^-$ is the distribution of the waiting time for the r th success in a series of random experiments. We come back to this in Example 3.4(iv).

(v) Let $\lambda \in [0, \infty)$ and let $X : \Omega \rightarrow \mathbb{N}_0$ be such that

$$\mathbf{P}[X = n] = e^{-\lambda} \frac{\lambda^n}{n!} \quad \text{for any } n \in \mathbb{N}_0.$$

Then $\mathbf{P}_X =: \text{Poi}_\lambda$ is called the Poisson distribution with parameter λ .

(vi) Consider an urn with $B \in \mathbb{N}$ black balls and $W \in \mathbb{N}$ white balls. Draw $n \in \mathbb{N}$ balls from the urn without replacement. A little bit of combinatorics shows that the probability of drawing exactly $b \in \{0, \dots, n\}$ black balls is given by the hypergeometric distribution with parameters $B, W, n \in \mathbb{N}$:

$$\text{Hyp}_{B,W;n}(\{b\}) = \frac{\binom{B}{b} \binom{W}{n-b}}{\binom{B+W}{n}} \quad \text{for } b \in \{0, \dots, n\}. \quad (1.18)$$

This generalizes easily to the situation of k colors and B_i balls of color $i = 1, \dots, k$. As above, we get that the probability of drawing out of n balls exactly b_i balls of color i for each $i = 1, \dots, k$ (with the restriction $b_1 + \dots + b_k = n$ and $b_i \leq B_i$ for all i) is given by the generalized hypergeometric distribution

$$\text{Hyp}_{B_1, \dots, B_k; n}(\{(b_1, \dots, b_k)\}) = \frac{\binom{B_1}{b_1} \dots \binom{B_k}{b_k}}{\binom{B_1 + \dots + B_k}{n}}. \quad (1.19)$$

(vii) Let $\mu \in \mathbb{R}, \sigma^2 > 0$ and let X be a real random variable with

$$\mathbf{P}[X \leq x] = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^x \exp\left(-\frac{(t-\mu)^2}{2\sigma^2}\right) dt \quad \text{for } x \in \mathbb{R}.$$

Then $\mathbf{P}_X =: \mathcal{N}_{\mu, \sigma^2}$ is called the Gaussian normal distribution with parameters μ and σ^2 . In particular, $\mathcal{N}_{0,1}$ is called the standard normal distribution.

(viii) Let $\theta > 0$ and let X be a nonnegative random variable such that

$$\mathbf{P}[X \leq x] = \mathbf{P}[X \in [0, x]] = \int_0^x \theta e^{-\theta t} dt \quad \text{for } x \geq 0$$

⁰² Warning: For some authors, the geometric distribution is shifted by one to the right; that is, it is a distribution on $\mathbb{N}$.

Then $\mathbf{P}_X$ is called the exponential distribution with parameter θ (in shorthand, $\exp_\theta$).

(ix) Let $\mu \in \mathbb{R}^d$ and let Σ be a positive definite symmetric $d \times d$ matrix. Let X be an $\mathbb{R}^d$ -valued random variable such that

$$\mathbf{P}[X \leq x] = \det(2\pi\Sigma)^{-1/2} \int_{(-\infty, x]} \exp\left(-\frac{1}{2}\langle t - \mu, \Sigma^{-1}(t - \mu) \rangle\right) \lambda^d(dt)$$

for $x \in \mathbb{R}^d$ (where $\langle \cdot, \cdot \rangle$ denotes the inner product in $\mathbb{R}^d$). Then $\mathbf{P}_X =: \mathcal{N}_{\mu, \Sigma}$ is the d -dimensional normal distribution with parameters μ and Σ .

Definition 1.106 If the distribution function $F : \mathbb{R}^n \rightarrow [0, 1]$ is of the form

$$F(x) = \int_{-\infty}^{x_1} dt_1 \dots \int_{-\infty}^{x_n} dt_n f(t_1, \dots, t_n) \quad \text{for } x = (x_1, \dots, x_n) \in \mathbb{R}^n$$

for some integrable function $f : \mathbb{R}^n \rightarrow [0, \infty)$, then f is called the density of the distribution.

Example 1.107

(i) Let $\theta, r > 0$ and let $\Gamma_{\theta, r}$ be the distribution on $[0, \infty)$ with density

$$x \mapsto \frac{\theta^r}{\Gamma(r)} x^{r-1} e^{-\theta x}$$

(Here Γ denotes the gamma function.) Then $\Gamma_{\theta, r}$ is called the Gamma distribution with scale parameter θ and shape parameter r .

(ii) Let $r, s > 0$ and let $\beta_{r, s}$ be the distribution on $[0, 1]$ with density

$$x \mapsto \frac{\Gamma(r+s)}{\Gamma(r)\Gamma(s)} x^{r-1} (1-x)^{s-1}$$

Then $\beta_{r, s}$ is called the Beta distribution with parameters r and s .

(iii) Let $a > 0$ and let Cau_a be the distribution on $\mathbb{R}$ with density

$$x \mapsto \frac{1}{a\pi} \frac{1}{1 + (x/a)^2}.$$

Then Cau_a is called the Cauchy distribution with parameter a .

Exercise 1.5.1 Use the identity $\binom{-n}{k}(-1)^k = \binom{n+k-1}{k}$ to deduce (1.17) by combinatorial means from its interpretation as a waiting time.

Exercise 1.5.2 Give an example of two normally distributed random variables X and Y such that (X, Y) is not (two-dimensional) normally distributed.

Exercise 1.5.3 Use the transformation formula (Theorem 1.101) to show the following statements.

(i) Let $X \sim \mathcal{N}_{\mu, \sigma^2}$ and let $a \in \mathbb{R} \setminus \{0\}$ and $b \in \mathbb{R}$. Then $(aX + b) \sim \mathcal{N}_{a\mu+b, a^2\sigma^2}$.

(ii) Let $X \sim \exp_\theta$ and $a > 0$. Then $aX \sim \exp_{\theta/a}$.

Exercise 1.5.4 Show that $F : \mathbb{R}^2 \rightarrow [0, 1]$ is the distribution function of a (uniquely determined) probability measure μ on $(\mathbb{R}^2, \mathcal{B}(\mathbb{R}^2))$ if and only if

(i) F is monotone increasing and right-continuous

(ii) $F(-x) \rightarrow 0$ and $F(x) \rightarrow 1$ as $x \rightarrow \infty$,

(iii) $F((y_1, y_2)) - F((y_1, x_2)) - F((x_1, y_2)) + F((x_1, x_2)) \geq 0$ for all $x_1 \leq y_1$ and $x_2 \leq y_2$.

Exercise 1.5.5

- (i) Let F and G be distribution functions on $\mathbb{R}$. Use Exercise 1.5.4 to show that $(x, y) \mapsto F(x) \wedge G(y)$ is a distribution function on $\mathbb{R}^2$.
- (ii) Give an example of two distribution functions F and G on $\mathbb{R}^2$ such that $(x, y) \mapsto F(x) \wedge G(y)$ is not a distribution function on $\mathbb{R}^4$.

Hint: First use the inclusion-exclusion formula (Theorem 1.33) to derive a criterion similar to that in Exercise 1.5.4(iii).

5.2 2 Independence

The measure theory from the preceding chapter is a linear theory that could not describe the dependence structure of events or random variables. We enter the realm of probability theory exactly at this point, where we define independence of events and random variables. Independence is a pivotal notion of probability theory, and the computation of dependencies is one of the theory's major tasks.

In the following, $(\Omega, \mathcal{A}, \mathbf{P})$ is a probability space and the sets $A \in \mathcal{A}$ are the events. As soon as constructing probability spaces has become routine, the concrete probability space will lose its importance and it will be only the random variables that will interest us. The bold font symbol $\mathbf{P}$ will then denote the universal object of a probability measure, and the probabilities $\mathbf{P}[\cdot]$ with respect to it will always be written in square brackets.

5.2.1 2.1 Independence of Events

We consider two events A and B as (stochastically) independent if the occurrence of A does not change the probability that B also occurs. Formally, we say that A and B are independent if

$$\mathbf{P}[A \cap B] = \mathbf{P}[A] \cdot \mathbf{P}[B] \quad (2.1)$$

Example 2.1 (Rolling a die twice) Consider the random experiment of rolling a die twice. Hence $\Omega = \{1, \dots, 6\}^2$ endowed with the σ -algebra $\mathcal{A} = 2^\Omega$ and the uniform distribution $\mathbf{P} = \mathcal{U}_\Omega$ (see Example 1.30(ii)).

- (i) Two events A and B should be independent, e.g., if A depends only on the outcome of the first roll and B depends only on the outcome of the second roll. Formally, we assume that there are sets $\tilde{A}, \tilde{B} \subset \{1, \dots, 6\}$ such that

$$A = \tilde{A} \times \{1, \dots, 6\} \quad \text{and} \quad B = \{1, \dots, 6\} \times \tilde{B}.$$

Now we check that A and B indeed fulfill (2.1). To this end, we compute $\mathbf{P}[A] = \frac{\#\tilde{A}}{36} = \frac{\#\tilde{A}}{6}$ and $\mathbf{P}[B] = \frac{\#\tilde{B}}{36} = \frac{\#\tilde{B}}{6}$. Furthermore,

$$\mathbf{P}[A \cap B] = \frac{\#(\tilde{A} \times \tilde{B})}{36} = \frac{\#\tilde{A}}{6} \cdot \frac{\#\tilde{B}}{6} = \mathbf{P}[A] \cdot \mathbf{P}[B]$$

- (ii) Stochastic independence can occur also in less obvious situations. For instance, let A be the event where the sum of the two rolls is odd,

$$A = \{(\omega_1, \omega_2) \in \Omega : \omega_1 + \omega_2 \in \{3, 5, 7, 9, 11\}\}$$

and let B be the event where the first roll gives at most a three

$$B = \{(\omega_1, \omega_2) \in \Omega : \omega_1 \in \{1, 2, 3\}\}.$$

Although it might seem that these two events are entangled in some way, they are stochastically independent. Indeed, it is easy to check that $\mathbf{P}[A] = \mathbf{P}[B] = \frac{1}{2}$ and $\mathbf{P}[A \cap B] = \frac{1}{4}$.

What is the condition for three events A_1, A_2, A_3 to be independent? Of course, any of the pairs (A_1, A_2) , (A_1, A_3) and (A_2, A_3) has to be independent. However, we have to make sure also that the simultaneous occurrence of A_1 and A_2 does not change the probability that A_3 occurs. Hence, it is not enough to consider pairs only.

Formally, we call three events A_1, A_2 and A_3 (stochastically) independent if

$$\mathbf{P}[A_i \cap A_j] = \mathbf{P}[A_i] \cdot \mathbf{P}[A_j] \quad \text{for all } i, j \in \{1, 2, 3\}, i \neq j, \quad (2.2)$$

and

$$\mathbf{P}[A_1 \cap A_2 \cap A_3] = \mathbf{P}[A_1] \cdot \mathbf{P}[A_2] \cdot \mathbf{P}[A_3] \quad (2.3)$$

Note that (2.2) does not imply (2.3) (and (2.3) does not imply (2.2)).

Example 2.2 (Rolling a die three times) We roll a die three times. Hence $\Omega = \{1, \dots, 6\}^3$ endowed with the discrete σ -algebra $\mathcal{A} = 2^\Omega$ and the uniform distribution $\mathbf{P} = \mathcal{U}_\Omega$ (see Example 1.30(ii)).

(i) If we assume that for any $i = 1, 2, 3$ the event A_i depends only on the outcome of the i th roll, then the events A_1, A_2 and A_3 are independent. Indeed, as in the preceding example, there are sets $\tilde{A}_1, \tilde{A}_2, \tilde{A}_3 \subset \{1, \dots, 6\}$ such that

$$\begin{aligned} A_1 &= \tilde{A}_1 \times \{1, \dots, 6\}^2 \\ A_2 &= \{1, \dots, 6\} \times \tilde{A}_2 \times \{1, \dots, 6\} \\ A_3 &= \{1, \dots, 6\}^2 \times \tilde{A}_3 \end{aligned}$$

The validity of (2.2) follows as in Example 2.1(i). In order to show (2.3), we compute

$$\mathbf{P}[A_1 \cap A_2 \cap A_3] = \frac{\#(\tilde{A}_1 \times \tilde{A}_2 \times \tilde{A}_3)}{216} = \prod_{i=1}^3 \frac{\#\tilde{A}_i}{6} = \prod_{i=1}^3 \mathbf{P}[A_i]$$

(ii) Consider now the events

$$\begin{aligned} A_1 &:= \{\omega \in \Omega : \omega_1 = \omega_2\}, \\ A_2 &:= \{\omega \in \Omega : \omega_2 = \omega_3\}, \\ A_3 &:= \{\omega \in \Omega : \omega_1 = \omega_3\}. \end{aligned}$$

Then $\#A_1 = \#A_2 = \#A_3 = 36$; hence $\mathbf{P}[A_1] = \mathbf{P}[A_2] = \mathbf{P}[A_3] = \frac{1}{6}$. Furthermore, $\#(A_i \cap A_j) = 6$ if $i \neq j$; hence $\mathbf{P}[A_i \cap A_j] = \frac{1}{36}$. Hence (2.2) holds. On the other hand, we have $\#(A_1 \cap A_2 \cap A_3) = 6$, thus $\mathbf{P}[A_1 \cap A_2 \cap A_3] = \frac{1}{36} \neq \frac{1}{6} \cdot \frac{1}{6} \cdot \frac{1}{6}$. Thus (2.3) does not hold and so the events A_1, A_2, A_3 are not independent.

In order to define independence of larger families of events, we have to request the validity of product formulas, such as (2.2) and (2.3), not only for pairs and triples but for all finite subfamilies of events. We thus make the following definition.

Definition 2.3 (Independence of events) Let I be an arbitrary index set and let $(A_i)_{i \in I}$ be an arbitrary family of events. The family $(A_i)_{i \in I}$ is called independent if for any finite subset $J \subset I$ the product formula holds:

$$\mathbf{P}\left[\bigcap_{j \in J} A_j\right] = \prod_{j \in J} \mathbf{P}[A_j].$$

The most prominent example of an independent family of infinitely many events is given by the perpetuated independent repetition of a random experiment.

Example 2.4 Let E be a finite set (the set of possible outcomes of the individual experiment) and let $(p_e)_{e \in E}$ be a probability vector on E . Equip (as in Theorem 1.64) the probability space $\Omega = E^{\mathbb{N}}$ with the σ -algebra $\mathcal{A} = \sigma(\{\omega_1, \dots, \omega_n\} : \omega_1, \dots, \omega_n \in E, n \in \mathbb{N}\})$ and with the product measure (or Bernoulli measure) $\mathbf{P} = (\sum_{e \in E} p_e \delta_e)^{\otimes \mathbb{N}}$; that is where $\mathbf{P}[\omega_1, \dots, \omega_n] = \prod_{i=1}^n p_{\omega_i}$. Let $\tilde{A}_i \subset E$ for any $i \in \mathbb{N}$, and let A_i be the event where $\tilde{A}_i$ occurs in the i th experiment; that is,

$$A_i = \left\{ \omega \in \Omega : \omega_i \in \tilde{A}_i \right\} = \bigcup_{(\omega_1, \dots, \omega_i) \in E^{i-1} \times \tilde{A}_i} [\omega_1, \dots, \omega_i]$$

Intuitively, the family $(A_i)_{i \in \mathbb{N}}$ should be independent if the definition of independence makes any sense at all. We check that this is indeed the case. Let $J \subset \mathbb{N}$ be finite and $n := \max J$. Formally, we define $B_j = A_j$ and $\tilde{B}_j = \tilde{A}_j$ for $j \in J$ and $B_j = \Omega$ and $\tilde{B}_j = E$ for $j \in \{1, \dots, n\} \setminus J$. Then

$$\begin{aligned} \mathbf{P} \left[\bigcap_{j \in J} A_j \right] &= \mathbf{P} \left[\bigcap_{j \in J} B_j \right] = \mathbf{P} \left[\bigcap_{j=1}^n B_j \right] \\ &= \sum_{e_1 \in \tilde{B}_1} \dots \sum_{e_n \in \tilde{B}_n} \prod_{j=1}^n p_{e_j} = \prod_{j=1}^n \left(\sum_{e \in \tilde{B}_j} p_e \right) = \prod_{j \in J} \left(\sum_{e \in \tilde{A}_j} p_e \right). \end{aligned}$$

This is true in particular for $\#J = 1$. Hence $\mathbf{P}[A_i] = \sum_{e \in \tilde{A}_i} p_e$ for all $i \in \mathbb{N}$, whence

$$\mathbf{P} \left[\bigcap_{j \in J} A_j \right] = \prod_{j \in J} \mathbf{P}[A_j]. \quad (2.4)$$

Since this holds for all finite $J \subset \mathbb{N}$, the family $(A_i)_{i \in \mathbb{N}}$ is independent.

If A and B are independent, then A^c and B also are independent since $\mathbf{P}[A^c \cap B] = \mathbf{P}[B] - \mathbf{P}[A \cap B] = \mathbf{P}[B] - \mathbf{P}[A]\mathbf{P}[B] = (1 - \mathbf{P}[A])\mathbf{P}[B] = \mathbf{P}[A^c]\mathbf{P}[B]$. We generalize this observation in the following theorem.

Theorem 2.5 Let I be an arbitrary index set and let $(A_i)_{i \in I}$ be a family of events. Define $B_i^0 = A_i$ and $B_i^1 = A_i^c$ for $i \in I$. Then the following three statements are equivalent.

- (i) The family $(A_i)_{i \in I}$ is independent.
- (ii) There is an $\alpha \in \{0, 1\}^I$ such that the family $(B_i^{\alpha_i})_{i \in I}$ is independent.
- (iii) For any $\alpha \in \{0, 1\}^I$, the family $(B_i^{\alpha_i})_{i \in I}$ is independent.

Proof This is left as an exercise.

Example 2.6 (Euler's prime number formula) The Riemann zeta function is defined by the Dirichlet series

$$\zeta(s) := \sum_{n=1}^{\infty} n^{-s} \quad \text{for } s \in (1, \infty)$$

Euler's prime number formula is a representation of the Riemann zeta function as an infinite product

$$\zeta(s) = \prod_{p \in \mathcal{P}} (1 - p^{-s})^{-1} \quad (2.5)$$

where $\mathcal{P} := \{p \in \mathbb{N} : p \text{ is prime}\}$.

We give a probabilistic proof for this formula. Let $\Omega = \mathbb{N}$, and for fixed $s > 1$ define $\mathbf{P}$ on 2^Ω by

$$\mathbf{P}[\{n\}] = \zeta(s)^{-1} n^{-s} \quad \text{for } n \in \mathbb{N}$$

Let $p\mathbb{N} = \{pn : n \in \mathbb{N}\}$ and $\mathcal{P}_n = \{p \in \mathcal{P} : p \leq n\}$. We consider $p\mathbb{N} \subset \Omega$ as an event. Note that $\mathbf{P}[p\mathbb{N}] = p^{-s}$ and that $(p\mathbb{N}, p \in \mathcal{P})$ is independent. Indeed, for $k \in \mathbb{N}$ and mutually distinct $p_1, \dots, p_k \in \mathcal{P}$, we have $\bigcap_{i=1}^k (p_i\mathbb{N}) = (p_1 \dots p_k)\mathbb{N}$. Thus

$$\begin{aligned} \mathbf{P} \left[\bigcap_{i=1}^k (p_i\mathbb{N}) \right] &= \sum_{n=1}^{\infty} \mathbf{P}[\{p_1 \dots p_k n\}] \\ &= \zeta(s)^{-1} (p_1 \dots p_k)^{-s} \sum_{n=1}^{\infty} n^{-s} \\ &= (p_1 \dots p_k)^{-s} = \prod_{i=1}^k \mathbf{P}[p_i\mathbb{N}] \end{aligned}$$

By Theorem 2.5, the family $((p\mathbb{N})^c, p \in \mathcal{P})$ is also independent, whence

$$\begin{aligned} \zeta(s)^{-1} &= \mathbf{P}[\{1\}] = \mathbf{P} \left[\bigcap_{p \in \mathcal{P}} (p\mathbb{N})^c \right] \\ &= \lim_{n \rightarrow \infty} \mathbf{P} \left[\bigcap_{p \in \mathcal{P}_n} (p\mathbb{N})^c \right] \\ &= \lim_{n \rightarrow \infty} \prod_{p \in \mathcal{P}_n} (1 - \mathbf{P}[p\mathbb{N}]) = \prod_{p \in \mathcal{P}} (1 - p^{-s}) \end{aligned}$$

This shows (2.5).

If we roll a die infinitely often, what is the chance that the face shows a six infinitely often? This probability should equal one. Otherwise there would be a last point in time when we see a six and after which the face only shows a number one to five. However, this is not very plausible.

Recall that we formalized the event where infinitely many of a series of events occur by means of the limes superior (see Definition 1.13). The following theorem confirms the conjecture mentioned above and also gives conditions under which we cannot expect that infinitely many of the events occur.

Theorem 2.7 (Borel-Cantelli lemma) Let $A_1, A_2, \dots$ be events and define $A^* = \limsup_{n \rightarrow \infty} A_n$.

(i) If $\sum_{n=1}^{\infty} \mathbf{P}[A_n] < \infty$, then $\mathbf{P}[A^*] = 0$. (Here $\mathbf{P}$ could be an arbitrary measure on $(\Omega, \mathcal{A})$.)

(ii) If $(A_n)_{n \in \mathbb{N}}$ is independent and $\sum_{n=1}^{\infty} \mathbf{P}[A_n] = \infty$, then $\mathbf{P}[A^*] = 1$.

Proof (i) $\mathbf{P}$ is upper semicontinuous and σ -subadditive; hence, by assumption,

$$\mathbf{P}[A^*] = \lim_{n \rightarrow \infty} \mathbf{P} \left[\bigcup_{m=n}^{\infty} A_m \right] \leq \lim_{n \rightarrow \infty} \sum_{m=n}^{\infty} \mathbf{P}[A_m] = 0$$

(ii) De Morgan's rule and the lower semicontinuity of $\mathbf{P}$ yield

$$\mathbf{P}[(A^*)^c] = \mathbf{P} \left[\bigcup_{m=1}^{\infty} \bigcap_{n=m}^{\infty} A_n^c \right] = \lim_{m \rightarrow \infty} \mathbf{P} \left[\bigcap_{n=m}^{\infty} A_n^c \right].$$

However, for every $m \in \mathbb{N}$ (since $\log(1-x) \leq -x$ for $x \in [0, 1]$), by upper continuity of $\mathbf{P}$

$$\begin{aligned} \mathbf{P} \left[\bigcap_{n=m}^{\infty} A_n^c \right] &= \lim_{N \rightarrow \infty} \mathbf{P} \left[\bigcap_{n=m}^N A_n^c \right] = \prod_{n=m}^{\infty} (1 - \mathbf{P}[A_n]) \\ &= \exp \left(\sum_{n=m}^{\infty} \log(1 - \mathbf{P}[A_n]) \right) \leq \exp \left(- \sum_{n=m}^{\infty} \mathbf{P}[A_n] \right) = 0 \end{aligned}$$

Example 2.8 We throw a die again and again and ask for the probability of seeing a six infinitely often. Hence $\Omega = \{1, \dots, 6\}^{\mathbb{N}}$, $\mathcal{A} = (2^{\{1, \dots, 6\}})^{\otimes \mathbb{N}}$ is the product σ algebra and $\mathbf{P} = \left(\sum_{e \in \{1, \dots, 6\}} \frac{1}{6} \delta_e \right)^{\otimes \mathbb{N}}$ is the Bernoulli measure (see Theorem 1.64). Furthermore, let $A_n = \{\omega \in \Omega : \omega_n = 6\}$ be the event where the n th roll shows a six. Then $A^* = \limsup_{n \rightarrow \infty} A_n$ is the event where we see a six infinitely often (see Remark 1.14). Furthermore, $(A_n)_{n \in \mathbb{N}}$ is an independent family with the property $\sum_{n=1}^{\infty} \mathbf{P}[A_n] = \sum_{n=1}^{\infty} \frac{1}{6} = \infty$. Hence the Borel-Cantelli lemma yields $\mathbf{P}[A^*] = 1$. $\diamond$

Example 2.9 We roll a die only once and define A_n for any $n \in \mathbb{N}$ as the event where in this one roll the face showed a six. Note that $A_1 = A_2 = A_3 = \dots$. Then $\sum_{n \in \mathbb{N}} \mathbf{P}[A_n] = \infty$; however, $\mathbf{P}[A^*] = \mathbf{P}[A_1] = \frac{1}{6}$. This shows that in Part (ii) of the Borel-Cantelli lemma, the assumption of independence is indispensable.

Example 2.10 Let $\Lambda \in (0, \infty)$ and $0 \leq \lambda_n \leq \Lambda$ for $n \in \mathbb{N}$. Let $X_n, n \in \mathbb{N}$, be Poisson random variables with parameters λ_n . Then

$$\mathbf{P}[X_n \geq n \text{ for infinitely many } n] = 0.$$

Indeed,

$$\begin{aligned} \sum_{n=1}^{\infty} \mathbf{P}[X_n \geq n] &= \sum_{n=1}^{\infty} \sum_{m=n}^{\infty} \mathbf{P}[X_n = m] = \sum_{m=1}^{\infty} \sum_{n=1}^m \mathbf{P}[X_n = m] \\ &= \sum_{m=1}^{\infty} \sum_{n=1}^m e^{-\lambda_n} \frac{\lambda_n^m}{m!} \leq \sum_{m=1}^{\infty} m \frac{\Lambda^m}{m!} = \Lambda e^{\Lambda} < \infty \end{aligned}$$

Note that in Theorem 2.7 in the case of independent events, only the probabilities $\mathbf{P}[A^*] = 0$ and $\mathbf{P}[A^*] = 1$ could show up. Thus the Borel-Cantelli lemma belongs to the class of so-called 0 – 1 laws. Later we will encounter more 0 – 1 laws (see, for example, Theorem 2.37).

Now we extend the notion of independence from families of events to families of classes of events.

Definition 2.11 (Independence of classes of events) Let I be an arbitrary index set and let $\mathcal{E}_i \subset \mathcal{A}$ for all $i \in I$. The family $(\mathcal{E}_i)_{i \in I}$ is called independent if, for any finite subset $J \subset I$ and any choice of $E_j \in \mathcal{E}_j, j \in J$, we have

$$\mathbf{P} \left[\bigcap_{j \in J} E_j \right] = \prod_{j \in J} \mathbf{P}[E_j]. \quad (2.6)$$

Example 2.12 As in Example 2.4, let $(\Omega, \mathcal{A}, \mathbf{P})$ be the product space of infinitely many repetitions of a random experiment whose possible outcomes e are the elements of the finite set E and have probabilities $p = (p_e)_{e \in E}$. For $i \in \mathbb{N}$, define

$$\mathcal{E}_i = \{ \{\omega \in \Omega : \omega_i \in A\} : A \subset E \}.$$

For any choice of sets $A_i \in \mathcal{E}_i, i \in \mathbb{N}$, the family $(A_i)_{i \in \mathbb{N}}$ is independent; hence $(\mathcal{E}_i)_{i \in \mathbb{N}}$ is independent.

Theorem 2.13

(i) Let I be finite, and for any $i \in I$ let $\mathcal{E}_i \subset \mathcal{A}$ with $\Omega \in \mathcal{E}_i$. Then

$$(\mathcal{E}_i)_{i \in I} \text{ is independent} \iff (2.6) \text{ holds for } J = I.$$

(ii) $(\mathcal{E}_i)_{i \in I}$ is independent $\iff ((\mathcal{E}_j)_{j \in J})$ is independent for all finite $J \subset I$.

(iii) If $(\mathcal{E}_i \cup \{\emptyset\})$ is $\cap$ -stable, then

$$(\mathcal{E}_i)_{i \in I} \text{ is independent} \iff (\sigma(\mathcal{E}_i))_{i \in I} \text{ is independent}.$$

(iv) Let K be an arbitrary set and let $(I_k)_{k \in K}$ be mutually disjoint subsets of I . If $(\mathcal{E}_i)_{i \in I}$ is independent, then $(\bigcup_{i \in I_k} \mathcal{E}_i)_{k \in K}$ is also independent.

Proof (i) " " This is trivial.

(i) " " For $J \subset I$ and $j \in I \setminus J$, choose $E_j = \Omega$.

(ii) This is trivial.

(iii) " " This is trivial.

(iii) " $\implies$ " Let $J \subset I$ be finite. We will show that for any two finite sets J and J' with $J \subset J' \subset I$,

$$\begin{aligned} \mathbf{P} \left[\bigcap_{i \in J'} E_i \right] &= \prod_{i \in J'} \mathbf{P} [E_i] \\ \text{for any choice } \begin{cases} E_i \in \sigma(\mathcal{E}_i), & i \in J, \\ E_i \in \mathcal{E}_i, & i \in J' \setminus J \end{cases} \end{aligned} \quad (2.7)$$

The case $J' = J$ is exactly the claim we have to show.

We carry out the proof of (2.7) by induction on $\#J$. For $\#J = 0$, the statement (2.7) holds by assumption of this theorem.

Now assume that (2.7) holds for every J with $\#J = n$ and for every finite $J' \supset J$. Fix such a J and let $j \in I \setminus J$. Choose $J' \supset \tilde{J} := J \cup \{j\}$. We show the validity of (2.7) with J replaced by $\tilde{J}$. Since $\#\tilde{J} = n + 1$, this verifies the induction step.

Let $E_i \in \sigma(\mathcal{E}_i)$ for any $i \in J$, and let $E_i \in \mathcal{E}_i$ for any $i \in J' \setminus (J \cup \{j\})$. Define two measures μ and ν on $(\Omega, \mathcal{A})$ by

$$\mu : E_j \mapsto \mathbf{P} \left[\bigcap_{i \in J'} E_i \right] \quad \text{and} \quad \nu : E_j \mapsto \prod_{i \in J'} \mathbf{P} [E_i]$$

By the induction hypothesis (2.7), we have $\mu(E_j) = \nu(E_j)$ for every $E_j \in \mathcal{E}_j \cup \{\emptyset, \Omega\}$. Since $\mathcal{E}_j \cup \{\emptyset\}$ is a π -system, Lemma 1.42 yields that $\mu(E_j) = \nu(E_j)$ for all $E_j \in \sigma(\mathcal{E}_j)$. That is, (2.7) holds with J replaced by $J \cup \{j\}$.

(iv) This is trivial, as (2.6) has to be checked only for $J \subset I$ with

$$\#(J \cap I_k) \leq 1 \quad \text{for any } k \in K.$$

5.2.2 2.2 Independent Random Variables

Now that we have studied independence of events, we want to study independence of random variables. Here also the definition ends up with a product formula. Formally, however, we can also define independence of random variables via independence of the σ -algebras they generate. This is the reason why we studied independence of classes of events in the last section.

Independent random variables allow for a rich calculus. For example, we can compute the distribution of a sum of two independent random variables by a simple convolution formula. Since we do not have a general notion of an integral at hand at this point, for the time being we restrict ourselves to presenting the convolution formula for integer-valued random variables only.

Let I be an arbitrary index set. For each $i \in I$, let $(\Omega_i, \mathcal{A}_i)$ be a measurable space and let $X_i : (\Omega, \mathcal{A}) \rightarrow (\Omega_i, \mathcal{A}_i)$ be a random variable with generated σ -algebra $\sigma(X_i) = X_i^{-1}(\mathcal{A}_i)$.

Definition 2.14 (Independent random variables) The family $(X_i)_{i \in I}$ of random variables is called independent if the family $(\sigma(X_i))_{i \in I}$ of σ -algebras is independent.

As a shorthand, we say that a family $(X_i)_{i \in I}$ is "i.i.d." (for "independent and identically distributed") if $(X_i)_{i \in I}$ is independent and if $\mathbf{P}_{X_i} = \mathbf{P}_{X_j}$ for all $i, j \in I$.

Remark 2.15

(i) Clearly, the family $(X_i)_{i \in I}$ is independent if and only if, for any finite set $J \subset I$ and any choice of $A_j \in \mathcal{A}_j, j \in J$, we have

$$\mathbf{P} \left[\bigcap_{j \in J} \{X_j \in A_j\} \right] = \prod_{j \in J} \mathbf{P}[X_j \in A_j].$$

The next theorem will show that it is enough to request the validity of such a product formula for A_j from an $\cap$ -stable generator of $\mathcal{A}_j$ only.

(ii) If $(\tilde{\mathcal{A}}_i)_{i \in I}$ is an independent family of σ -algebras and if each X_i is $\tilde{\mathcal{A}}_i - \mathcal{A}_i$ measurable, then $(X_i)_{i \in I}$ is independent. This is a direct consequence of the fact that $\sigma(X_i) \subset \tilde{\mathcal{A}}_i$.

(iii) For each $i \in I$, let $(\Omega'_i, \mathcal{A}'_i)$ be another measurable space and assume that $f_i : (\Omega_i, \mathcal{A}_i) \rightarrow (\Omega'_i, \mathcal{A}'_i)$ is a measurable map. If $(X_i)_{i \in I}$ is independent, then $(f_i \circ X_i)_{i \in I}$ is independent. This statement is a special case of (ii) since $f_i \circ X_i$ is $\sigma(X_i) - \mathcal{A}'_i$ -measurable (see Theorem 1.80).

Theorem 2.16 (Independent generators) For any $i \in I$, let $\mathcal{E}_i \subset \mathcal{A}_i$ be a π -system that generates $\mathcal{A}_i$. If $(X_i^{-1}(\mathcal{E}_i))_{i \in I}$ is independent, then $(X_i)_{i \in I}$ is independent.

Proof By Theorem 1.81, $X_i^{-1}(\mathcal{E}_i)$ is a π -system that generates the σ -algebra $X_i^{-1}(\mathcal{A}_i) = \sigma(X_i)$. Hence the statement follows from Theorem 2.13.

Example 2.17 Let E be a countable set and let $(X_i)_{i \in I}$ be random variables with values in $(E, 2^E)$. In this case, $(X_i)_{i \in I}$ is independent if and only if, for any finite $J \subset I$ and any choice of $x_j \in E, j \in J$,

$$\mathbf{P}[X_j = x_j \text{ for all } j \in J] = \prod_{j \in J} \mathbf{P}[X_j = x_j]$$

This is obvious since $\{\{x\} : x \in E\} \cup \{\emptyset\}$ is a π -system that generates 2^E , thus $\{X_i^{-1}(\{x_i\}) : x_i \in E\} \cup \{\emptyset\}$ is a π -system that generates $\sigma(X_i)$ (Theorem 1.81).

Example 2.18 Let E be a finite set and let $p = (p_e)_{e \in E}$ be a probability vector. Repeat a random experiment with possible outcomes $e \in E$ and probabilities p_e for $e \in E$ infinitely often (see Example 1.40 and Theorem 1.64). Let $\Omega = E^{\mathbb{N}}$ be the infinite product space and let $\mathcal{A}$ be the σ -algebra generated by the cylinder sets (see (1.8)). Let $\mathbf{P} = (\sum_{e \in E} p_e \delta_e)^{\otimes \mathbb{N}}$ be the Bernoulli measure. Further, for any $n \in \mathbb{N}$, let

$$X_n : \Omega \rightarrow E, \quad (\omega_m)_{m \in \mathbb{N}} \mapsto \omega_n,$$

be the projection on the n th coordinate. In other words: For any simple event $\omega \in \Omega$, $X_n(\omega)$ yields the result of the n th experiment. Then, by (2.4) (in Example 2.4), for $n \in \mathbb{N}$ and $x \in E^n$, we have

$$\begin{aligned}
& \mathbf{P}[X_j = x_j \text{ for all } j = 1, \dots, n] \\
&= \mathbf{P}[[x_1, \dots, x_n]] = \mathbf{P}\left[\bigcap_{j=1}^n X_j^{-1}(\{x_j\})\right] \\
&= \prod_{j=1}^n \mathbf{P}[X_j^{-1}(\{x_j\})] = \prod_{j=1}^n \mathbf{P}[X_j = x_j]
\end{aligned}$$

and $\mathbf{P}[X_j = x_j] = p_{x_j}$. By virtue of Theorem 2.13(i), this implies that the family $(X_1, \dots, X_n)$ is independent and hence, by Theorem 2.13(ii), $(X_n)_{n \in \mathbb{N}}$ is independent as well.

In particular, we have shown the following theorem.

Theorem 2.19 Let E be a finite set and let $(p_e)_{e \in E}$ be a probability vector on E . Then there exists a probability space $(\Omega, \mathcal{A}, \mathbf{P})$ and an independent family $(X_n)_{n \in \mathbb{N}}$ of E -valued random variables on $(\Omega, \mathcal{A}, \mathbf{P})$ such that $\mathbf{P}[X_n = e] = p_e$ for any $e \in E$.

Later we will see that the assumption that E is finite can be dropped. Also one can allow for different distributions in the respective factors. For the time being, however, this theorem gives us enough examples of interesting families of independent random variables.

Our next goal is to deduce simple criteria in terms of distribution functions and densities for checking whether a family of random variables is independent or not.

Definition 2.20 For any $i \in I$, let X_i be a real random variable. For any finite subset $J \subset I$, let

$$F_J := F_{(X_j)_{j \in J}} : \mathbb{R}^J \rightarrow [0, 1]$$

$$x \mapsto \mathbf{P}[X_j \leq x_j \text{ for all } j \in J] = \mathbf{P}\left[\bigcap_{j \in J} X_j^{-1}((-\infty, x_j])\right].$$

Then F_J is called the joint distribution function of $(X_j)_{j \in J}$. The probability measure $\mathbf{P}_{(X_j)_{j \in J}}$ on $\mathbb{R}^J$ is called the joint distribution of $(X_j)_{j \in J}$.

Theorem 2.21 A family $(X_i)_{i \in I}$ of real random variables is independent if and only if, for every finite $J \subset I$ and every $x = (x_j)_{j \in J} \in \mathbb{R}^J$,

$$F_J(x) = \prod_{j \in J} F_{\{j\}}(x_j) \quad (2.8)$$

Proof The class of sets $\{(-\infty, b], b \in \mathbb{R}\}$ is an $\cap$ -stable generator of the Borel σ -algebra $\mathcal{B}(\mathbb{R})$ (see Theorem 1.23). Equation (2.8) says that, for any choice of real numbers $(x_i)_{i \in I}$, the events $(X_i^{-1}((-\infty, x_i]))_{i \in I}$ are independent. Hence Theorem 2.16 yields the claim.

Corollary 2.22 In addition to the assumptions of Theorem 2.21, we assume that any F_J has a continuous density $f_J = f_{(X_j)_{j \in J}}$ (the joint density of $(X_j)_{j \in J}$). That is, there exists a continuous map $f_J : \mathbb{R}^J \rightarrow [0, \infty)$ such that

$$F_J(x) = \int_{-\infty}^{x_{j_1}} dt_1 \dots \int_{-\infty}^{x_{j_n}} dt_n f_J(t_1, \dots, t_n) \quad \text{for all } x \in \mathbb{R}^J$$

(where $J = \{j_1, \dots, j_n\}$). In this case, the family $(X_i)_{i \in I}$ is independent if and only if, for any finite $J \subset I$

$$f_J(x) = \prod_{j \in J} f_j(x_j) \quad \text{for all } x \in \mathbb{R}^J \quad (2.9)$$

Corollary 2.23 Let $n \in \mathbb{N}$ and let $\mu_1, \dots, \mu_n$ be probability measures on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$. Then there exists a probability space $(\Omega, \mathcal{A}, \mathbf{P})$ and an independent family of random variables $(X_i)_{i=1, \dots, n}$ on $(\Omega, \mathcal{A}, \mathbf{P})$ with $\mathbf{P}_{X_i} = \mu_i$ for each $i = 1, \dots, n$.

Proof Let $\Omega = \mathbb{R}^n$ and $\mathcal{A} = \mathcal{B}(\mathbb{R}^n)$. Let $\mathbf{P} = \bigotimes_{i=1}^n \mu_i$ be the product measure of the μ_i (see Theorem 1.61). Further, let $X_i : \mathbb{R}^n \rightarrow \mathbb{R}, (x_1, \dots, x_n) \mapsto x_i$ be the projection on the i th coordinate for each $i = 1, \dots, n$. Then, for any $i = 1, \dots, n$,

$$\begin{aligned} F_{\{i\}}(x) &= \mathbf{P}[X_i \leq x] = \mathbf{P}[\mathbb{R}^{i-1} \times (-\infty, x] \times \mathbb{R}^{n-i}] \\ &= \mu_i((-\infty, x]) \cdot \prod_{j \neq i} \mu_j(\mathbb{R}) = \mu_i((-\infty, x]) \end{aligned}$$

Hence indeed $\mathbf{P}_{X_i} = \mu_i$. Furthermore, for all $x_1, \dots, x_n \in \mathbb{R}$,

$$\begin{aligned} F_{\{1, \dots, n\}}((x_1, \dots, x_n)) &= \mathbf{P}\left[\bigcap_{i=1}^n (-\infty, x_i]\right] = \prod_{i=1}^n \mu_i((-\infty, x_i]) \\ &= \prod_{i=1}^n F_{\{i\}}(x_i) \end{aligned}$$

Hence Theorem 2.21 (and Theorem 2.13(i)) yields the independence of $(X_i)_{i=1, \dots, n}$.

Example 2.24 Let $X_1, \dots, X_n$ be independent exponentially distributed random variables with parameters $\theta_1, \dots, \theta_n \in (0, \infty)$. Then

$$F_{\{i\}}(x) = \int_0^x \theta_i e^{-\theta_i t} dt = 1 - e^{-\theta_i x} \quad \text{for } x \geq 0$$

and hence

$$F_{\{1, \dots, n\}}((x_1, \dots, x_n)) = \prod_{i=1}^n (1 - e^{-\theta_i x_i}).$$

Consider now the random variable $Y = \max(X_1, \dots, X_n)$. Then

$$\begin{aligned} F_Y(x) &= \mathbf{P}[X_i \leq x \text{ for all } i = 1, \dots, n] \\ &= F_{\{1, \dots, n\}}((x, \dots, x)) = \prod_{i=1}^n (1 - e^{-\theta_i x}) \end{aligned}$$

The distribution function of the random variable $Z := \min(X_1, \dots, X_n)$ has a nice closed form:

$$\begin{aligned} F_Z(x) &= 1 - \mathbf{P}[Z > x] \\ &= 1 - \mathbf{P}[X_i > x \text{ for all } i = 1, \dots, n] \\ &= 1 - \prod_{i=1}^n e^{-\theta_i x} = 1 - \exp(-(\theta_1 + \dots + \theta_n)x) \end{aligned}$$

In other words, Z is exponentially distributed with parameter $\theta_1 + \dots + \theta_n$.

Example 2.25 Let $\mu_i \in \mathbb{R}$ and $\sigma_i^2 > 0$ for $i \in I$. Let $(X_i)_{i \in I}$ be real random variables with joint density functions (for finite $J \subset I$)

$$f_J(x) = \prod_{j \in J} (2\pi\sigma_j^2)^{-\frac{1}{2}} \exp\left(-\sum_{j \in J} \frac{(x_j - \mu_j)^2}{2\sigma_j^2}\right) \quad \text{for } x \in \mathbb{R}^J$$

Then $(X_i)_{i \in I}$ is independent and X_i is normally distributed with parameters (μ_i, σ_i^2) .

For any finite $I = \{i_1, \dots, i_n\}$ (with mutually distinct $i_1, \dots, i_n$), the vector $Y = (X_{i_1}, \dots, X_{i_n})$ has the n -dimensional normal distribution with $\mu = \mu^I := (\mu_{i_1}, \dots, \mu_{i_n})$ and with $\Sigma = \Sigma^I$ the diagonal matrix with entries $\sigma_{i_1}^2, \dots, \sigma_{i_n}^2$ (see Example 1.105(ix)).

Theorem 2.26 Let K be an arbitrary set and $I_k, k \in K$, arbitrary mutually disjoint index sets. Define $I = \bigcup_{k \in K} I_k$.

If the family $(X_i)_{i \in I}$ is independent, then the family of σ -algebras $(\sigma(X_j, j \in I_k))_{k \in K}$ is independent.

Proof For $k \in K$, let

$$\mathcal{Z}_k = \left\{ \bigcap_{j \in I_k} A_j : A_j \in \sigma(X_j), \# \{j \in I_k : A_j \neq \Omega\} < \infty \right\}$$

be the semiring of finite-dimensional rectangular cylinder sets. Clearly, $\mathcal{Z}_k$ is a π system and $\sigma(\mathcal{Z}_k) = \sigma(X_j, j \in I_k)$. Hence, by Theorem 2.13(iii), it is enough to show that $(\mathcal{Z}_k)_{k \in K}$ is independent. By Theorem 2.13(ii), we can even assume that K is finite.

For $k \in K$, let $B_k \in \mathcal{Z}_k$ and $J_k \subset I_k$ be finite with $B_k = \bigcap_{j \in J_k} A_j$ for certain $A_j \in \sigma(X_j)$. Define $J = \bigcup_{k \in K} J_k$. Then

$$\mathbf{P} \left[\bigcap_{k \in K} B_k \right] = \mathbf{P} \left[\bigcap_{j \in J} A_j \right] = \prod_{j \in J} \mathbf{P}[A_j] = \prod_{k \in K} \prod_{j \in J_k} \mathbf{P}[A_j] = \prod_{k \in K} \mathbf{P}[B_k]$$

Example 2.27 If $(X_n)_{n \in \mathbb{N}}$ is an independent family of real random variables, then also $(Y_n)_{n \in \mathbb{N}} = (X_{2n} - X_{2n-1})_{n \in \mathbb{N}}$ is independent. Indeed, for any $n \in \mathbb{N}$, the random variable Y_n is $\sigma(X_{2n}, X_{2n-1})$ -measurable by Theorem 1.91, and $(\sigma(X_{2n}, X_{2n-1}))_{n \in \mathbb{N}}$ is independent by Theorem 2.26.

Example 2.28 Let $(X_{m,n})_{(m,n) \in \mathbb{N}^2}$ be an independent family of Bernoulli random variables with parameter $p \in (0, 1)$. Define the waiting time for the first "success" in the m th row of the matrix $(X_{m,n})_{m,n}$ by

$$Y_m := \inf \{n \in \mathbb{N} : X_{m,n} = 1\} - 1$$

Then $(Y_m)_{m \in \mathbb{N}}$ are independent geometric random variables with parameter p (see Example 1.105(iii)). Indeed,

$$\{Y_m \leq k\} = \bigcup_{l=1}^{k+1} \{X_{m,l} = 1\} \in \sigma(X_{m,l}, l = 1, \dots, k+1) \subset \sigma(X_{m,l}, l \in \mathbb{N})$$

Hence Y_m is $\sigma(X_{m,l}, l \in \mathbb{N})$ -measurable and thus $(Y_m)_{m \in \mathbb{N}}$ is independent. Furthermore,

$$\mathbf{P}[Y_m > k] = \mathbf{P}[X_{m,l} = 0, l = 1, \dots, k+1] = \prod_{l=1}^{k+1} \mathbf{P}[X_{m,l} = 0] = (1-p)^{k+1}$$

Concluding, we get $\mathbf{P}[Y_m = k] = \mathbf{P}[Y_m > k-1] - \mathbf{P}[Y_m > k] = p(1-p)^k$.

Definition 2.29 (Convolution) Let μ and ν be probability measures on $(\mathbb{Z}, 2^{\mathbb{Z}})$. The convolution $\mu * \nu$ is defined as the probability measure on $(\mathbb{Z}, 2^{\mathbb{Z}})$ such that

$$(\mu * \nu)(\{n\}) = \sum_{m=-\infty}^{\infty} \mu(\{m\})\nu(\{n-m\})$$

We define the n th convolution power recursively by $\mu^{*1} = \mu$ and

$$\mu^{*(n+1)} = \mu^{*n} * \mu.$$

Remark 2.30 The convolution is a symmetric operation: $\mu * \nu = \nu * \mu$. $\square$

Theorem 2.31 If X and Y are independent $\mathbb{Z}$ -valued random variables, then $\mathbf{P}_{X+Y} = \mathbf{P}_X * \mathbf{P}_Y$.

Proof For any $n \in \mathbb{Z}$,

$$\begin{aligned}
 \mathbf{P}_{X+Y}(\{n\}) &= \mathbf{P}[X + Y = n] \\
 &= \mathbf{P} \left[\biguplus_{m \in \mathbb{Z}} (\{X = m\} \cap \{Y = n - m\}) \right] \\
 &= \sum_{m \in \mathbb{Z}} \mathbf{P}[\{X = m\} \cap \{Y = n - m\}] \\
 &= \sum_{m \in \mathbb{Z}} \mathbf{P}_X[\{m\}] \mathbf{P}_Y[\{n - m\}] = (\mathbf{P}_X * \mathbf{P}_Y)[\{n\}].
 \end{aligned}$$

□

Owing to the last theorem, it is natural to define the convolution of two probability measures on $\mathbb{R}^n$ (or more generally on an Abelian group) as the distribution of the sum of two independent random variables with the corresponding distributions. Later we will encounter a different (but equivalent) definition that will, however, rely on the notion of an integral that is not yet available to us at this point (see Definition 14.17).

Definition 2.32 (Convolution of measures) Let μ and ν be probability measures on $\mathbb{R}^n$ and let X and Y be independent random variables with $\mathbf{P}_X = \mu$ and $\mathbf{P}_Y = \nu$. We define the convolution of μ and ν as $\mu * \nu = \mathbf{P}_{X+Y}$.

Recursively, we define the convolution powers μ^{*k} for all $k \in \mathbb{N}$ and let $\mu^{*0} = \delta_0$.

Example 2.33 Let X and Y be independent Poisson random variables with parameters μ and $\lambda \geq 0$. Then

$$\begin{aligned}
 \mathbf{P}[X + Y = n] &= e^{-\mu} e^{-\lambda} \sum_{m=0}^n \frac{\mu^m}{m!} \frac{\lambda^{n-m}}{(n-m)!} \\
 &= e^{-(\mu+\lambda)} \frac{1}{n!} \sum_{m=0}^n \binom{n}{m} \mu^m \lambda^{n-m} = e^{-(\mu+\lambda)} \frac{(\mu + \lambda)^n}{n!}.
 \end{aligned}$$

Hence $\text{Poi}_\mu * \text{Poi}_\lambda = \text{Poi}_{\mu+\lambda}$. □

Exercise 2.2.1 Let X and Y be independent random variables with $X \sim \exp_\theta$ and $Y \sim \exp_\rho$ for certain $\theta, \rho > 0$. Show that

$$\mathbf{P}[X < Y] = \frac{\theta}{\theta + \rho}$$

Exercise 2.2.2 (Box-Muller method) Let U and V be independent random variables that are uniformly distributed on $[0, 1]$. Define

$$X := \sqrt{-2 \log(U)} \cos(2\pi V) \quad \text{and} \quad Y := \sqrt{-2 \log(U)} \sin(2\pi V).$$

Show that X and Y are independent and $\mathcal{N}_{0,1}$ -distributed.

Hint: First compute the distribution of $\sqrt{-2 \log(U)}$ and then use the transformation formula (Theorem 1.101) as well as polar coordinates.

Exercise 2.2.3 (Multinomial distribution) Let $m \in \mathbb{N}$ and let $p = (p_1, \dots, p_m)$ be a probability vector on $\{1, \dots, m\}$. Let $X_1, \dots, X_n$ be independent random variables with values in $1, \dots, m$ and distribution p . We define an $\mathbb{N}_0^m$ -valued random variable $Y = (Y_1, \dots, Y_m)$ by

$$Y_i := \# \{k = 1, \dots, n : X_k = i\} \quad \text{for } i = 1, \dots, m.$$

Show that for $k = (k_1, \dots, k_m) \in \mathbb{N}_0^m$ with $k_1 + \dots + k_m = n$, we have

$$\mathbf{P}[Y = k] = \text{Mul}_{n,p}(\{k\}) := \binom{n}{k} p^k. \quad (2.10)$$

Here

$$\binom{n}{k} = \binom{n}{k_1, \dots, k_m} = \frac{n!}{k_1! \dots k_m!}$$

is the multinomial coefficient and $p^k = p_1^{k_1} \dots p_m^{k_m}$. The distribution $\text{Mul}_{n,p}$ on $\mathbb{N}_0^m$ is called multinomial distribution with parameters n and p .

5.2.3 2.3 Kolmogorov's 0-1 Law

With the Borel-Cantelli lemma, we have seen a first 0 – 1 law for independent events. We now come to another 0-1 law for independent events and for independent σ -algebras. To this end, we first introduce the notion of the tail σ -algebra.

Definition 2.34 (Tail σ -algebra) Let I be a countably infinite index set and let $(\mathcal{A}_i)_{i \in I}$ be a family of σ -algebras. Then

$$\mathcal{T}((\mathcal{A}_i)_{i \in I}) := \bigcap_{\substack{J \subset I \\ \#J < \infty}} \sigma \left(\bigcup_{j \in I \setminus J} \mathcal{A}_j \right)$$

is called the tail σ -algebra of $(\mathcal{A}_i)_{i \in I}$. If $(A_i)_{i \in I}$ is a family of events, then we define

$$\mathcal{T}((A_i)_{i \in I}) := \mathcal{T}((\{\emptyset, A_i, A_i^c, \Omega\})_{i \in I}).$$

If $(X_i)_{i \in I}$ is a family of random variables, then we define $\mathcal{T}((X_i)_{i \in I}) := \mathcal{T}((\sigma(X_i))_{i \in I})$.

The tail σ -algebra contains those events A whose occurrence is independent of any fixed finite subfamily of the X_i . To put it differently, for any finite subfamily of the X_i , we can change the values of the X_i arbitrarily without changing whether A occurs or not.

Theorem 2.35 Let $J_1, J_2, \dots$ be finite sets with $J_n \uparrow I$. Then

$$\mathcal{T}((\mathcal{A}_i)_{i \in I}) = \bigcap_{n=1}^{\infty} \sigma \left(\bigcup_{m \in I \setminus J_n} \mathcal{A}_m \right).$$

In the particular case $I = \mathbb{N}$, this reads $\mathcal{T}((\mathcal{A}_n)_{n \in \mathbb{N}}) = \bigcap_{n=1}^{\infty} \sigma \left(\bigcup_{m=n}^{\infty} \mathcal{A}_m \right)$.

The name "tail σ -algebra" is due to the interpretation of $I = \mathbb{N}$ as a set of times. As is made clear in the theorem, any event in $\mathcal{T}$ does not depend on the first finitely many time points.

Proof " $\subset$ " This is clear.

" $\supset$ " Let $J_n \subset I, n \in \mathbb{N}$, be finite sets with $J_n \uparrow I$. Let $J \subset I$ be finite. Then there exists an $N \in \mathbb{N}$ with $J \subset J_N$ and

$$\begin{aligned} \bigcap_{n=1}^{\infty} \sigma \left(\bigcup_{m \in I \setminus J_n} \mathcal{A}_m \right) &\subset \bigcap_{n=1}^N \sigma \left(\bigcup_{m \in I \setminus J_n} \mathcal{A}_m \right) \\ &= \sigma \left(\bigcup_{m \in I \setminus J_N} \mathcal{A}_m \right) \subset \sigma \left(\bigcup_{m \in I \setminus J} \mathcal{A}_m \right) \end{aligned}$$

The left-hand side does not depend on J . Hence we can form the intersection over all finite J and obtain

$$\bigcap_{n=1}^{\infty} \sigma \left(\bigcup_{m \in I \setminus J_n} \mathcal{A}_m \right) \subset \mathcal{T}((\mathcal{A}_i)_{i \in I})$$

Maybe at first glance it is not evident that there are any interesting events in the tail σ -algebra at all. It might not even be clear that we do not have $\mathcal{T} = \{\emptyset, \Omega\}$. Hence we now present simple examples of tail events and tail σ -algebra measurable random variables. In Section 2.4, we will study a more complex example.

Example 2.36

(i) Let $A_1, A_2, \dots$ be events. Then the events $A_* := \liminf_{n \rightarrow \infty} A_n$ and $A^* := \limsup_{n \rightarrow \infty} A_n$ are in $\mathcal{T}((A_n)_{n \in \mathbb{N}})$. Indeed, if we define $B_n := \bigcap_{m=n}^{\infty} A_m$ for $n \in \mathbb{N}$, then $B_n \uparrow A_*$ and $B_n \in \sigma((A_m)_{m \geq N})$ for any $n \geq N$. Thus $A_* \in \sigma((A_m)_{m \geq N})$ for any $N \in \mathbb{N}$ and hence $A_* \in \mathcal{T}((A_n)_{n \in \mathbb{N}})$. The case A^* is similar.

(ii) Let $(X_n)_{n \in \mathbb{N}}$ be a family of $\mathbb{R}$ -valued random variables. Then the maps $X_* := \liminf_{n \rightarrow \infty} X_n$ and $X^* := \limsup_{n \rightarrow \infty} X_n$ are $\mathcal{T}((X_n)_{n \in \mathbb{N}})$ -measurable. Indeed, if we define $Y_n := \sup_{m \geq n} X_m$, then for any $N \in \mathbb{N}$, the random variable $X^* = \inf_{n \geq 1} Y_n = \inf_{n \geq N} Y_n$ is $\mathcal{T}_N := \sigma(X_n, n \geq N)$ -measurable and hence also measurable with respect to $\mathcal{T}((X_n)_{n \in \mathbb{N}}) = \bigcap_{n=1}^{\infty} \mathcal{T}_n$.

The case X_* is similar.

(iii) Let $(X_n)_{n \in \mathbb{N}}$ be real random variables. Then the Cesàro limits

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n X_i \quad \text{and} \quad \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n X_i$$

are $\mathcal{T}((X_n)_{n \in \mathbb{N}})$ -measurable. In order to show this, choose $N \in \mathbb{N}$ and note that

$$X_* := \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n X_i = \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=N}^n X_i$$

is $\sigma((X_n)_{n \geq N})$ -measurable. Since this holds for any N , X_* is $\mathcal{T}((X_n)_{n \in \mathbb{N}})$ measurable. The case of the limes superior is similar.

Theorem 2.37 (Kolmogorov's 0-1 Law) Let I be a countably infinite index set and let $(\mathcal{A}_i)_{i \in I}$ be an independent family of σ -algebras. Then the tail σ -algebra is $\mathbf{P}$ -trivial, that is,

$$\mathbf{P}[A] \in \{0, 1\} \quad \text{for any } A \in \mathcal{T}((\mathcal{A}_i)_{i \in I}).$$

Proof It is enough to consider the case $I = \mathbb{N}$. For $n \in \mathbb{N}$, let

$$\mathcal{F}_n := \left\{ \bigcap_{k=1}^n A_k : A_1 \in \mathcal{A}_1, \dots, A_n \in \mathcal{A}_n \right\}.$$

Then $\mathcal{F} := \bigcup_{n=1}^{\infty} \mathcal{F}_n$ is a semiring and $\sigma(\mathcal{F}) = \sigma(\bigcup_{n \in \mathbb{N}} \mathcal{A}_n)$. Indeed, for any $n \in \mathbb{N}$ and $A_n \in \mathcal{A}_n$, we have $A_n \in \mathcal{F}$; hence $\sigma(\bigcup_{n \in \mathbb{N}} \mathcal{A}_n) \subset \sigma(\mathcal{F})$. On the other hand, we have $\mathcal{F}_m \subset \sigma(\bigcup_{n=1}^m \mathcal{A}_n) \subset \sigma(\bigcup_{n \in \mathbb{N}} \mathcal{A}_n)$ for any $m \in \mathbb{N}$; hence $\mathcal{F} \subset \sigma(\bigcup_{n \in \mathbb{N}} \mathcal{A}_n)$.

Let $A \in \mathcal{T}((\mathcal{A}_n)_{n \in \mathbb{N}})$ and $\varepsilon > 0$. By the approximation theorem for measures (Theorem 1.65), there exists an $N \in \mathbb{N}$ and mutually disjoint sets $F_1, \dots, F_N \in \mathcal{F}$ such that $\mathbf{P}[A \Delta (F_1 \cup \dots \cup F_N)] < \varepsilon$. Clearly, there is an $n \in \mathbb{N}$ such that $F_1, \dots, F_N \in \mathcal{F}_n$ and thus $F := F_1 \cup \dots \cup F_N \in \sigma(\mathcal{A}_1 \cup \dots \cup \mathcal{A}_n)$. Obviously, $A \in \sigma(\bigcup_{m=n+1}^{\infty} \mathcal{A}_m)$; hence A is independent of F . Thus

$$\varepsilon > \mathbf{P}[A \setminus F] = \mathbf{P}[A \cap (\Omega \setminus F)] = \mathbf{P}[A](1 - \mathbf{P}[F]) \geq \mathbf{P}[A](1 - \mathbf{P}[A] - \varepsilon)$$

Letting $\varepsilon \downarrow 0$ yields $0 = \mathbf{P}[A](1 - \mathbf{P}[A])$.

Corollary 2.38 Let $(A_n)_{n \in \mathbb{N}}$ be a sequence of independent events. Then

$$\mathbf{P} \left[\limsup_{n \rightarrow \infty} A_n \right] \in \{0, 1\} \quad \text{and} \quad \mathbf{P} \left[\liminf_{n \rightarrow \infty} A_n \right] \in \{0, 1\}.$$

Proof Essentially this is a simple conclusion of the Borel-Cantelli lemma. However, the statement can also be deduced from Kolmogorov's 0–1 law as limes superior and limes inferior are in the tail σ -algebra.

Corollary 2.39 Let $(X_n)_{n \in \mathbb{N}}$ be an independent family of $\overline{\mathbb{R}}$ -valued random variables. Then $X_* := \liminf_{n \rightarrow \infty} X_n$ and $X^* := \limsup_{n \rightarrow \infty} X_n$ are almost surely constant. That is, there exist $x_*, x^* \in \overline{\mathbb{R}}$ such that $\mathbf{P}[X_* = x_*] = 1$ and $\mathbf{P}[X^* = x^*] = 1$.

If all X_i are real-valued, then the Cesàro limits

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n X_i \quad \text{and} \quad \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n X_i$$

are also almost surely constant.

Proof Let $X_* := \liminf_{n \rightarrow \infty} X_n$. For any $x \in \overline{\mathbb{R}}$, we have $\{X_* \leq x\} \in \mathcal{T}((X_n)_{n \in \mathbb{N}})$; hence $\mathbf{P}[X_* \leq x] \in \{0, 1\}$. Define

$$x_* := \inf \{x \in \mathbb{R} : \mathbf{P}[X_* \leq x] = 1\} \in \overline{\mathbb{R}}$$

If $x_* = \infty$, then evidently

$$\mathbf{P}[X_* < \infty] = \lim_{n \rightarrow \infty} \mathbf{P}[X_* \leq n] = 0$$

If $x_* \in \mathbb{R}$, then

$$\mathbf{P}[X_* \leq x_*] = \lim_{n \rightarrow \infty} \mathbf{P} \left[X_* \leq x_* + \frac{1}{n} \right] = 1$$

and

$$\mathbf{P}[X_* < x_*] = \lim_{n \rightarrow \infty} \mathbf{P} \left[X_* \leq x_* - \frac{1}{n} \right] = 0$$

If $x_* = -\infty$, then

$$\mathbf{P}[X_* > -\infty] = \lim_{n \rightarrow \infty} \mathbf{P}[X_* > -n] = 0$$

The cases of the limes superior and the Cesàro limits are similar.

Exercise 2.3.1 Let $(X_n)_{n \in \mathbb{N}}$ be an independent family of $\text{Rad}_{1/2}$ random variables (i.e., $\mathbf{P}[X_n = -1] = \mathbf{P}[X_n = +1] = \frac{1}{2}$) and let $S_n = X_1 + \dots + X_n$ for any $n \in \mathbb{N}$. Show that $\limsup_{n \rightarrow \infty} S_n = \infty$ almost surely.

5.2.4 2.4 Example: Percolation

Consider the d -dimensional integer lattice $\mathbb{Z}^d$, where any point is connected to any of its $2d$ nearest neighbors by an edge. If $x, y \in \mathbb{Z}^d$ are nearest neighbors (that is, $\|x - y\|_2 = 1$), then we denote by $e = \langle x, y \rangle = \langle y, x \rangle$ the edge that connects x and y . Formally, the set of edges is a subset of the set of subsets of $\mathbb{Z}^d$ with two elements:

$$E = \{ \{x, y\} : x, y \in \mathbb{Z}^d \text{ with } \|x - y\|_2 = 1 \}.$$

Somewhat more generally, an undirected graph G is a pair $G = (V, E)$, where V is a set (the set of "vertices" or nodes) and $E \subset \{\{x, y\} : x, y \in V, x \neq y\}$ is a subset of the set of subsets of V of cardinality two (the set of edges or bonds).

Our intuitive understanding of an edge is a connection between two points x and y and not an (unordered) pair $\{x, y\}$. To stress this notion of a connection, we use a different symbol from the set brackets. That is, we denote the edge that connects x and y by $\langle x, y \rangle = \langle y, x \rangle$ instead of $\{x, y\}$.

Our graph (V, E) is the starting point for a stochastic model of a porous medium. We interpret the edges as tubes along which water can flow. However, we want the medium not to have a homogeneous structure, such as $\mathbb{Z}^d$, but an amorphous structure. In order to model this, we randomly destroy a certain fraction $1 - p$ of the tubes (with $p \in [0, 1]$ a parameter) and keep the others. Water can flow only through the remaining tubes. The destroyed tubes will be called "closed", the others "open". The fundamental question is: For which values of p is there a connected infinite system of tubes along which water can flow? The physical interpretation is that if we throw a block of the considered material into a bathtub, then the block will soak up water; that is, it will be wetted inside. If there is no infinite open component, then the water may wet only a thin layer at the surface. See Fig. 2.1 for a computer simulation of the percolation model.

We now come to a formal description of the model. Choose a parameter $p \in [0, 1]$ and an independent family of identically distributed random variables $(X_e^p)_{e \in E}$ with $X_e^p \sim \text{Ber}_p$; that is, $\mathbf{P}[X_e^p = 1] = 1 - \mathbf{P}[X_e^p = 0] = p$ for any $e \in E$. We define the set of open edges as

$$E^p := \{e \in E : X_e^p = 1\}. \quad (2.11)$$

Consequently, the edges in $E \setminus E^p$ are called closed. Hence we have constructed a (random) subgraph $(\mathbb{Z}^d, E^p)$ of $(\mathbb{Z}^d, E)$. We call $(\mathbb{Z}^d, E^p)$ a percolation model (more precisely, a model for bond percolation, in contrast to site percolation, where vertices can be open or closed). An (open) path (of length n) in this subgraph is a sequence $\pi = (x_0, x_1, \dots, x_n)$ of points in $\mathbb{Z}^d$ with $\langle x_{i-1}, x_i \rangle \in E^p$ for all $i = 1, \dots, n$. We say that two points $x, y \in \mathbb{Z}^d$ are connected by an open path if there is an $n \in \mathbb{N}$ and an open path $(x_0, x_1, \dots, x_n)$ with $x_0 = x$ and $x_n = y$. In this case, we write $x \longleftrightarrow_p y$. Note that " $\longleftrightarrow_p$ " is an equivalence relation; however, a random one, as it depends on the values of the random variables $(X_e^p)_{e \in E}$. For every $x \in \mathbb{Z}^d$, we define the (random) open cluster of x ; that is, the connected component of x in the graph $(\mathbb{Z}^d, E^p)$:

$$C^p(x) := \{y \in \mathbb{Z}^d : x \longleftrightarrow_p y\}. \quad (2.12)$$

Lemma 2.40 Let $x, y \in \mathbb{Z}^d$. Then $\mathbb{1}_{\{x \longleftrightarrow_p y\}}$ is a random variable. In particular, $\#C^p(x)$ is a random variable for any $x \in \mathbb{Z}^d$.

Proof We may assume $x = 0$. Let $f_{y,n} = 1$ if there exists an open path of length at most n that connects 0 to y , and $f_{y,n} = 0$ otherwise. Clearly, $f_{y,n} \uparrow \mathbb{1}_{\{0 \longleftrightarrow_p y\}}$ for $n \rightarrow \infty$; hence it suffices to show that each $f_{y,n}$ is measurable. Let $B_n := \{-n, -n+1, \dots, n-1, n\}^d$ and $E_n := \{e \in E : e \cap B_n \neq \emptyset\}$. Then $Y_n := (X_e^p : e \in E_n) : \Omega \rightarrow \{0, 1\}^{E_n}$ is measurable (with respect to $2^{\{0,1\}^{E_n}}$) by Theorem 1.90. However, $f_{y,n}$ is a function of Y_n , say $f_{y,n} = g_{y,n} \circ Y_n$ for some map $g_{y,n} : \{0, 1\}^{E_n} \rightarrow \{0, 1\}$. By the composition theorem for maps (Theorem 1.80), $f_{y,n}$ is measurable.

The additional statement holds since $\#C^p(x) = \sum_{y \in \mathbb{Z}^d} \mathbb{1}_{\{x \longleftrightarrow_p y\}}$.

Definition 2.41 We say that percolation occurs if there exists an infinitely large open cluster. We call

$$\begin{aligned}\psi(p) &:= \mathbf{P}[\text{there exists an infinite open cluster}] \\ &= \mathbf{P}\left[\bigcup_{x \in \mathbb{Z}^d} \{\#C^p(x) = \infty\}\right]\end{aligned}$$

the probability of percolation. We define

$$\theta(p) := \mathbf{P}[\#C^p(0) = \infty]$$

as the probability that the origin is in an infinite open cluster.

By the translation invariance of the lattice, we have

$$\theta(p) = \mathbf{P}[\#C^p(y) = \infty] \quad \text{for any } y \in \mathbb{Z}^d \quad (2.13)$$

The fundamental question is: How large are $\theta(p)$ and $\psi(p)$ depending on p ? We make the following simple observation.




Fig. 2.1 Percolation on a 15×15 grid, $p = 0.42$

Theorem 2.42 The map $[0, 1] \rightarrow [0, 1], p \mapsto \theta(p)$ is monotone increasing.

Proof Although the statement is intuitively so clear that it might not need a proof, we give a formal proof in order to introduce a technique called coupling. Let $p, p' \in [0, 1]$ with $p < p'$. Let $(Y_e)_{e \in E}$ be an independent family of random variables with $\mathbf{P}[Y_e \leq q] = q$ for any $e \in E$ and $q \in \{p, p', 1\}$. At this point, we could, for example, assume that $Y_e \sim \mathcal{U}_{[0,1]}$ is uniformly distributed on $[0,1]$. Since we have not yet shown the existence of an independent family with this distribution, we content ourselves with Y_e that assume only three values $\{p, p', 1\}$. Hence

$$\mathbf{P}[Y_e = q] = \begin{cases} p, & \text{if } q = p, \\ p' - p, & \text{if } q = p', \\ 1 - p', & \text{if } q = 1. \end{cases}$$

Such a family $(Y_e)_{e \in E}$ exists by Theorem 2.19. For $q \in \{p, p'\}$ and $e \in E$, we define

$$X_e^q := \begin{cases} 1, & \text{if } Y_e \leq q, \\ 0, & \text{else.} \end{cases}$$

Clearly, for any $q \in \{p, p'\}$, the family $(X_e^q)_{e \in E}$ is independent (see Remark 2.15(iii)) and $X_e^q \sim \text{Ber}_q$. Furthermore, $X_e^p \leq X_e^{p'}$ for any $e \in E$. The procedure of defining two families of random variables that are related in a specific way (here " $\leq$ ") on one probability space is called a coupling.

Clearly, $C^p(x) \subset C^{p'}(x)$ for any $x \in \mathbb{Z}^d$; hence $\theta(p) \leq \theta(p')$.

With the aid of Kolmogorov's 0-1 law, we can infer the following theorem.

Theorem 2.43 For any $p \in [0, 1]$, we have

$$\psi(p) = \begin{cases} 0, & \text{if } \theta(p) = 0, \\ 1, & \text{if } \theta(p) > 0. \end{cases}$$

Proof If $\theta(p) = 0$, then by (2.13)

$$\psi(p) \leq \sum_{y \in \mathbb{Z}^d} \mathbf{P}[\#C^p(y) = \infty] = \sum_{y \in \mathbb{Z}^d} \theta(p) = 0.$$

Now let $A = \bigcup_{y \in \mathbb{Z}^d} \{\#C^p(y) = \infty\}$. Clearly, A remains unchanged if we change the state of finitely many edges. That is, $A \in \sigma((X_e^p)_{e \in E \setminus F})$ for every finite $F \subset E$. Hence A is in the tail σ -algebra $\mathcal{T}((X_e^p)_{e \in E})$ by Theorem 2.35. Kolmogorov's 0-1 law (Theorem 2.37) implies that $\psi(p) = \mathbf{P}[A] \in \{0, 1\}$. If $\theta(p) > 0$, then $\psi(p) \geq \theta(p)$ implies $\psi(p) = 1$.

Due to the monotonicity, we can make the following definition.

Definition 2.44 The critical value p_c for percolation is defined as

$$\begin{aligned} p_c &= \inf\{p \in [0, 1] : \theta(p) > 0\} = \sup\{p \in [0, 1] : \theta(p) = 0\} \\ &= \inf\{p \in [0, 1] : \psi(p) = 1\} = \sup\{p \in [0, 1] : \psi(p) = 0\}. \end{aligned}$$

We come to the main theorem of this section.

Theorem 2.45 For $d = 1$, we have $p_c = 1$. For $d \geq 2$, we have $p_c(d) \in [\frac{1}{2d-1}, \frac{2}{3}]$.

Proof First consider $d = 1$ and $p < 1$. Let $A^- := \{X_{\langle n, n+1 \rangle}^p = 0 \text{ for some } n < 0\}$ and $A^+ := \{X_{\langle n, n+1 \rangle}^p = 0 \text{ for some } n > 0\}$. Let $A = A^- \cap A^+$. By the Borel-Cantelli lemma, we get $\mathbf{P}[A^-] = \mathbf{P}[A^+] = 1$. Hence $\theta(p) = \mathbf{P}[A^c] = 0$.

Now assume $d \geq 2$.

Lower bound. First we show $p_c \geq \frac{1}{2d-1}$. Clearly, for any $n \in \mathbb{N}$,

$$\mathbf{P}[\#C^p(0) = \infty] \leq \mathbf{P}[\text{there is an } x \in C^p(0) \text{ with } \|x\|_\infty = n]$$

We want to estimate the probability that there exists a point $x \in C^p(0)$ with distance n from the origin. Any such point is connected to the origin by a path without selfintersections π that starts at 0 and has length $m \geq n$. Let $\Pi_{0,m}$ be the set of such paths. Clearly, $\#\Pi_{0,m} \leq 2d \cdot (2d-1)^{m-1}$ since there are $2d$ choices for the first step and at most $2d-1$ choices for any further step. For any $\pi \in \Pi_{0,m}$, the probability that π uses only open edges is

$$\mathbf{P}[\pi \text{ is open}] = p^m.$$

Hence, for $p < \frac{1}{2d-1}$,

$$\begin{aligned}
 \theta(p) &\leq \sum_{m=n}^{\infty} \sum_{\pi \in \Pi_{0,m}} \mathbf{P}[\pi \text{ is open}] \\
 &\leq \frac{2d}{2d-1} \sum_{m=n}^{\infty} ((2d-1)p)^m \\
 &= \frac{2d}{(2d-1)(1-(2d-1)p)} ((2d-1)p)^n \xrightarrow{n \rightarrow \infty} 0
 \end{aligned}$$

We conclude that $p_c \geq \frac{1}{2d-1}$.

Upper bound. We can consider $\mathbb{Z}^d$ as a subset of $\mathbb{Z}^d \times \{0\} \subset \mathbb{Z}^{d+1}$. Hence, if percolation occurs for p in $\mathbb{Z}^d$, then it also occurs for p in $\mathbb{Z}^{d+1}$. Hence the corresponding critical values are ordered $p_c(d+1) \leq p_c(d)$.

Thus, it is enough to consider the case $d = 2$. Here we show $p_c \leq \frac{2}{3}$ by using a contour argument due to Peierls [127], originally designed for the Ising model of a ferromagnet, see Example 18.16 and (18.9).

For $N \in \mathbb{N}$, we define (compare (2.12) with $x = (i, 0)$)

$$C_N := \bigcup_{i=0}^N C^p((i, 0))$$

as the set of points that are connected (along open edges) to at least one of the points in $\{0, \dots, N\} \times \{0\}$. Due to the subadditivity of probability (and since $\mathbf{P}[\#C^p((i, 0)) = \infty] = \theta(p)$ for any $i \in \mathbb{Z}$), we have

$$\theta(p) = \frac{1}{N+1} \sum_{i=0}^N \mathbf{P}[\#C^p((i, 0)) = \infty] \geq \frac{1}{N+1} \mathbf{P}[\#C_N = \infty]$$

Now consider those closed contours in the dual graph $(\tilde{\mathbb{Z}}^2, \tilde{E})$ that surrounds C_N if $\#C_N < \infty$. Here the dual graph is defined by

$$\begin{aligned}
 \tilde{\mathbb{Z}}^2 &= \left(\frac{1}{2}, \frac{1}{2} \right) + \mathbb{Z}^2 \\
 \tilde{E} &= \left\{ \{x, y\} : x, y \in \tilde{\mathbb{Z}}^2, \|x - y\|_2 = 1 \right\}
 \end{aligned}$$

An edge $\tilde{e}$ in the dual graph $(\tilde{\mathbb{Z}}^2, \tilde{E})$ crosses exactly one edge e in $(\mathbb{Z}^2, E)$. We call $\tilde{e}$ open if e is open and closed otherwise. A circle γ is a self-intersection free




Fig. 2.2 Contour of the cluster $C_{-}\{5\}$

path in $(\tilde{\mathbb{Z}}^2, \tilde{E})$ that starts and ends at the same point. A contour of the set C_N is a minimal circle that surrounds C_N . Minimal means that the enclosed area is minimal (see Fig. 2.2). For $n \geq 2N$, let

$$\Gamma_n = \{\gamma : \gamma \text{ is a circle of length } n \text{ that surrounds } \{0, \dots, N\} \times \{0\}\}.$$

We want to deduce an upper bound for $\#\Gamma_n$. Let $\gamma \in \Gamma_n$ and fix one point of γ . For definiteness, choose the upper point $(m + \frac{1}{2}, \frac{1}{2})$ of the rightmost edge of γ that crosses the horizontal axis (in Fig. 2.2 this is the point $(5 + \frac{1}{2}, \frac{1}{2})$). Clearly, $m \geq N$ and $m \leq n$ since γ surrounds the origin. Starting from $(m + \frac{1}{2}, \frac{1}{2})$, for any further edge of γ , there are at most three possibilities. Hence

$$\#\Gamma_n \leq n \cdot 3^n$$

We say that γ is closed if it uses only closed edges (in $\tilde{E}$). A contour of C_N is automatically closed and has a length of at least $2N$. Hence for $p > \frac{2}{3}$

$$\begin{aligned} \mathbf{P}[\#C_N < \infty] &= \sum_{n=2N}^{\infty} \mathbf{P}[\text{there is a closed circle } \gamma \in \Gamma_n] \\ &\leq \sum_{n=2N}^{\infty} n \cdot (3(1-p))^n \xrightarrow{N \rightarrow \infty} 0 \end{aligned}$$

We conclude $p_c \leq \frac{2}{3}$. $\square$

In general, the value of p_c is not known and is extremely hard to determine. In the case of bond percolation on $\mathbb{Z}^2$, however, the exact value of p_c can be determined due to the self-duality of the planar graph $(\mathbb{Z}^2, E)$. (If $G = (V, E)$ is a planar graph; that is, a graph that can be embedded into $\mathbb{R}^2$ without self-intersections, then the vertex set of the dual graph is the set of faces of G . Two such vertices are connected by exactly one edge; that is, by the edge in E that separates the two faces. Evidently, the two-dimensional integer lattice is isomorphic to its dual graph. Note that the contour in Fig. 2.2 can be considered as a closed path in the dual graph.) We cite a theorem of Kesten [95].

Theorem 2.46 (Kesten 1980) For bond percolation in $\mathbb{Z}^2$, the critical value is $p_c = \frac{1}{2}$ and $\theta(p_c) = 0$.

Proof See, for example, the book of Grimmett [63, pp. 287ff].

It is conjectured that $\theta(p_c) = 0$ holds in any dimension $d \geq 2$. However, rigorous proofs are known only for $d = 2$ and $d \geq 19$ (see [67]).

Uniqueness of the Infinite Open Cluster*

Fix a p such that $\theta(p) > 0$. We saw that with probability one there is at least one infinite open cluster. Now we want to show that there is exactly one.

Denote by $N \in \{0, 1, \dots, \infty\}$ the (random) number of infinite open clusters.

Theorem 2.47 (Uniqueness of the infinite open cluster) For any $p \in [0, 1]$, we have $\mathbf{P}_p[N \leq 1] = 1$.

Proof This theorem was first proved by Aizenman, Kesten and Newman [2, 3]. Here we follow the proof of Burton and Keane [23] as described in [63, Section 8.2].

The cases $p = 1$ and $\theta(p) = 0$ (hence in particular the case $p = 0$) are trivial. Hence we assume now that $p \in (0, 1)$ and $\theta(p) > 0$.

Step 1. We first show that

$$\mathbf{P}_p[N = m] = 1 \quad \text{for some } m = 0, 1, \dots, \infty. \quad (2.14)$$

We need a 0 – 1 law similar to that of Kolmogorov. However, N is not measurable with respect to the tail σ -algebra. Hence we have to find a more subtle argument. Let $u_1 = (1, 0, \dots, 0)$ be the first unit vector in $\mathbb{Z}^d$. On the edge set E , define the translation $\tau : E \rightarrow E$ by $\tau(\langle x, y \rangle) = \langle x + u_1, y + u_1 \rangle$. Let

$$E_0 := \{ \langle (x_1, \dots, x_d), (y_1, \dots, y_d) \rangle \in E : x_1 = 0, y_1 \geq 0 \}$$

be the set of all edges in $\mathbb{Z}^d$ that either connect two points from $\{0\} \times \mathbb{Z}^{d-1}$ or one point of $\{0\} \times \mathbb{Z}^{d-1}$ with one point of $\{1\} \times \mathbb{Z}^{d-1}$. Clearly, the sets $(\tau^n(E_0), n \in \mathbb{Z})$ are disjoint and $E = \biguplus_{n \in \mathbb{Z}} \tau^n(E_0)$. Hence the random variables $Y_n := \left(X_{\tau^n(e)}^p \right)_{e \in E_0}$, $n \in \mathbb{Z}$, are independent and identically distributed (with values in $\{0, 1\}^{E_0}$). Define $Y = (Y_n)_{n \in \mathbb{Z}}$ and $\tau(Y) = (Y_{n+1})_{n \in \mathbb{Z}}$. Define $A_m \in \{0, 1\}^E$ by

$$\{Y \in A_m\} = \{N = m\}$$

Clearly, the value of N does not change if we shift all edges simultaneously. That is, $\{Y \in A_m\} = \{\tau(Y) \in A_m\}$. An event with this property is called invariant or shift invariant. Using an argument similar to that in the proof of Kolmogorov's 0-1 law, one can show that invariant events (defined by i.i.d. random variables) have probability either 0 or 1 (see Example 20.26 for a proof).

Step 2. We will show that

$$\mathbf{P}_p[N = m] = 0 \quad \text{for any } m \in \mathbb{N} \setminus \{1\} \quad (2.15)$$

Accordingly, let $m = 2, 3, \dots$. We assume that $\mathbf{P}[N = m] = 1$ and show that this leads to a contradiction.

For $L \in \mathbb{N}$, let $B_L := \{-L, \dots, L\}^d$ and denote by $E_L = \{e = \langle x, y \rangle \in E : x, y \in B_L\}$ the set of those edges with both vertices lying in B_L . For $i = 0, 1$, let $D_L^i := \{X_e^p = i \text{ for all } e \in E_L\}$. Let N_L^1 be the number of infinite open clusters if we consider all edges e in E_L as open (independently of the value of X_e^p). Similarly define N_L^0 where we consider all edges in E_L as closed. Since $\mathbf{P}_p[D_L^i] > 0$, and since $N = m$ almost surely, we have $N_L^i = m$ almost surely for $i = 0, 1$.

Let

$$A_L^2 := \bigcup_{x^1, x^2 \in B_L \setminus B_{L-1}} \{C^p(x^1) \cap C^p(x^2) = \emptyset\} \cap \{\#C^p(x^1) = \#C_p(x^2) = \infty\}$$

be the event where there exist two points on the boundary of B_L that lie in different infinite open clusters. Clearly, $A_L^2 \uparrow \{N \geq 2\}$ for $L \rightarrow \infty$.

Define $A_{L,0}^2$ in a similarly way to A_L^2 ; however, we now consider all edges $e \in E_L$ as closed, irrespective of whether $X_e^p = 1$ or $X_e^p = 0$. If A_L^2 occurs, then there are two points x^1, x^2 on the boundary of B_L such that for any $i = 1, 2$, there is an infinite self-intersection free open path π_{x^i} starting at x^i that avoids x^{3-i} . Hence $A_L^2 \subset A_{L,0}^2$. Now choose L large enough for $\mathbf{P}[A_{L,0}^2] > 0$.

If $A_{L,0}^2$ occurs and if we open all edges in B_L , then at least two of the infinite open clusters get connected by edges in B_L . Hence the total number of infinite open clusters decreases by at least one. We infer $\mathbf{P}_p[N_L^1 \leq N_L^0 - 1] \geq \mathbf{P}_p[A_{L,0}^2] > 0$, which leads to a contradiction.

Step 3. In Step 2, we have shown already that N does not assume a finite value larger than 1. Hence it remains to show that almost surely N does not assume the value ∞ . Indeed, we show that

$$\mathbf{P}_p[N \geq 3] = 0 \quad (2.16)$$

This part of the proof is the most difficult one. We assume that $\mathbf{P}_p[N \geq 3] > 0$ and show that this leads to a contradiction.

We say that a point $x \in \mathbb{Z}^d$ is a trifurcation point if

- x is in an infinite open cluster $C^p(x)$,
- there are exactly three open edges with endpoint x , and
- removing all of these three edges splits $C^p(x)$ into three mutually disjoint infinite open clusters.

By T we denote the set of trifurcation points, and let $T_L := T \cap B_L$. Let $r := \mathbf{P}_p[0 \in T]$. Due to translation invariance, we have $(\#B_L)^{-1} \mathbf{E}_p[\#T_L] = r$ for any L . (Here $\mathbf{E}_p[\#T_L]$ denotes the expected value of $\#T_L$, which we define formally in Chapter 5.) Let

$$A_L^3 := \bigcup_{x^1, x^2, x^3 \in B_L \setminus B_{L-1}} \left(\bigcap_{i \neq j} \{C^p(x^i) \cap C^p(x^j) = \emptyset\} \right) \cap \left(\bigcap_{i=1}^3 \{\#C^p(x^i) = \infty\} \right)$$

be the event where there are three points on the boundary of B_L that lie in different infinite open clusters. Clearly, $A_L^3 \uparrow \{N \geq 3\}$ for $L \rightarrow \infty$.

As for $A_{L,0}^2$, we define $A_{L,0}^3$ as the event where there are three distinct points on the boundary of B_L that lie in different infinite open clusters if we consider all edges in E_L as closed. As above, we have $A_L^3 \subset A_{L,0}^3$.

For three distinct points $x^1, x^2, x^3 \in B_L \setminus B_{L-1}$, let F_{x^1, x^2, x^3} be the event where for any $i = 1, 2, 3$, there exists an infinite self-intersection free open path π_{x^i} starting at x^i that uses only edges in $E^p \setminus E_L$ and that avoids the points $x^j, j \neq i$. Then

$$A_{L,0}^3 \subset \bigcup_{\substack{x^1, x^2, x^3 \in B_L \setminus B_{L-1} \\ \text{mutually distinct}}} F_{x^1, x^2, x^3}$$

Let L be large enough for $\mathbf{P}_p[A_{L,0}^3] \geq \mathbf{P}_p[N \geq 3]/2 > 0$. Choose three pairwise distinct points $x^1, x^2, x^3 \in B_L \setminus B_{L-1}$ with $\mathbf{P}_p[F_{x^1, x^2, x^3}] > 0$.

If F_{x^1, x^2, x^3} occurs, then we can find a point $y \in B_L$ that is the starting point of three mutually disjoint (not necessarily open) paths π_1, π_2 and π_3 that end at x^1, x^2 and x^3 . Let G_{y, x^1, x^2, x^3} be the event where in E_L exactly those edges are open that belong to these three paths (that is,

all other edges in E_L are closed). The events F_{x^1, x^2, x^3} and G_{y, x^1, x^2, x^3} are independent, and if both of them occur, then y is a trifurcation point. Hence

$$r = \mathbf{P}_p[y \in T] \geq \mathbf{P}_p[F_{x^1, x^2, x^3}] \cdot (p \wedge (1 - p))^{\#E_L} > 0$$

Now we show that r must equal 0, which contradicts the assumption $\mathbf{P}_p[N \geq 3] > 0$. Let K_L be the set of all edges which have at least one endpoint in B_L . We consider two edges in K_L as equivalent if there exists a path in B_L along open edges that does not hit any trifurcation point and which joins at least one endpoint of each of the two edges. We denote the equivalence relation by R and let $U_L = K_L/R$ be the set of equivalence classes. (Note that the three neighboring edges of a trifurcation point are in different equivalence classes.) We turn the set $H_L := U_L \cup T_L$ into a graph by considering two points $x \in T_L$ and $u \in U_L$ as neighbors if there exists an edge $k \in u$ which is incident to x . Note that each point $x \in T_L$ has exactly three neighbors which are in U_L . The points in U_L can be isolated (that is, without neighbors) or can be joined to arbitrarily many points in T_L but not in U_L .

A circle is a self-avoiding (finite) path that ends at its starting point. Note that the graph H_L has no circles. To show this assume there was a self-avoiding path $(h_0, h_1, \dots, h_n)$ starting and ending in some point $h_0 = h_n = x \in T_L$. Then $h_1, h_{n-1} \in U_L$ are distinct but connected in K^p even if we remove x . However, by the definition of the trifurcation point x , this is impossible. On the other hand, if there was a self-avoiding path $(g_0, \dots, g_m)$ starting and ending in some point $g_0 = g_m = u \in U_L$, then $(g_1, g_2, \dots, g_m, g_1)$ is a self-avoiding path starting and ending in $g_1 \in T_L$. However, we have just shown that such a path could not exist.

Write $\deg_{H_L}(h)$ for the degree of $h \in H_L$; that is, the number of neighbors of h in H_L . A point h with $\deg_{H_L}(h) = 1$ is called a leaf of H_L . Obviously, only points of U_L can be leaves. Let Z be a connected component of H_L that contains at least one point $x \in T_L$. Since Z is a tree (that is, it is connected and contains no circles), we have

$$\#Z - 1 = \frac{1}{2} \sum_{h \in Z} \deg_{H_L}(h)$$

Rearranging this formula yields an expression for the number of leaves:

$$\begin{aligned} \#\{u \in Z : \deg_{H_L}(u) = 1\} &= 2 + \sum_{h \in Z} (\deg_{H_L}(h) - 2)^+ \\ &\geq 2 + \#\{h \in Z : \deg_{H_L}(h) \geq 3\} \\ &= 2 + \#(Z \cap T_L) \end{aligned}$$

Summing over the connected components Z of H_L with at least one point in T_L , we obtain

$$\#\{u \in H_L : \deg_{H_L}(u) = 1\} \geq \#T_L.$$

Observe that any leaf $u \in H_L$ contains an edge that is incident to a point $x \in T_L$. Hence the edges of u lie in an infinite open cluster of K^p and there is at least one edge $k \in u$ incident to a point at the boundary $B_L \setminus B_{L-1}$ of B_L . For distinct leaves these are distinct points since the leaves belong to disjoint open clusters. Hence we get the bound

$$\#T_L \leq \#(B_L \setminus B_{L-1})$$

and thus

$$\frac{\#T_L}{\#B_L} \leq \frac{\#(B_L \setminus B_{L-1})}{\#B_L} \leq \frac{d}{L} \xrightarrow{L \rightarrow \infty} 0$$

Now $r = (\#B_L)^{-1} \mathbf{E}_p[\#T_L] \leq d/L$ implies $r = 0$. (Note that in the argument we used the notion of the expected value $\mathbf{E}_p[\#T_L]$ that will be formally introduced only in Chapter 5.)

5.3 3 Generating Functions

It is a fundamental principle of mathematics to map a class of objects that are of interest into a class of objects where computations are easier. This map can be one to one, as with linear maps and matrices, or it may map only some properties uniquely, as with matrices and determinants.

In probability theory, in the second category fall quantities such as the median, mean and variance of random variables. In the first category, we have characteristic functions, Laplace transforms and probability generating functions. These are useful mostly because addition of independent random variables leads to multiplication of the transforms. Before we introduce characteristic functions (and Laplace transforms) later in the book, we want to illustrate the basic idea with probability generating functions that are designed for $\mathbb{N}_0$ -valued random variables.

In the first section, we give the basic definitions and derive simple properties. The next two sections are devoted to two applications: The Poisson approximation theorem and a simple investigation of Galton-Watson branching processes.

5.3.1 3.1 Definition and Examples

Definition 3.1 (Probability generating function) Let X be an $\mathbb{N}_0$ -valued random variable. The (probability) generating function (p.g.f.) of $\mathbf{P}_X$ (or, loosely speaking, of X) is the map $\psi_{\mathbf{P}_X} = \psi_X$ defined by (with the understanding that $0^0 = 1$)

$$\psi_X : [0, 1] \rightarrow [0, 1], \quad z \mapsto \sum_{n=0}^{\infty} \mathbf{P}[X = n] z^n \quad (3.1)$$

Theorem 3.2

(i) ψ_X is continuous on $[0, 1]$ and infinitely often continuously differentiable on $(0, 1)$. For $n \in \mathbb{N}$, the n th derivative $\psi_X^{(n)}$ fulfills

$$\lim_{z \uparrow 1} \psi_X^{(n)}(z) = \sum_{k=n}^{\infty} \mathbf{P}[X = k] \cdot k(k-1)\dots(k-n+1), \quad (3.2)$$

where both sides can equal ∞ .

(ii) The distribution $\mathbf{P}_X$ of X is uniquely determined by ψ_X .

(iii) For any $r \in (0, 1)$, ψ_X is uniquely determined by countably many values $\psi_X(x_i)$, $x_i \in [0, r]$, $i \in \mathbb{N}$. If the series in (3.1) converges for some $z > 1$, then the statement is also true for any $r \in (0, z)$ and we have

$$\lim_{x \uparrow 1} \psi_X^{(n)}(x) = \psi_X^{(n)}(1) < \infty \quad \text{for } n \in \mathbb{N}.$$

In this case, ψ_X is uniquely determined by the derivatives $\psi_X^{(n)}(1)$, $n \in \mathbb{N}$.

Proof The statements follow from the elementary theory of power series. For the first part of (iii), see, e.g. [148, Theorem 8.5].

Theorem 3.3 (Multiplicativity of generating functions) If $X_1, \dots, X_n$ are independent and $\mathbb{N}_0$ -valued random variables, then

$$\psi_{X_1+\dots+X_n} = \prod_{i=1}^n \psi_{X_i}.$$

Proof Let $z \in [0, 1)$ and write $\psi_{X_1}(z)\psi_{X_2}(z)$ as a Cauchy product

$$\begin{aligned} \psi_{X_1}(z)\psi_{X_2}(z) &= \left(\sum_{n=0}^{\infty} \mathbf{P}[X_1 = n] z^n \right) \left(\sum_{n=0}^{\infty} \mathbf{P}[X_2 = n] z^n \right) \\ &= \sum_{n=0}^{\infty} z^n \left(\sum_{m=0}^n \mathbf{P}[X_1 = m] \mathbf{P}[X_2 = n - m] \right) \\ &= \sum_{n=0}^{\infty} z^n \sum_{m=0}^n \mathbf{P}[X_1 = m, X_2 = n - m] \\ &= \sum_{n=0}^{\infty} \mathbf{P}[X_1 + X_2 = n] z^n = \psi_{X_1+X_2}(z) \end{aligned}$$

Inductively, the claim follows for all $n \geq 2$.

Example 3.4

(i) Let X be $b_{n,p}$ -distributed for some $n \in \mathbb{N}$ and let $p \in [0, 1]$. Then

$$\psi_X(z) = \sum_{m=0}^n \binom{n}{m} p^m (1-p)^{n-m} z^m = (pz + (1-p))^n. \quad (3.3)$$

(ii) If X, Y are independent, $X \sim b_{m,p}$ and $Y \sim b_{n,p}$, then, by Theorem 3.3,

$$\psi_{X+Y}(z) = (pz + (1-p))^m (pz + (1-p))^n = (pz + (1-p))^{m+n}$$

Hence, by Theorem 3.2(ii), $X + Y$ is $b_{m+n,p}$ -distributed and thus (by Theorem 2.31)

$$b_{m,p} * b_{n,p} = b_{m+n,p}$$

(iii) Let X and Y be independent Poisson random variables with parameters $\lambda \geq 0$ and $\mu \geq 0$, respectively. That is, $\mathbf{P}[X = n] = e^{-\lambda} \lambda^n / n!$ for $n \in \mathbb{N}_0$. Then

$$\psi_{\text{Poi}_\lambda}(z) = \sum_{n=0}^{\infty} e^{-\lambda} \frac{(\lambda z)^n}{n!} = e^{\lambda(z-1)} \quad (3.4)$$

Hence $X + Y$ has probability generating function

$$\psi_{\text{Poi}_\lambda}(z) \cdot \psi_{\text{Poi}_\mu}(z) = e^{\lambda(z-1)} e^{\mu(z-1)} = \psi_{\text{Poi}_{\lambda+\mu}}(z)$$

Thus $X + Y \sim \text{Poi}_{\lambda+\mu}$. We conclude that

$$\text{Poi}_\lambda * \text{Poi}_\mu = \text{Poi}_{\lambda+\mu}. \quad (3.5)$$

(iv) Let $X_1, \dots, X_n \sim \gamma_p$ be independent geometrically distributed random variables with parameter $p \in (0, 1)$. Define $Y = X_1 + \dots + X_n$. Then, for any $z \in [0, 1]$,

$$\psi_{X_1}(z) = \sum_{k=0}^{\infty} p(1-p)^k z^k = \frac{p}{1 - (1-p)z} \quad (3.6)$$

By the generalized binomial theorem (see Lemma 3.5 with $\alpha = -n$), Theorem 3.3 and (3.6), we have

$$\begin{aligned}\psi_Y(z) &= \psi_{X_1}(z)^n = \frac{p^n}{(1 - (1-p)z)^n} \\ &= \sum_{k=0}^{\infty} p^n \binom{-n}{k} (-1)^k (1-p)^k z^k \\ &= \sum_{k=0}^{\infty} b_{n,p}^-(\{k\}) z^k.\end{aligned}$$

Here, for $r \in (0, \infty)$ and $p \in (0, 1]$,

$$b_{r,p}^- = \sum_{k=0}^{\infty} \binom{-r}{k} (-1)^k p^r (1-p)^k \delta_k \quad (3.7)$$

is the negative binomial distribution with parameters r and p . By the uniqueness theorem for probability generating functions, we get $Y \sim b_{n,p}^-$; hence (see Definition 2.29 for the n th convolution power) $b_{n,p}^- = \gamma_p^{*n}$.

Lemma 3.5 (Generalized binomial theorem) For $\alpha \in \mathbb{R}$ and $k \in \mathbb{N}_0$, we define the binomial coefficient

$$\binom{\alpha}{k} := \frac{\alpha \cdot (\alpha - 1) \cdots (\alpha - k + 1)}{k!}. \quad (3.8)$$

Then the generalized binomial theorem holds:

$$(1+x)^\alpha = \sum_{k=0}^{\infty} \binom{\alpha}{k} x^k \quad \text{for all } x \in \mathbb{C} \text{ with } |x| < 1 \quad (3.9)$$

In particular, we have

$$\frac{1}{\sqrt{1-x}} = \sum_{n=0}^{\infty} \binom{2n}{n} 4^{-n} x^n \quad \text{for all } x \in \mathbb{C} \text{ with } |x| < 1 \quad (3.10)$$

Proof The map $f : x \mapsto (1+x)^\alpha$ is holomorphic up to possibly a singularity at $x = -1$. Hence it can be developed in a power series about 0 with radius of convergence at least 1:

$$f(x) = \sum_{k=0}^{\infty} \frac{f^{(k)}(0)}{k!} x^k \quad \text{for } |x| < 1$$

For $k \in \mathbb{N}_0$, the k th derivative is $f^{(k)}(0) = \alpha(\alpha-1)\cdots(\alpha-k+1)$. Hence (3.9) holds.

The additional claim follows by the observation that (for $\alpha = -1/2$) we have $\binom{-1/2}{n} = \binom{2n}{n}(-4)^{-n}$.

Exercise 3.1.1 Show that $b_{r,p}^- * b_{s,p}^- = b_{r+s,p}^-$ for $r, s \in (0, \infty)$ and $p \in (0, 1]$.

Exercise 3.1.2 Give an example for two different probability generating functions that coincide at countably many points $x_i \in (0, 1)$, $i \in \mathbb{N}$. (That is, in Theorem 3.2(iii), the assumption $\psi(z) < \infty$ for some $z > 1$ cannot be dropped.)

5.3.2 3.2 Poisson Approximation

Lemma 3.6 Let μ and $(\mu_n)_{n \in \mathbb{N}}$ be probability measures on $(\mathbb{N}_0, 2^{\mathbb{N}_0})$ with generating functions ψ and ψ_n , $n \in \mathbb{N}$. Then the following statements are equivalent.

- (i) $\mu_n(\{k\}) \xrightarrow{n \rightarrow \infty} \mu(\{k\})$ for all $k \in \mathbb{N}_0$.
- (ii) $\mu_n(A) \xrightarrow{n \rightarrow \infty} \mu(A)$ for all $A \subset \mathbb{N}_0$.
- (iii) $\psi_n(z) \xrightarrow{n \rightarrow \infty} \psi(z)$ for all $z \in [0, 1]$.
- (iv) $\psi_n(z) \xrightarrow{n \rightarrow \infty} \psi(z)$ for all $z \in [0, \eta]$ for some $\eta \in (0, 1)$.

We write $\mu_n \xrightarrow{n \rightarrow \infty} \mu$ if any of the four conditions holds and say that $(\mu_n)_{n \in \mathbb{N}}$ converges weakly to μ .

Proof (i) $\implies$ (ii) Fix $\varepsilon > 0$ and choose $N \in \mathbb{N}$ such that

$$\mu(\{N+1, N+2, \dots\}) < \frac{\varepsilon}{4}.$$

For sufficiently large $n_0 \in \mathbb{N}$, we have

$$\sum_{k=0}^N |\mu_n(\{k\}) - \mu(\{k\})| < \frac{\varepsilon}{4} \quad \text{for all } n \geq n_0.$$

In particular, for any $n \geq n_0$, we have $\mu_n(\{N+1, N+2, \dots\}) < \frac{\varepsilon}{2}$. Hence, for $n \geq n_0$,

$$\begin{aligned} |\mu_n(A) - \mu(A)| &\leq \mu_n(\{N+1, N+2, \dots\}) + \mu(\{N+1, N+2, \dots\}) \\ &\quad + \sum_{k \in A \cap \{0, \dots, N\}} |\mu_n(\{k\}) - \mu(\{k\})| \\ &< \varepsilon. \end{aligned}$$

(ii) $\implies$ (i) This is trivial.

(i) $\iff$ (iii) $\iff$ (iv) This follows from the elementary theory of power series.

Let $(p_{n,k})_{n,k \in \mathbb{N}}$ be numbers with $p_{n,k} \in [0, 1]$ such that the limit

$$\lambda := \lim_{n \rightarrow \infty} \sum_{k=1}^{\infty} p_{n,k} \in (0, \infty) \quad (3.11)$$

exists and such that $\lim_{n \rightarrow \infty} \sum_{k=1}^{\infty} p_{n,k}^2 = 0$ (e.g., $p_{n,k} = \lambda/n$ for $k \leq n$ and $p_{n,k} = 0$ for $k > n$). For each $n \in \mathbb{N}$, let $(X_{n,k})_{k \in \mathbb{N}}$ be an independent family of random variables with $X_{n,k} \sim \text{Ber}_{p_{n,k}}$.

Define

$$S^n := \sum_{l=1}^{\infty} X_{n,l} \quad \text{and} \quad S_k^n := \sum_{l=1}^k X_{n,l} \quad \text{for } k \in \mathbb{N}.$$

Theorem 3.7 (Poisson approximation) Under the above assumptions, the distributions $(\mathbf{P}_{S^n})_{n \in \mathbb{N}}$ converge weakly to the Poisson distribution Poi_λ .

Proof The p.g.f. of the Poisson distribution is $\psi(z) = e^{\lambda(z-1)}$ (see (3.4)). On the other hand, $S^n - S_k^n$ and S_k^n are independent for any $k \in \mathbb{N}$; hence $\psi_{S^n} = \psi_{S_k^n} \cdot \psi_{S^n - S_k^n}$. Now, for any $z \in [0, 1]$,

$$1 \geq \frac{\psi_{S^n}(z)}{\psi_{S_k^n}(z)} = \psi_{S^n - S_k^n}(z) \geq 1 - \mathbf{P}[S^n - S_k^n \geq 1] \geq 1 - \sum_{l=k+1}^{\infty} p_{n,l} \xrightarrow{k \rightarrow \infty} 1,$$

hence

$$\begin{aligned} \psi_{S^n}(z) &= \lim_{k \rightarrow \infty} \psi_{S_k^n}(z) = \prod_{l=1}^{\infty} (p_{n,l}z + (1 - p_{n,l})) \\ &= \exp \left(\sum_{l=1}^{\infty} \log(1 + p_{n,l}(z-1)) \right) \end{aligned}$$

Note that $|\log(1+x) - x| \leq x^2$ for $|x| < \frac{1}{2}$. By assumption, $\max_{l \in \mathbb{N}} p_{n,l} \rightarrow 0$ for $n \rightarrow \infty$; hence, for sufficiently large n ,

$$\begin{aligned} & \left| \left(\sum_{l=1}^{\infty} \log(1 + p_{n,l}(z-1)) \right) - \left((z-1) \sum_{l=1}^{\infty} p_{n,l} \right) \right| \\ & \leq \sum_{l=1}^{\infty} p_{n,l}^2 \leq \left(\sum_{l=1}^{\infty} p_{n,l} \right) \max_{l \in \mathbb{N}} p_{n,l} \xrightarrow{n \rightarrow \infty} 0 \end{aligned}$$

Together with (3.11), we infer

$$\lim_{n \rightarrow \infty} \psi_{S^n}(z) = \lim_{n \rightarrow \infty} \exp \left((z-1) \sum_{l=1}^{\infty} p_{n,l} \right) = e^{\lambda(z-1)}$$

5.3.3 3.3 Branching Processes

Let $T, X_1, X_2, \dots$ be independent $\mathbb{N}_0$ -valued random variables. What is the distribution of $S := \sum_{n=1}^T X_n$? First of all, note that S is measurable since

$$\{S = k\} = \bigcup_{n=0}^{\infty} \{T = n\} \cap \{X_1 + \dots + X_n = k\}$$

Theorem 3.8 If the random variables $X_1, X_2, \dots$ are also identically distributed, then the probability generating function of S is given by $\psi_S(z) = \psi_T(\psi_{X_1}(z))$.

Proof We compute

$$\begin{aligned} \psi_S(z) &= \sum_{k=0}^{\infty} \mathbf{P}[S = k] z^k \\ &= \sum_{k=0}^{\infty} \sum_{n=0}^{\infty} \mathbf{P}[T = n] \mathbf{P}[X_1 + \dots + X_n = k] z^k \\ &= \sum_{n=0}^{\infty} \mathbf{P}[T = n] \psi_{X_1}(z)^n = \psi_T(\psi_{X_1}(z)) \end{aligned}$$

Now assume that $p_0, p_1, p_2, \dots \in [0, 1]$ are such that $\sum_{k=0}^{\infty} p_k = 1$. Let $(X_{n,i})_{n,i \in \mathbb{N}_0}$ be an independent family of random variables with $\mathbf{P}[X_{n,i} = k] = p_k$ for all $i, k, n \in \mathbb{N}_0$.

Let $Z_0 = 1$ and

$$Z_n = \sum_{i=1}^{Z_{n-1}} X_{n-1,i} \quad \text{for } n \in \mathbb{N}$$

Z_n can be interpreted as the number of individuals in the n th generation of a randomly developing population. The i th individual in the n th generation has $X_{n,i}$ offspring (in the $(n+1)$ th generation).

Definition 3.9 $(Z_n)_{n \in \mathbb{N}_0}$ is called a Galton-Watson process or branching process with offspring distribution $(p_k)_{k \in \mathbb{N}_0}$.

Probability generating functions are an important tool for the investigation of branching processes. Hence, let

$$\psi(z) = \sum_{k=0}^{\infty} p_k z^k$$

be the p.g.f. of the offspring distribution and let ψ' be its derivative. Recursively, define the n th iterate of ψ by

$$\psi_1 := \psi \quad \text{and} \quad \psi_n := \psi \circ \psi_{n-1} \quad \text{for } n = 2, 3, \dots$$

Finally, let ψ_{Z_n} be the p.g.f. of Z_n .

Lemma 3.10 $\psi_n = \psi_{Z_n}$ for all $n \in \mathbb{N}$.

Proof For $n = 1$, the statement is true by definition. For $n \in \mathbb{N}$, we conclude inductively by Theorem 3.8 that $\psi_{Z_{n+1}} = \psi \circ \psi_{Z_n} = \psi \circ \psi_n = \psi_{n+1}$.

Clearly, the probability $q_n := \mathbf{P}[Z_n = 0]$ that Z is extinct by time n is monotone increasing in n . We denote by

$$q := \lim_{n \rightarrow \infty} \mathbf{P}[Z_n = 0]$$

the extinction probability; that is, the probability that the population will eventually die out.

Under what conditions do we have $q = 0$, $q = 1$, or $q \in (0, 1)$? Clearly, $q \geq p_0$. On the other hand, if $p_0 = 0$, then Z_n is monotone in n ; hence $q = 0$.

Theorem 3.11 (Extinction probability of the Galton-Watson process) Assume $p_1 \neq 1$. Then:

- (i) $F := \{r \in [0, 1] : \psi(r) = r\} = \{q, 1\}$.
- (ii) The following equivalences hold:

$$q < 1 \iff \lim_{z \uparrow 1} \psi'(z) > 1 \iff \sum_{k=1}^{\infty} k p_k > 1.$$

Proof (i) We have $\psi(1) = 1$; hence $1 \in F$. Note that

$$q_n = \psi_n(0) = \psi(q_{n-1}) \quad \text{for all } n \in \mathbb{N}$$

and $q_n \uparrow q$. Since ψ is continuous, we infer

$$\psi(q) = \psi\left(\lim_{n \rightarrow \infty} q_n\right) = \lim_{n \rightarrow \infty} \psi(q_n) = \lim_{n \rightarrow \infty} q_{n+1} = q.$$

Thus $q \in F$. If $r \in F$ is an arbitrary fixed point of ψ , then $r \geq 0 = q_0$. Since ψ is monotone increasing, it follows that $r = \psi(r) \geq \psi(q_0) = q_1$. Inductively, we get $r \geq q_n$ for all $n \in \mathbb{N}_0$; that is, $r \geq q$. We conclude $q = \min F$.

(ii) For the first equivalence, we distinguish two cases.

Case 1: $\lim_{z \uparrow 1} \psi'(z) \leq 1$. Since ψ is strictly convex, in this case, we have $\psi(z) > z$ for all $z \in [0, 1)$; hence $F = \{1\}$. We conclude $q = 1$.

Case 2: $\lim_{z \uparrow 1} \psi'(z) > 1$. As ψ is strictly convex and since $\psi(0) \geq 0$, there is a unique $r \in [0, 1)$ such that $\psi(r) = r$. Hence $F = \{r, 1\}$ and $q = \min F = r$.

The second equivalence in (ii) follows by (3.2).

For further reading, we refer to [5].

5.4 4 The Integral

Based on the notions of measure spaces and measurable maps, we introduce the integral of a measurable map with respect to a general measure. This generalizes the Lebesgue integral that can be found in textbooks on calculus. Furthermore, the integral is a cornerstone in a systematic theory of probability that allows for the definition and investigation of expected values and higher moments of random variables.

In this chapter, we define the integral by an approximation scheme with simple functions. Then we deduce basic statements such as Fatou's lemma. Other important convergence theorems for integrals follow in Chapters 6 and 7.

5.4.1 4.1 Construction and Simple Properties

In the following, $(\Omega, \mathcal{A}, \mu)$ will always be a measure space. We denote by $\mathbb{E}$ the vector space of simple functions (see Definition 1.93) on $(\Omega, \mathcal{A})$ and by

$$\mathbb{E}^+ := \{f \in \mathbb{E} : f \geq 0\}$$

the cone (why this name?) of nonnegative simple functions. If

$$f = \sum_{i=1}^m \alpha_i \mathbb{1}_{A_i} \quad (4.1)$$

for some $m \in \mathbb{N}$ and for $\alpha_1, \dots, \alpha_m \in (0, \infty)$, and for mutually disjoint sets $A_1, \dots, A_m \in \mathcal{A}$, then (4.1) is said to be a normal representation of f .

Lemma 4.1 If $f = \sum_{i=1}^m \alpha_i \mathbb{1}_{A_i}$ and $f = \sum_{j=1}^n \beta_j \mathbb{1}_{B_j}$ are two normal representations of $f \in \mathbb{E}^+$, then

$$\sum_{i=1}^m \alpha_i \mu(A_i) = \sum_{j=1}^n \beta_j \mu(B_j).$$

Proof If $\mu(A_i \cap B_j) > 0$ for some i and j , then $A_i \cap B_j \neq \emptyset$, and $f(\omega) = \alpha_i = \beta_j$ for any $\omega \in A_i \cap B_j$. Furthermore, clearly $A_i \subset \bigcup_{j=1}^n B_j$ if $\alpha_i \neq 0$, and $B_j \subset \bigcup_{i=1}^m A_i$ if $\beta_j \neq 0$. We conclude that

$$\begin{aligned} \sum_{i=1}^m \alpha_i \mu(A_i) &= \sum_{i=1}^m \sum_{j=1}^n \alpha_i \mu(A_i \cap B_j) \\ &= \sum_{i=1}^m \sum_{j=1}^n \beta_j \mu(A_i \cap B_j) = \sum_{j=1}^n \beta_j \mu(B_j). \end{aligned}$$

This lemma allows us to make the following definition (since the value of $I(f)$ does not depend on the choice of the normal representation).

Definition 4.2 Define the map $I : \mathbb{E}^+ \rightarrow [0, \infty]$ by

$$I(f) = \sum_{i=1}^m \alpha_i \mu(A_i)$$

if f has the normal representation $f = \sum_{i=1}^m \alpha_i \mathbb{1}_{A_i}$.

Lemma 4.3 The map I is positive linear and monotone increasing: Let $f, g \in \mathbb{E}^+$ and $\alpha \geq 0$. Then the following statements hold.

- (i) $I(\alpha f) = \alpha I(f)$.
- (ii) $I(f + g) = I(f) + I(g)$.
- (iii) If $f \leq g$, then $I(f) \leq I(g)$.

Proof This is left as an exercise.

Definition 4.4 (Integral) If $f : \Omega \rightarrow [0, \infty]$ is measurable, then we define the integral of f with respect to μ by

$$\int f d\mu := \sup \{I(g) : g \in \mathbb{E}^+, g \leq f\}$$

Remark 4.5 By Lemma 4.3(iii), we have $I(f) = \int f d\mu$ for any $f \in \mathbb{E}^+$. Hence the integral is an extension of the map I from $\mathbb{E}^+$ to the set of nonnegative measurable functions.

If $f, g : \Omega \rightarrow \overline{\mathbb{R}}$ with $f(\omega) \leq g(\omega)$ for any $\omega \in \Omega$, then we write $f \leq g$. Analogously, we write $f \geq 0$ and so on. On the other hand, we write " $f \leq g$ almost everywhere" if the weaker condition holds that there exists a μ -null set N such that $f(\omega) \leq g(\omega)$ for any $\omega \in N^c$.

Lemma 4.6 Let $f, g, f_1, f_2, \dots$ be measurable maps $\Omega \rightarrow [0, \infty]$. Then:

- (i) (Monotonicity) If $f \leq g$, then $\int f d\mu \leq \int g d\mu$.
- (ii) (Monotone convergence) If $f_n \uparrow f$, then the integrals also converge: $\int f_n d\mu \uparrow \int f d\mu$.
- (iii) (Linearity) If $\alpha, \beta \in [0, \infty]$, then

$$\int (\alpha f + \beta g) d\mu = \alpha \int f d\mu + \beta \int g d\mu$$

where we use the convention $\infty \cdot 0 := 0$.

Proof (i) This is immediate from the definition of the integral.

(ii) By (i), we have

$$\lim_{n \rightarrow \infty} \int f_n d\mu = \sup_{n \in \mathbb{N}} \int f_n d\mu \leq \int f d\mu$$

Hence we only have to show $\int f d\mu \leq \sup_{n \in \mathbb{N}} \int f_n d\mu$.

Let $g \in \mathbb{E}^+$ with $g \leq f$. It is enough to show that

$$\sup_{n \in \mathbb{N}} \int f_n d\mu \geq \int g d\mu \quad (4.2)$$

Assume that the simple function g has the normal representation $g = \sum_{i=1}^N \alpha_i \mathbb{1}_{A_i}$ for some $\alpha_1, \dots, \alpha_N \in (0, \infty)$ and mutually disjoint sets $A_1, \dots, A_N \in \mathcal{A}$. For any $\varepsilon > 0$ and $n \in \mathbb{N}$, define the set

$$B_n^\varepsilon = \{f_n \geq (1 - \varepsilon)g\}$$

Since $f_n \uparrow f \geq g$, we have $B_n^\varepsilon \uparrow \Omega$ for any $\varepsilon > 0$. Hence, by (i), for any $\varepsilon > 0$,

$$\begin{aligned} \int f_n d\mu &\geq \int ((1 - \varepsilon)g \mathbb{1}_{B_n^\varepsilon}) d\mu \\ &= \sum_{i=1}^N (1 - \varepsilon) \alpha_i \mu(A_i \cap B_n^\varepsilon) \xrightarrow{n \rightarrow \infty} \sum_{i=1}^N (1 - \varepsilon) \alpha_i \mu(A_i) \\ &= (1 - \varepsilon) \int g d\mu \end{aligned}$$

Letting $\varepsilon \downarrow 0$ implies (4.2) and hence the claim (ii).

(iii) By Theorem 1.96, any nonnegative measurable map is a monotone limit of simple functions. Hence there are sequences $(f_n)_{n \in \mathbb{N}}$ and $(g_n)_{n \in \mathbb{N}}$ in $\mathbb{E}^+$ such that $f_n \uparrow f$ and $g_n \uparrow g$. Thus also $(\alpha f_n + \beta g_n) \uparrow \alpha f + \beta g$. By (ii) and Lemma 4.3, this implies

$$\begin{aligned} \int (\alpha f + \beta g) d\mu &= \lim_{n \rightarrow \infty} \int (\alpha f_n + \beta g_n) d\mu \\ &= \alpha \lim_{n \rightarrow \infty} \int f_n d\mu + \beta \lim_{n \rightarrow \infty} \int g_n d\mu = \alpha \int f d\mu + \beta \int g d\mu \end{aligned}$$

For any measurable map $f : \Omega \rightarrow \overline{\mathbb{R}}$, we have $f^+ \leq |f|$ and $f^- \leq |f|$, which implies $\int f^\pm d\mu \leq \int |f| d\mu$. In particular, if $\int |f| d\mu < \infty$, then also $\int f^- d\mu < \infty$ and $\int f^+ d\mu < \infty$.

∞ . Thus we can make the following definition that is the final definition for the integral of measurable functions.

Definition 4.7 (Integral of measurable functions) A measurable function $f : \Omega \rightarrow \overline{\mathbb{R}}$ is called μ -integrable if $\int |f| d\mu < \infty$. We write

$$\mathcal{L}^1(\mu) := \mathcal{L}^1(\Omega, \mathcal{A}, \mu) := \left\{ f : \Omega \rightarrow \overline{\mathbb{R}} : f \text{ is measurable and } \int |f| d\mu < \infty \right\}.$$

For $f \in \mathcal{L}^1(\mu)$, we define the integral of f with respect to μ by

$$\int f(\omega) \mu(d\omega) := \int f d\mu := \int f^+ d\mu - \int f^- d\mu \quad (4.3)$$

If we only have $\int f^- d\mu < \infty$ or $\int f^+ d\mu < \infty$, then we also define $\int f d\mu$ by (4.3). Here the values $+\infty$ and $-\infty$, respectively, are possible.

For $A \in \mathcal{A}$, we define $\int_A f d\mu := \int (f \mathbb{1}_A) d\mu$.

Theorem 4.8 Let $f : \Omega \rightarrow [0, \infty]$ be a measurable map.

(i) We have $f = 0$ almost everywhere if and only if $\int f d\mu = 0$.

(ii) If $\int f d\mu < \infty$, then $f < \infty$ almost everywhere.

Proof (i) " $\implies$ " Assume $f = 0$ almost everywhere. Let $N = \{\omega : f(\omega) > 0\}$. Then $f \leq \infty \cdot \mathbb{1}_N$ and $n\mathbb{1}_N \uparrow \infty \cdot \mathbb{1}_N$. From Lemma 4.6(i) and (ii), we infer

$$0 \leq \int f d\mu \leq \int (\infty \cdot \mathbb{1}_N) d\mu = \lim_{n \rightarrow \infty} \int n\mathbb{1}_N d\mu = 0$$

" $\Leftarrow$ " Let $N_n = \{f \geq \frac{1}{n}\}$, $n \in \mathbb{N}$. Then $N_n \uparrow N$ and

$$0 = \int f d\mu \geq \int \frac{1}{n} \mathbb{1}_{N_n} d\mu = \frac{\mu(N_n)}{n}$$

Hence $\mu(N_n) = 0$ for any $n \in \mathbb{N}$ and thus $\mu(N) = 0$.

(ii) Let $A = \{\omega : f(\omega) = \infty\}$. For $n \in \mathbb{N}$, we have $\frac{1}{n} f \mathbb{1}_{\{f \geq n\}} \geq \mathbb{1}_{\{f \geq n\}}$. Hence Lemma 4.6(i) implies

$$\mu(A) = \int \mathbb{1}_A d\mu \leq \int \mathbb{1}_{\{f \geq n\}} d\mu \leq \frac{1}{n} \int f \mathbb{1}_{\{f \geq n\}} d\mu \leq \frac{1}{n} \int f d\mu \xrightarrow{n \rightarrow \infty} 0$$

Theorem 4.9 (Properties of the integral) Let $f, g \in \mathcal{L}^1(\mu)$.

(i) (Monotonicity) If $f \leq g$ almost everywhere, then $\int f d\mu \leq \int g d\mu$. In particular, if $f = g$ almost everywhere, then $\int f d\mu = \int g d\mu$.

(ii) (Triangle inequality) $|\int f d\mu| \leq \int |f| d\mu$.

(iii) (Linearity) If $\alpha, \beta \in \mathbb{R}$, then $\alpha f + \beta g \in \mathcal{L}^1(\mu)$ and

$$\int (\alpha f + \beta g) d\mu = \alpha \int f d\mu + \beta \int g d\mu$$

This equation also holds if at most one of the integrals $\int f d\mu$ and $\int g d\mu$ is infinite.

Proof (i) Clearly, $f^+ \leq g^+$ and $f^- \geq g^-$ a.e. Hence, by Lemma 4.6(i),

$$\int f^+ d\mu \leq \int g^+ d\mu \quad \text{and} \quad \int f^- d\mu \geq \int g^- d\mu$$

This implies

$$\int f d\mu = \int f^+ d\mu - \int f^- d\mu \leq \int g^+ d\mu - \int g^- d\mu = \int g d\mu.$$

(ii) Since $f^+ + f^- = |f|$, Lemma 4.6(iii) yields

$$\begin{aligned} \left| \int f d\mu \right| &= \left| \int f^+ d\mu - \int f^- d\mu \right| \leq \int f^+ d\mu + \int f^- d\mu \\ &= \int (f^+ + f^-) d\mu = \int |f| d\mu \end{aligned}$$

(iii) Since $|\alpha f + \beta g| \leq |\alpha| \cdot |f| + |\beta| \cdot |g|$, Lemma 4.6(i) and (iii) yield that $\alpha f + \beta g \in \mathcal{L}^1(\mu)$. In order to show linearity, it is enough to check the following three properties.

(a) $\int (f + g) d\mu = \int f d\mu + \int g d\mu$.

(b) $\int \alpha f d\mu = \alpha \int f d\mu$ for $\alpha \geq 0$.

(c) $\int (-f) d\mu = - \int f d\mu$.

(a) We have $(f + g)^+ - (f + g)^- = f + g = f^+ - f^- + g^+ - g^-$; hence $(f + g)^+ + f^- + g^- = (f + g)^- + f^+ + g^+$. By Lemma 4.6(iii), we infer

$$\int (f + g)^+ d\mu + \int f^- d\mu + \int g^- d\mu = \int (f + g)^- d\mu + \int f^+ d\mu + \int g^+ d\mu.$$

Hence

$$\begin{aligned} \int (f + g) d\mu &= \int (f + g)^+ d\mu - \int (f + g)^- d\mu \\ &= \int f^+ d\mu - \int f^- d\mu + \int g^+ d\mu - \int g^- d\mu \\ &= \int f d\mu + \int g d\mu \end{aligned}$$

(b) For $\alpha \geq 0$, we have

$$\int \alpha f d\mu = \int \alpha f^+ d\mu - \int \alpha f^- d\mu = \alpha \int f^+ d\mu - \alpha \int f^- d\mu = \alpha \int f d\mu.$$

(c) We have

$$\begin{aligned} \int (-f) d\mu &= \int (-f)^+ d\mu - \int (-f)^- d\mu \\ &= \int f^- d\mu - \int f^+ d\mu = - \int f d\mu \end{aligned}$$

The supplementary statement is simple and is left as an exercise.

Theorem 4.10 (Image measure) Let $(\Omega, \mathcal{A})$ and $(\Omega', \mathcal{A}')$ be measurable spaces, let μ be a measure on $(\Omega, \mathcal{A})$ and let $X : \Omega \rightarrow \Omega'$ be measurable. Let $\mu' = \mu \circ X^{-1}$ be the image measure of μ under the map X . Assume that $f : \Omega' \rightarrow \overline{\mathbb{R}}$ is μ' -integrable. Then $f \circ X \in \mathcal{L}^1(\mu)$ and

$$\int (f \circ X) d\mu = \int f d(\mu \circ X^{-1})$$

In particular, if X is a random variable on $(\Omega, \mathcal{A}, \mathbf{P})$, then

$$\int f(x) \mathbf{P}[X \in dx] := \int f(x) \mathbf{P}_X[dx] = \int f d\mathbf{P}_X = \int f(X(\omega)) \mathbf{P}[d\omega]$$

Proof This is left as an exercise.

Example 4.11 (Discrete measure space) Let $(\Omega, \mathcal{A})$ be a discrete measurable space and let $\mu = \sum_{\omega \in \Omega} \alpha_\omega \delta_\omega$ for certain numbers $\alpha_\omega \geq 0, \omega \in \Omega$. A map $f : \Omega \rightarrow \mathbb{R}$ is integrable if and only if $\sum_{\omega \in \Omega} |f(\omega)| \alpha_\omega < \infty$. In this case,

$$\int f d\mu = \sum_{\omega \in \Omega} f(\omega) \alpha_\omega.$$

Definition 4.12 (Lebesgue integral) Let λ be the Lebesgue measure on $\mathbb{R}^n$ and let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be measurable with respect to $\mathcal{B}^*(\mathbb{R}^n) - \mathcal{B}(\mathbb{R})$ (here $\mathcal{B}^*(\mathbb{R}^n)$ is the Lebesgue σ -algebra; see Example 1.71) and λ -integrable. Then we call

$$\int f d\lambda$$

the Lebesgue integral of f . If $A \in \mathcal{B}(\mathbb{R}^n)$ and $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is measurable (or $f : A \rightarrow \mathbb{R}$ is $\mathcal{B}^*(\mathbb{R}^n)|_A - \mathcal{B}(\mathbb{R})$ -measurable and hence f_1 is $\mathcal{B}^*(\mathbb{R}^n) - \mathcal{B}(\mathbb{R})$ measurable), then we write

$$\int_A f d\lambda := \int f \mathbb{1}_A d\lambda$$

Definition 4.13 Let μ be a measure on $(\Omega, \mathcal{A})$ and let $f : \Omega \rightarrow [0, \infty)$ be a measurable map. Define the measure ν by

$$\nu(A) := \int (\mathbb{1}_A f) d\mu \quad \text{for } A \in \mathcal{A}$$

We say that $f\mu := \nu$ has density f with respect to μ .

Remark 4.14 We still have to show that ν is a measure. To this end, we check the conditions of Theorem 1.36. Clearly, $\nu(\emptyset) = 0$. Finite additivity follows from additivity of the integral (Lemma 4.6(iii)). Lower semicontinuity follows from the monotone convergence theorem (Theorem 4.20).

Theorem 4.15 We have $g \in \mathcal{L}^1(f\mu)$ if and only if $(gf) \in \mathcal{L}^1(\mu)$. In this case,

$$\int g d(f\mu) = \int (gf) d\mu$$

Proof First note that the statement holds for indicator functions. Then, with the usual arguments, extend it step by step first to simple functions, then to nonnegative measurable functions and finally to signed measurable functions.

Definition 4.16 For measurable $f : \Omega \rightarrow \overline{\mathbb{R}}$, define

$$\|f\|_p := \left(\int |f|^p d\mu \right)^{1/p}, \quad \text{if } p \in [1, \infty)$$

and

$$\|f\|_\infty := \inf \{K \geq 0 : \mu(\{|f| > K\}) = 0\}.$$

Further, for any $p \in [1, \infty]$, define the vector space

$$\mathcal{L}^p(\mu) := \{f : \Omega \rightarrow \overline{\mathbb{R}} \text{ is measurable and } \|f\|_p < \infty\}$$

Theorem 4.17 The map $\|\cdot\|_1$ is a seminorm on $\mathcal{L}^1(\mu)$; that is, for all $f, g \in \mathcal{L}^1(\mu)$ and $\alpha \in \mathbb{R}$,

$$\begin{aligned} \|\alpha f\|_1 &= |\alpha| \cdot \|f\|_1, \\ \|f + g\|_1 &\leq \|f\|_1 + \|g\|_1, \\ \|f\|_1 &\geq 0 \quad \text{for all } f \quad \text{and} \quad \|f\|_1 = 0 \quad \text{if } f = 0 \text{ a.e.} \end{aligned} \tag{4.4}$$

Proof The first and the third statements follow from Theorem 4.9(iii) and Theorem 4.8(i). The second statement follows from Theorem 4.9(i) since $|f + g| \leq |f| + |g|$; hence

$$\|f + g\|_1 = \int |f + g| d\mu \leq \int |f| d\mu + \int |g| d\mu = \|f\|_1 + \|g\|_1$$

Remark 4.18 In fact, $\|\cdot\|_p$ is a seminorm on $\mathcal{L}^p(\mu)$ for all $p \in [1, \infty]$. Linearity and positivity are obvious, and the triangle inequality is a consequence of Minkowski's inequality, which we will show in Theorem 7.17.

Theorem 4.19 Let $\mu(\Omega) < \infty$ and $1 \leq p' \leq p \leq \infty$. Then $\mathcal{L}^p(\mu) \subset \mathcal{L}^{p'}(\mu)$ and the canonical inclusion $i : \mathcal{L}^p(\mu) \hookrightarrow \mathcal{L}^{p'}(\mu)$, $f \mapsto f$ is continuous.

Proof Let $f \in \mathcal{L}^\infty(\mu)$ and $p' \in [1, \infty)$. Then $|f|^{p'} \leq \|f\|_\infty^{p'}$ almost everywhere; hence

$$\int |f|^{p'} d\mu \leq \int \|f\|_\infty^{p'} d\mu = \|f\|_\infty^{p'} \cdot \mu(\Omega) < \infty$$

Thus $\|f - g\|_{p'} \leq \mu(\Omega)^{1/p'} \|f - g\|_\infty$ for $f, g \in \mathcal{L}^\infty(\mu)$ and hence i is continuous. Now let $p, p' \in [1, \infty)$ with $p' < p$ and let $f \in \mathcal{L}^p(\mu)$. Then $|f|^{p'} \leq 1 + |f|^p$; hence

$$\int |f|^{p'} d\mu \leq \mu(\Omega) + \int |f|^p d\mu < \infty$$

Finally, let $f, g \in \mathcal{L}^p(\mu)$. For any $c > 0$, we have

$$|f - g|^{p'} = |f - g|^{p'} \mathbb{1}_{\{|f-g| \leq c\}} + |f - g|^{p'} \mathbb{1}_{\{|f-g| > c\}} \leq c^{p'} + c^{p'-p} |f - g|^p.$$

In particular, letting $c = \|f - g\|_p$ we obtain

$$\|f - g\|_{p'} \leq \left(c^{p'} \mu(\Omega) + c^{p'-p} \|f - g\|_p^p \right)^{1/p'} = (1 + \mu(\Omega))^{1/p'} \|f - g\|_p$$

Hence, also in this case, i is continuous.

Exercise 4.1.1 (Sequence spaces) Now we do not assume $\mu(\Omega) < \infty$. Assume there exists an $a > 0$ such that for any $A \in \mathcal{A}$ either $\mu(A) = 0$ or $\mu(A) \geq a$. Show that the reverse inclusion to Theorem 4.19 holds,

$$\mathcal{L}^{p'}(\mu) \subset \mathcal{L}^p(\mu) \quad \text{if } 1 \leq p' \leq p \leq \infty \quad (4.5)$$

Exercise 4.1.2 Let $1 \leq p' < p \leq \infty$ and let μ be σ -finite but not finite. Show that $\mathcal{L}^p(\mu) \setminus \mathcal{L}^{p'}(\mu) \neq \emptyset$.

5.4.2 4.2 Monotone Convergence and Fatou's Lemma

What are the conditions that allow the interchange of limit and integral? In this section, we derive two simple criteria that prepare us for important applications such as the law of large numbers (Chapter 5). More general criteria will be presented in Chapter 6.

Theorem 4.20 (Monotone convergence, Beppo Levi theorem) Let $f_1, f_2, \dots \in \mathcal{L}^1(\mu)$ and let $f : \Omega \rightarrow \overline{\mathbb{R}}$ be measurable. Assume $f_n \uparrow f$ a.e. for $n \rightarrow \infty$. Then

$$\lim_{n \rightarrow \infty} \int f_n d\mu = \int f d\mu$$

where both sides can equal $+\infty$.

Proof Let $N \subset \Omega$ be a null set such that $f_n(\omega) \uparrow f(\omega)$ for all $\omega \in N^c$. The functions $f'_n := (f_n - f_1) \mathbb{1}_{N^c}$ and $f' := (f - f_1) \mathbb{1}_{N^c}$ are nonnegative and fulfill $f'_n \uparrow f'$. By Lemma 4.6(ii), we have $\int f'_n d\mu \xrightarrow{n \rightarrow \infty} \int f' d\mu$. Since $f_n = f'_n + f_1$ a.e. and $f = f' + f_1$ a.e., Theorem 4.9(iii) implies

$$\int f_n d\mu = \int f_1 d\mu + \int f'_n d\mu \xrightarrow{n \rightarrow \infty} \int f_1 d\mu + \int f' d\mu = \int f d\mu$$

Theorem 4.21 (Fatou's lemma) Let $f \in \mathcal{L}^1(\mu)$ and let $f_1, f_2, \dots$ be measurable with $f_n \geq f$ a.e. for all $n \in \mathbb{N}$. Then

$$\int \left(\liminf_{n \rightarrow \infty} f_n \right) d\mu \leq \liminf_{n \rightarrow \infty} \int f_n d\mu$$

Proof By considering $(f_n - f)_{n \in \mathbb{N}}$, we may assume $f_n \geq 0$ a.e. for all $n \in \mathbb{N}$. Define

$$g_n := \inf_{m \geq n} f_m$$

Then $g_n \uparrow \liminf_{m \rightarrow \infty} f_m$ as $n \rightarrow \infty$, and hence by the monotone convergence theorem (Lemma 4.6(ii)) and by monotonicity, $g_n \leq f_n$ (thus $\int g_n d\mu \leq \int f_n d\mu$),

$$\int \liminf_{n \rightarrow \infty} f_n d\mu = \lim_{n \rightarrow \infty} \int g_n d\mu \leq \liminf_{n \rightarrow \infty} \int f_n d\mu$$

Example 4.22 (Petersburg game) By a concrete example, we show that in Fatou's lemma the assumption of an integrable minorant is essential. Consider a gamble in a casino where in each round the player's bet either gets doubled or lost. For example, roulette is such a game. If the player bets on "red", she gets the stake back doubled if the ball lands in a red pocket. Otherwise the bet is lost (for the player, not for the casino). There are 37 pockets (in European roulettes), 18 of which are red, 18 are black and one is green (the zero). Hence, by symmetry, the chance of winning should be $p = 18/37 < \frac{1}{2}$. Now assume the gamble is played again and again. We can model this on a probability space $(\Omega, \mathcal{A}, \mathbf{P})$ where $\Omega = \{-1, 1\}^{\mathbb{N}}$, $\mathcal{A} = (2^{\{-1, 1\}})^{\otimes \mathbb{N}}$ is the σ -algebra generated by the cylinder sets $[\omega_1, \dots, \omega_n]$ and $\mathbf{P} = ((1-p)\delta_{-1} + p\delta_1)^{\otimes \mathbb{N}}$ is the product measure. Denote by $D_n : \Omega \rightarrow \{-1, 1\}$, $\omega \mapsto \omega_n$ the result of the n th game (for $n \in \mathbb{N}$). If in the i th game the player makes a (random) stake of H_i euros, then the cumulative profit after the n th game is

$$S_n = \sum_{i=1}^n H_i D_i$$

Now assume the gambler adopts the following doubling strategy. In the first round, the stake is $H_1 = 1$. If she wins, then she does not bet any money in the subsequent games; that is, $H_n = 0$ for all $n \geq 2$ if $D_1 = 1$. On the other hand, if she loses, then in the second game she doubles the stake; that is, $H_2 = 2$ if $D_1 = -1$. If she wins the second game, she leaves the casino and otherwise doubles the stake again and so on. Hence we can describe the strategy by the formula

$$H_n = \begin{cases} 0, & \text{if there is an } i \in \{1, \dots, n-1\} \text{ with } D_i = 1 \\ 2^{n-1}, & \text{else.} \end{cases}$$

Note that H_n depends on $D_1, \dots, D_{n-1}$ only. That is, it is measurable with respect to $\sigma(D_1, \dots, D_{n-1})$. Clearly, it is a crucial requirement for any strategy that the decision for the next stake depend only on the information available at that time and not depend on the future results of the gamble.

The probability of no win until the n th game is $(1-p)^n$; hence $\mathbf{P}[S_n = 1 - 2^n] = (1-p)^n$ and $\mathbf{P}[S_n = 1] = 1 - (1-p)^n$. Hence we expect an average gain of

$$\int S_n d\mathbf{P} = (1-p)^n (1 - 2^n) + (1 - (1-p)^n) = 1 - (2(1-p))^n \leq 0$$

since $p \leq \frac{1}{2}$ (in the profitable casinos). We define

$$S = \begin{cases} -\infty, & \text{if } -1 = D_1 = D_2 = \dots \\ 1, & \text{else} \end{cases}$$

Then $S_n \xrightarrow{n \rightarrow \infty} S$ a.s. but $\lim_{n \rightarrow \infty} \int S_n d\mathbf{P} < \int S d\mathbf{P} = 1$ since $S = 1$ a.s. By Fatou's lemma, this is possible only if there is no integrable minorant for the sequence $(S_n)_{n \in \mathbb{N}}$. If we define $\tilde{S} := \inf \{S_n : n \in \mathbb{N}\}$, then indeed

$$\mathbf{P} [\tilde{S} = 1 - 2^{n-1}] = \mathbf{P} [D_1 = \dots = D_{n-1} = -1 \quad \text{and} \quad D_n = 1] = p(1-p)^{n-1}$$

Hence $\int \tilde{S} d\mathbf{P} = \sum_{n=1}^{\infty} (1 - 2^{n-1}) p(1-p)^{n-1} = -\infty$ since $p \leq \frac{1}{2}$.

Exercise 4.2.1 Let $(\Omega, \mathcal{A}, \mu)$ be a measure space and let $f \in \mathcal{L}^1(\mu)$. Show that for any $\varepsilon > 0$, there is an $A \in \mathcal{A}$ with $\mu(A) < \infty$ and $|\int_A f d\mu - \int f d\mu| < \varepsilon$.

Exercise 4.2.2 Let $f_1, f_2, \dots \in \mathcal{L}^1(\mu)$ be nonnegative and such that $\lim_{n \rightarrow \infty} \int f_n d\mu$ exists. Assume there exists a measurable f with $f_n \xrightarrow{n \rightarrow \infty} f$ μ -almost everywhere. Show that $f \in \mathcal{L}^1(\mu)$ and

$$\lim_{n \rightarrow \infty} \int |f_n - f| d\mu = \lim_{n \rightarrow \infty} \int f_n d\mu - \int f d\mu$$

Exercise 4.2.3 Let $f \in \mathcal{L}^1([0, \infty), \lambda)$ be a Lebesgue integrable function on $[0, \infty)$. Show that for λ -almost all $t \in [0, \infty)$ the series $\sum_{n=1}^{\infty} f(nt)$ converges absolutely.

Exercise 4.2.4 Let λ be the Lebesgue measure on $\mathbb{R}$ and let A be a Borel set with $\lambda(A) < \infty$. Show that for any $\varepsilon > 0$, there is a compact set $C \subset A$, a closed set $D \subset \mathbb{R} \setminus A$ and a continuous map $\varphi : \mathbb{R} \rightarrow [0, 1]$ with $\mathbb{1}_C \leq \varphi \leq \mathbb{1}_{\mathbb{R} \setminus D}$ and such that $\|\mathbb{1}_A - \varphi\|_1 < \varepsilon$.

Hint: Use the regularity of Lebesgue measure (Remark 1.67).

Exercise 4.2.5 Let λ be the Lebesgue measure on $\mathbb{R}$, $p \in [1, \infty)$ and let $f \in \mathcal{L}^p(\lambda)$. Show that for any $\varepsilon > 0$, there is a continuous function $h : \mathbb{R} \rightarrow \mathbb{R}$ such that $\|f - h\|_p < \varepsilon$.

Hint: Use Exercise 4.2.4 to show the assertion first for indicator functions, then for simple functions and finally for general $f \in \mathcal{L}^p(\lambda)$.

Exercise 4.2.6 Let λ be the Lebesgue measure on $\mathbb{R}$, $p \in [1, \infty)$ and let $f \in \mathcal{L}^p(\lambda)$. A map $h : \mathbb{R} \rightarrow \mathbb{R}$ is called a step function if there exist $n \in \mathbb{N}$ and numbers $t_0 < t_1 < \dots < t_n$ and $\alpha_1, \dots, \alpha_n$ such that $h = \sum_{k=1}^n \alpha_k \mathbb{1}_{(t_{k-1}, t_k]}$.

Show that for any $\varepsilon > 0$, there exists a step function h such that $\|f - h\|_p < \varepsilon$.

Hint: Use the approximation theorem for measures (Theorem 1.65) with the semiring of left open intervals to show the assertion first for measurable indicator functions. Then use the approximation arguments as in Exercise 4.2.5.

5.4.3 4.3 Lebesgue Integral Versus Riemann Integral

We show that for Riemann integrable functions the Lebesgue integral and the Riemann integral coincide.

Let $I = [a, b] \subset \mathbb{R}$ be an interval and let λ be the Lebesgue measure on I . Further, consider sequences $t = (t^n)_{n \in \mathbb{N}}$ of partitions $t^n = (t_i^n)_{i=0, \dots, n}$ of I (i.e., $a = t_0^n < t_1^n < \dots < t_n^n = b$) that get finer and finer. That is,

$$|t^n| := \max \{t_i^n - t_{i-1}^n : i = 1, \dots, n\} \xrightarrow{n \rightarrow \infty} 0.$$

Assume that for any $n \in \mathbb{N}$, the partition t^{n+1} is a refinement of t^n ; that is, $\{t_0^n, \dots, t_n^n\} \subset \{t_0^{n+1}, \dots, t_{n+1}^{n+1}\}$.

For any function $f : I \rightarrow \mathbb{R}$ and any $n \in \mathbb{N}$, define the n th lower sum and upper sum, respectively, by

$$L_n^t(f) := \sum_{i=1}^n (t_i^n - t_{i-1}^n) \inf f([t_{i-1}^n, t_i^n])$$

$$U_n^t(f) := \sum_{i=1}^n (t_i^n - t_{i-1}^n) \sup f([t_{i-1}^n, t_i^n])$$

A function $f : I \rightarrow \mathbb{R}$ is called Riemann integrable if there exists a t such that the limits of the lower sums and upper sums are finite and coincide. In this case, the value of the limit does not depend on the choice of t , and the Riemann integral of f is defined as (see, e.g., [148])

$$\int_a^b f(x)dx := \lim_{n \rightarrow \infty} L_n^t(f) = \lim_{n \rightarrow \infty} U_n^t(f) \quad (4.6)$$

Theorem 4.23 (Riemann integral and Lebesgue integral) Let $f : I \rightarrow \mathbb{R}$ be Riemann integrable on $I = [a, b]$. Then f is Lebesgue integrable on I with integral

$$\int_I f d\lambda = \int_a^b f(x)dx$$

Proof Choose t such that (4.6) holds. By assumption, there is an $n \in \mathbb{N}$ with $|L_n^t(f)| < \infty$ and $|U_n^t(f)| < \infty$. Hence f is bounded. We can thus replace f by $f + \|f\|_\infty$ and hence assume that $f \geq 0$. Define

$$g_n := f(b)\mathbb{1}_{\{b\}} + \sum_{i=1}^n (\inf f([t_{i-1}^n, t_i^n])) \mathbb{1}_{[t_{i-1}^n, t_i^n]}$$

$$h_n := f(b)\mathbb{1}_{\{b\}} + \sum_{i=1}^n (\sup f([t_{i-1}^n, t_i^n])) \mathbb{1}_{[t_{i-1}^n, t_i^n]}$$

As t^{n+1} is a refinement of t^n , we have $g_n \leq g_{n+1} \leq h_{n+1} \leq h_n$. Hence there exist g and h with $g_n \uparrow g$ and $h_n \downarrow h$. By construction, we have $g \leq h$ and

$$\begin{aligned} \int_I g d\lambda &= \lim_{n \rightarrow \infty} \int_I g_n d\lambda = \lim_{n \rightarrow \infty} L_n^t(f) \\ &= \lim_{n \rightarrow \infty} U_n^t(f) = \lim_{n \rightarrow \infty} \int_I h_n d\lambda = \int_I h d\lambda \end{aligned}$$

Hence $h = g$ λ -a.e. By construction, $g \leq f \leq h$, and as limits of simple functions, g and h are $\mathcal{B}(I) - \mathcal{B}(\mathbb{R})$ -measurable. This implies that, for any $\alpha \in \mathbb{R}$, the set

$$\{f \leq \alpha\} = (\{g \leq \alpha\} \cap \{g = h\}) \uplus (\{f \leq \alpha\} \cap \{g \neq h\})$$

is the union of a $\mathcal{B}(I)$ -set with a subset of a null set and is hence in $\mathcal{B}(I)^*$ (the Lebesgue completion of $\mathcal{B}(I)$). Hence f is $\mathcal{B}(I)^*$ -measurable. By the monotone convergence theorem (Theorem 4.20), we conclude

$$\int_I f d\lambda = \lim_{n \rightarrow \infty} \int_I g_n d\lambda = \int_a^b f(x)dx$$

Example 4.24 Let $f : [0, 1] \rightarrow \mathbb{R}, x \mapsto \mathbb{1}_{\mathbb{Q}}$. Then clearly f is not Riemann integrable since $L_n(f) = 0$ and $U_n(f) = 1$ for all $n \in \mathbb{N}$. On the other hand, f is Lebesgue integrable with integral $\int_{[0,1]} f d\lambda = 0$ because $\mathbb{Q} \cap [0, 1]$ is a null set.

Remark 4.25 An improperly Riemann integrable function f on a one-sided open interval $I = (a, b]$ or $I = [0, \infty)$ is not necessarily Lebesgue integrable. Indeed, the improper integral

$\int_0^\infty f(x)dx := \lim_{n \rightarrow \infty} \int_0^n f(x)dx$ is defined by a limit procedure that respects the geometry of $\mathbb{R}$. The Lebesgue integral does not do that. For example, the function $f : [0, \infty) \rightarrow \mathbb{R}, x \mapsto \frac{1}{1+x} \sin(x)$ is improperly Riemann integrable but is not Lebesgue integrable since $\int_{[0, \infty)} |f| d\lambda = \infty$.

On the one hand, improperly Riemann integrable functions need not be Lebesgue integrable. On the other hand, there are Lebesgue integrable functions that are not Riemann integrable (such as $\mathbb{1}_{\mathbb{Q}}$). The geometric interpretation is that the Riemann integral respects the geometry of the integration domain by being defined via thinner and thinner vertical rectangles. On the other hand, the Lebesgue integral respects the geometry of the range by being defined via thinner and thinner horizontal strips. In particular, the Lebesgue integral does not make any assumption on the geometry of the domain and is thus more universal than the Riemann integral. In order to underline this, we present the following theorem that will also be useful later.

Theorem 4.26 Let $f : \Omega \rightarrow \mathbb{R}$ be measurable and $f \geq 0$ almost everywhere. Then

$$\sum_{n=1}^{\infty} \mu(\{f \geq n\}) \leq \int f d\mu \leq \sum_{n=0}^{\infty} \mu(\{f > n\}) \quad (4.7)$$

and

$$\int f d\mu = \int_0^\infty \mu(\{f \geq t\}) dt \quad (4.8)$$

Proof Define $f' = \lfloor f \rfloor$ and $f'' = \lceil f \rceil$. Then $f' \leq f \leq f''$ and hence $\int f' d\mu \leq \int f d\mu \leq \int f'' d\mu$. Now the first inequality of (4.7) follows from

$$\begin{aligned} \int f' d\mu &= \sum_{k=1}^{\infty} \mu(\{f' = k\}) \cdot k = \sum_{k=1}^{\infty} \sum_{n=1}^k \mu(\{f' = k\}) \\ &= \sum_{n=1}^{\infty} \sum_{k=n}^{\infty} \mu(\{f' = k\}) \\ &= \sum_{n=1}^{\infty} \mu(\{f' \geq n\}) = \sum_{n=1}^{\infty} \mu(\{f \geq n\}). \end{aligned}$$

Similarly, we infer the second inequality in (4.7) from

$$\int f'' d\mu = \sum_{n=1}^{\infty} \mu(\{f'' \geq n\}) = \sum_{n=1}^{\infty} \mu(\{f > n-1\})$$

If $g(t) := \mu(\{f \geq t\}) = \infty$ for some $t > 0$, then both sides in (4.8) equal ∞ . Hence, in the following, assume $g(t) < \infty$ for all $t > 0$.

For $\varepsilon > 0$ and $k \in \mathbb{N}$, define $g^\varepsilon := g \wedge g(\varepsilon)$, $f^\varepsilon := f \mathbb{1}_{\{f \geq \varepsilon\}}$ and $f_k^\varepsilon = 2^k f^\varepsilon$ as well as

$$\alpha_k^\varepsilon := 2^{-k} \sum_{n=1}^{\infty} \mu(\{f^\varepsilon \geq n 2^{-k}\}).$$

Then $\alpha_k^\varepsilon \xrightarrow{k \rightarrow \infty} \int_0^\infty g^\varepsilon(t) dt$. Furthermore, by (4.7) (with f replaced by f_k^ε), we have

$$\begin{aligned} \alpha_k^\varepsilon &= 2^{-k} \sum_{n=1}^{\infty} \mu(\{f_k^\varepsilon \geq n\}) \leq \int f_k^\varepsilon d\mu \\ &\leq 2^{-k} \sum_{n=0}^{\infty} \mu(\{f_k^\varepsilon > n\}) = 2^{-k} \sum_{n=0}^{\infty} \mu(\{f^\varepsilon > n 2^{-k}\}) \leq \alpha_k^\varepsilon + 2^{-k} g(\varepsilon) \end{aligned}$$

Since $2^{-k}g(\varepsilon) \xrightarrow{k \rightarrow \infty} 0$, we get

$$\int_0^\infty g^\varepsilon(t)dt = \int f^\varepsilon d\mu$$

Since $f^\varepsilon \uparrow f$ and $g^\varepsilon \uparrow g$ for $\varepsilon \downarrow 0$, the monotone convergence theorem implies (4.8).

Exercise 4.3.1 Let $f : [0, 1] \rightarrow \mathbb{R}$ be bounded. Show that f is (properly) Riemann integrable if and only if f is λ -a.e. continuous.

Exercise 4.3.2 If $f : [0, 1] \rightarrow \mathbb{R}$ is Riemann integrable, then f is Lebesgue measurable. Give an example that shows that f need not be Borel measurable.

Hint: Without proof, use the existence of a subset of $[0, 1]$ that is not Borel measurable. Based on this, construct a set that is not Borel and whose closure is a null set.

Exercise 4.3.3 Let $f : [0, 1] \rightarrow (0, \infty)$ be Riemann integrable. Without using the equivalence of the Lebesgue integral and the Riemann integral, show that $\int_0^1 f(x)dx > 0$.

5.5 5 Moments and Laws of Large Numbers

The most important characteristic quantities of random variables are the median, expectation and variance. For large n , the expectation describes the typical approximate value of the arithmetic mean $(X_1 + \dots + X_n)/n$ of i.i.d. random variables (law of large numbers). In Chapter 15, we will see how the variance determines the size of the typical deviations of the arithmetic mean from the expectation.

5.5.1 5.1 Moments

In the following, let $(\Omega, \mathcal{A}, \mathbf{P})$ be a probability space.

Definition 5.1 Let X be a real-valued random variable.

(i) If $X \in \mathcal{L}^1(\mathbf{P})$, then X is called integrable and we call

$$\mathbf{E}[X] := \int X d\mathbf{P}$$

the expectation or mean of X . If $\mathbf{E}[X] = 0$, then X is called centered. More generally, we also write $\mathbf{E}[X] = \int X d\mathbf{P}$ if only X^- or X^+ is integrable.

(ii) If $n \in \mathbb{N}$ and $X \in \mathcal{L}^n(\mathbf{P})$, then the quantities

$$m_k := \mathbf{E}[X^k], \quad M_k := \mathbf{E}[|X|^k] \quad \text{for any } k = 1, \dots, n,$$

are called the k th moments and k th absolute moments, respectively, of X .

(iii) If $X \in \mathcal{L}^2(\mathbf{P})$, then X is called square integrable and

$$\text{Var}[X] := \mathbf{E}[X^2] - \mathbf{E}[X]^2$$

is the variance of X . The number $\sigma := \sqrt{\text{Var}[X]}$ is called the standard deviation of X . Formally, we sometimes write $\text{Var}[X] = \infty$ if $\mathbf{E}[X^2] = \infty$.

(iv) If $X, Y \in \mathcal{L}^2(\mathbf{P})$, then we define the covariance of X and Y by

$$\text{Cov}[X, Y] := \mathbf{E}[(X - \mathbf{E}[X])(Y - \mathbf{E}[Y])].$$

X and Y are called uncorrelated if $\text{Cov}[X, Y] = 0$ and correlated otherwise.

Remark 5.2

- (i) The definition in (ii) is sensible since, by virtue of Theorem 4.19, $X \in \mathcal{L}^n(\mathbf{P})$ implies that $M_k < \infty$ for all $k = 1, \dots, n$.
- (ii) If $X, Y \in \mathcal{L}^2(\mathbf{P})$, then $XY \in \mathcal{L}^1(\mathbf{P})$ since $|XY| \leq X^2 + Y^2$. Hence the definition in (iv) makes sense and we have

$$\text{Cov}[X, Y] = \mathbf{E}[XY] - \mathbf{E}[X]\mathbf{E}[Y].$$

In particular, $\text{Var}[X] = \text{Cov}[X, X]$.

We collect the most important rules of expectations in a theorem. All of these properties are direct consequences of the corresponding properties of the integral.

Theorem 5.3 (Rules for expectations) Let $X, Y, X_n, Z_n, n \in \mathbb{N}$, be real integrable random variables on $(\Omega, \mathcal{A}, \mathbf{P})$.

- (i) If $\mathbf{P}_X = \mathbf{P}_Y$, then $\mathbf{E}[X] = \mathbf{E}[Y]$.
- (ii) (Linearity) Let $c \in \mathbb{R}$. Then $cX \in \mathcal{L}^1(\mathbf{P})$ and $X + Y \in \mathcal{L}^1(\mathbf{P})$ as well as

$$\mathbf{E}[cX] = c\mathbf{E}[X] \text{ and } \mathbf{E}[X + Y] = \mathbf{E}[X] + \mathbf{E}[Y].$$

- (iii) If $X \geq 0$ almost surely, then

$$\mathbf{E}[X] = 0 \iff X = 0 \text{ almost surely.}$$

- (iv) (Monotonicity) If $X \leq Y$ almost surely, then $\mathbf{E}[X] \leq \mathbf{E}[Y]$ with equality if and only if $X = Y$ almost surely.

- (v) (Triangle inequality) $|\mathbf{E}[X]| \leq \mathbf{E}[|X|]$.

- (vi) If $X_n \geq 0$ almost surely for all $n \in \mathbb{N}$, then $\mathbf{E}[\sum_{n=1}^{\infty} X_n] = \sum_{n=1}^{\infty} \mathbf{E}[X_n]$.

- (vii) If $Z_n \uparrow Z$ for some Z , then $\mathbf{E}[Z] = \lim_{n \rightarrow \infty} \mathbf{E}[Z_n] \in (-\infty, \infty]$.

Again probability theory comes into play when independence enters the stage; that is, when we exit the realm of linear integration theory.

Theorem 5.4 (Independent random variables are uncorrelated) Let $X, Y \in \mathcal{L}^1(\mathbf{P})$ be independent. Then $(XY) \in \mathcal{L}^1(\mathbf{P})$ and $\mathbf{E}[XY] = \mathbf{E}[X]\mathbf{E}[Y]$. In particular, independent random variables are uncorrelated.

Proof Assume first that X and Y take only finitely many values. Then XY also takes only finitely many values and thus $XY \in \mathcal{L}^1(\mathbf{P})$. It follows that

$$\begin{aligned} \mathbf{E}[XY] &= \sum_{z \in \mathbb{R} \setminus \{0\}} z \mathbf{P}[XY = z] \\ &= \sum_{z \in \mathbb{R} \setminus \{0\}} \sum_{x \in \mathbb{R} \setminus \{0\}} x \frac{z}{x} \mathbf{P}[X = x, Y = z/x] \\ &= \sum_{y \in \mathbb{R} \setminus \{0\}} \sum_{x \in \mathbb{R} \setminus \{0\}} xy \mathbf{P}[X = x] \mathbf{P}[Y = y] \\ &= \left(\sum_{x \in \mathbb{R}} x \mathbf{P}[X = x] \right) \left(\sum_{y \in \mathbb{R}} y \mathbf{P}[Y = y] \right) \\ &= \mathbf{E}[X] \mathbf{E}[Y] \end{aligned}$$

For $N \in \mathbb{N}$, the random variables $X_N := (2^{-N} \lfloor 2^N |X| \rfloor) \wedge N$ and $Y_N := (2^{-N} \lfloor 2^N |Y| \rfloor) \wedge N$ take only finitely many values and are independent as well. Furthermore, $X_N \uparrow |X|$ and $Y_N \uparrow |Y|$. By the monotone convergence theorem (Theorem 4.20), we infer

$$\begin{aligned}\mathbf{E}[|XY|] &= \lim_{N \rightarrow \infty} \mathbf{E}[X_N Y_N] = \lim_{N \rightarrow \infty} \mathbf{E}[X_N] \mathbf{E}[Y_N] \\ &= \left(\lim_{N \rightarrow \infty} \mathbf{E}[X_N] \right) \left(\lim_{N \rightarrow \infty} \mathbf{E}[Y_N] \right) = \mathbf{E}[|X|] \mathbf{E}[|Y|] < \infty\end{aligned}$$

Hence $XY \in \mathcal{L}^1(\mathbf{P})$. Furthermore, we have shown the claim in the case where X and Y are nonnegative. Hence (and since each of the families $\{X^+, Y^+\}$, $\{X^-, Y^+\}$, $\{X^+, Y^-\}$ and $\{X^-, Y^-\}$ is independent) we obtain

$$\begin{aligned}\mathbf{E}[XY] &= \mathbf{E}[(X^+ - X^-)(Y^+ - Y^-)] \\ &= \mathbf{E}[X^+ Y^+] - \mathbf{E}[X^- Y^+] - \mathbf{E}[X^+ Y^-] + \mathbf{E}[X^- Y^-] \\ &= \mathbf{E}[X^+] \mathbf{E}[Y^+] - \mathbf{E}[X^-] \mathbf{E}[Y^+] - \mathbf{E}[X^+] \mathbf{E}[Y^-] + \mathbf{E}[X^-] \mathbf{E}[Y^-] \\ &= \mathbf{E}[X^+ - X^-] \mathbf{E}[Y^+ - Y^-] = \mathbf{E}[X] \mathbf{E}[Y].\end{aligned}$$

Theorem 5.5 (Wald's identity) Let $T, X_1, X_2, \dots$ be independent real random variables in $\mathcal{L}^1(\mathbf{P})$. Let $\mathbf{P}[T \in \mathbb{N}_0] = 1$ and assume that $X_1, X_2, \dots$ are identically distributed. Define

$$S_T := \sum_{i=1}^T X_i$$

Then $S_T \in \mathcal{L}^1(\mathbf{P})$ and $\mathbf{E}[S_T] = \mathbf{E}[T] \mathbf{E}[X_1]$.

Proof Define $S_n = \sum_{i=1}^n X_i$ for $n \in \mathbb{N}_0$. Then $S_T = \sum_{n=1}^{\infty} S_n \mathbb{1}_{\{T=n\}}$. By Remark 2.15, the random variables S_n and $\mathbb{1}_{\{T=n\}}$ are independent for any $n \in \mathbb{N}$ and thus uncorrelated. This implies (using the triangle inequality; see Theorem 5.3(v))

$$\begin{aligned}\mathbf{E}[|S_T|] &= \sum_{n=1}^{\infty} \mathbf{E}[|S_n| \mathbb{1}_{\{T=n\}}] = \sum_{n=1}^{\infty} \mathbf{E}[|S_n|] \mathbf{E}[\mathbb{1}_{\{T=n\}}] \\ &\leq \sum_{n=1}^{\infty} \mathbf{E}[|X_1|] n \mathbf{P}[T=n] = \mathbf{E}[|X_1|] \mathbf{E}[T]\end{aligned}$$

The same computation without absolute values yields the remaining part of the claim.

We collect some basic properties of the variance.

Theorem 5.6 Let $X \in \mathcal{L}^2(\mathbf{P})$. Then:

- (i) $\text{Var}[X] = \mathbf{E}[(X - \mathbf{E}[X])^2] \geq 0$.
- (ii) $\text{Var}[X] = 0 \iff X = \mathbf{E}[X]$ almost surely.
- (iii) The map $f : \mathbb{R} \rightarrow \mathbb{R}, x \mapsto \mathbf{E}[(X - x)^2]$ is minimal at $x_0 = \mathbf{E}[X]$ with $f(\mathbf{E}[X]) = \text{Var}[X]$.

Proof (i) This is a direct consequence of Remark 5.2(ii).

(ii) By Theorem 5.3(iii), we have $\mathbf{E}[(X - \mathbf{E}[X])^2] = 0 \iff (X - \mathbf{E}[X])^2 = 0$ a.s.

(iii) Clearly, $f(x) = \mathbf{E}[X^2] - 2x\mathbf{E}[X] + x^2 = \text{Var}[X] + (x - \mathbf{E}[X])^2$.

Theorem 5.7 The map $\mathbf{Cov} : \mathcal{L}^2(\mathbf{P}) \times \mathcal{L}^2(\mathbf{P}) \rightarrow \mathbb{R}$ is a positive semidefinite symmetric bilinear form and $\text{Cov}[X, Y] = 0$ if Y is almost surely constant. The detailed version of this concise statement is: Let $X_1, \dots, X_m, Y_1, \dots, Y_n \in \mathcal{L}^2(\mathbf{P})$ and $\alpha_1, \dots, \alpha_m, \beta_1, \dots, \beta_n \in \mathbb{R}$ as well as $d, e \in \mathbb{R}$. Then

$$\text{Cov} \left[d + \sum_{i=1}^m \alpha_i X_i, e + \sum_{j=1}^n \beta_j Y_j \right] = \sum_{i,j} \alpha_i \beta_j \text{Cov}[X_i, Y_j] \quad (5.1)$$

In particular, $\text{Var}[\alpha X] = \alpha^2 \text{Var}[X]$ and the Bienaymé formula holds,

$$\text{Var} \left[\sum_{i=1}^m X_i \right] = \sum_{i=1}^m \text{Var} [X_i] + \sum_{\substack{i,j=1 \\ i \neq j}}^m \text{Cov} [X_i, X_j]. \quad (5.2)$$

For uncorrelated $X_1, \dots, X_m$, we have $\text{Var} [\sum_{i=1}^m X_i] = \sum_{i=1}^m \text{Var} [X_i]$.

Proof

$$\begin{aligned} & \text{Cov} \left[d + \sum_{i=1}^m \alpha_i X_i, e + \sum_{j=1}^n \beta_j Y_j \right] \\ &= \mathbf{E} \left[\left(\sum_{i=1}^m \alpha_i (X_i - \mathbf{E}[X_i]) \right) \left(\sum_{j=1}^n \beta_j (Y_j - \mathbf{E}[Y_j]) \right) \right] \\ &= \sum_{i=1}^m \sum_{j=1}^n \alpha_i \beta_j \mathbf{E} [(X_i - \mathbf{E}[X_i]) (Y_j - \mathbf{E}[Y_j])] \\ &= \sum_{i=1}^m \sum_{j=1}^n \alpha_i \beta_j \mathbf{Cov} [X_i, Y_j] \end{aligned}$$

Theorem 5.8 (Cauchy-Schwarz inequality) If $X, Y \in \mathcal{L}^2(\mathbf{P})$, then

$$(\text{Cov}[X, Y])^2 \leq \text{Var}[X] \text{Var}[Y]$$

Equality holds if and only if there are $a, b, c \in \mathbb{R}$ with $|a| + |b| + |c| > 0$ and such that $aX + bY + c = 0$ a.s.

Proof The Cauchy-Schwarz inequality holds for any positive semidefinite bilinear form and hence in particular for the covariance map. Using the notation of variance and covariance, a simple proof looks like this:

Case 1: $\text{Var}[Y] = 0$. Here the statement is trivial (choose $a = 0, b = 1$ and $c = -\mathbf{E}[Y]$).

Case 2: $\text{Var}[Y] > 0$. Let $\theta := -\frac{\text{Cov}[X, Y]}{\text{Var}[Y]}$. Then, by Theorem 5.6(i),

$$\begin{aligned} 0 &\leq \text{Var}[X + \theta Y] \text{Var}[Y] \\ &= (\text{Var}[X] + 2\theta \text{Cov}[X, Y] + \theta^2 \text{Var}[Y]) \text{Var}[Y] \\ &= \text{Var}[X] \text{Var}[Y] - \text{Cov}[X, Y]^2 \end{aligned}$$

with equality if and only if $X + \theta Y$ is a.s. constant. Now let $a = 1, b = \theta$ and $c = -\mathbf{E}[X] - b\mathbf{E}[Y]$.

Example 5.9

(i) Let $p \in [0, 1]$ and $X \sim \text{Ber}_p$. Then

$$\mathbf{E}[X^2] = \mathbf{E}[X] = \mathbf{P}[X = 1] = p$$

and thus $\text{Var}[X] = p(1 - p)$.

(ii) Let $n \in \mathbb{N}$ and $p \in [0, 1]$. Let X be binomially distributed, $X \sim b_{n,p}$. Then

$$\begin{aligned} \mathbf{E}[X] &= \sum_{k=0}^n k \mathbf{P}[X = k] = \sum_{k=0}^n k \binom{n}{k} p^k (1 - p)^{n-k} \\ &= np \cdot \sum_{k=1}^n \binom{n-1}{k-1} p^{k-1} (1 - p)^{(n-1)-(k-1)} = np. \end{aligned}$$

Furthermore,

$$\begin{aligned}
\mathbf{E}[X(X-1)] &= \sum_{k=0}^n k(k-1)\mathbf{P}[X=k] \\
&= \sum_{k=0}^n k(k-1)\binom{n}{k}p^k(1-p)^{n-k} \\
&= np \cdot \sum_{k=1}^n (k-1)\binom{n-1}{k-1}p^{k-1}(1-p)^{(n-1)-(k-1)} \\
&= n(n-1)p^2 \cdot \sum_{k=2}^n \binom{n-2}{k-2}p^{k-2}(1-p)^{(n-2)-(k-2)} \\
&= n(n-1)p^2
\end{aligned}$$

Hence $\mathbf{E}[X^2] = \mathbf{E}[X(X-1)] + \mathbf{E}[X] = n^2p^2 + np(1-p)$ and thus $\text{Var}[X] = np(1-p)$.

The statement can be derived more simply than by direct computation if we make use of the fact that $b_{n,p} = b_{1,p}^{*n}$ (see Example 3.4(ii)). That is (see Theorem 2.31), $\mathbf{P}_X = \mathbf{P}_{Y_1+\dots+Y_n}$, where $Y_1, \dots, Y_n$ are independent and $Y_i \sim \text{Ber}_p$ for any $i = 1, \dots, n$. Hence

$$\begin{aligned}
\mathbf{E}[X] &= n\mathbf{E}[Y_1] = np, \\
\text{Var}[X] &= n\text{Var}[Y_1] = np(1-p).
\end{aligned} \tag{5.3}$$

(iii) Let $\mu \in \mathbb{R}$ and $\sigma^2 > 0$, and let X be normally distributed, $X \sim \mathcal{N}_{\mu, \sigma^2}$. Then

$$\begin{aligned}
\mathbf{E}[X] &= \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} xe^{-(x-\mu)^2/(2\sigma^2)} dx \\
&= \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} (x+\mu)e^{-x^2/(2\sigma^2)} dx \\
&= \mu + \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} xe^{-x^2/(2\sigma^2)} dx = \mu
\end{aligned} \tag{5.4}$$

Similarly, we get $\text{Var}[X] = \mathbf{E}[X^2] - \mu^2 = \dots = \sigma^2$.

(iv) Let $\theta > 0$ and let X be exponentially distributed, $X \sim \exp_{\theta}$. Then

$$\begin{aligned}
\mathbf{E}[X] &= \theta \int_0^{\infty} xe^{-\theta x} dx = \frac{1}{\theta} \\
\text{Var}[X] &= -\theta^{-2} + \theta \int_0^{\infty} x^2 e^{-\theta x} dx = \theta^{-2} \left(-1 + \int_0^{\infty} x^2 e^{-x} dx \right) = \theta^{-2}
\end{aligned}$$

Theorem 5.10 (Blackwell-Girshick) Let $T, X_1, X_2, \dots$ be independent real random variables in $\mathcal{L}^2(\mathbf{P})$. Let $\mathbf{P}[T \in \mathbb{N}_0] = 1$ and let $X_1, X_2, \dots$ be identically distributed. Define

$$S_T := \sum_{i=1}^T X_i$$

Then $S_T \in \mathcal{L}^2(\mathbf{P})$ and

$$\text{Var}[S_T] = \mathbf{E}[X_1]^2 \text{Var}[T] + \mathbf{E}[T] \text{Var}[X_1]$$

Proof Define $S_n = \sum_{i=1}^n X_i$ for $n \in \mathbb{N}$. Then (as in the proof of Wald's identity) S_n and $\mathbb{1}_{\{T=n\}}$ are independent; hence S_n^2 and $\mathbb{1}_{\{T=n\}}$ are uncorrelated and thus

$$\begin{aligned}
\mathbf{E}[S_T^2] &= \sum_{n=0}^{\infty} \mathbf{E}[\mathbb{1}_{\{T=n\}} S_n^2] \\
&= \sum_{n=0}^{\infty} \mathbf{E}[\mathbb{1}_{\{T=n\}}] \mathbf{E}[S_n^2] \\
&= \sum_{n=0}^{\infty} \mathbf{P}[T=n] (\text{Var}[S_n] + \mathbf{E}[S_n]^2) \\
&= \sum_{n=0}^{\infty} \mathbf{P}[T=n] (n \text{Var}[X_1] + n^2 \mathbf{E}[X_1]^2) \\
&= \mathbf{E}[T] \text{Var}[X_1] + \mathbf{E}[T^2] \mathbf{E}[X_1]^2.
\end{aligned}$$

By Wald's identity (Theorem 5.5), we have $\mathbf{E}[S_T] = \mathbf{E}[T] \mathbf{E}[X_1]$; hence

$$\text{Var}[S_T] = \mathbf{E}[S_T^2] - \mathbf{E}[S_T]^2 = \mathbf{E}[T] \text{Var}[X_1] + (\mathbf{E}[T^2] - \mathbf{E}[T]^2) \mathbf{E}[X_1]^2$$

as claimed.

Exercise 5.1.1 Let X be an integrable real random variable whose distribution $\mathbf{P}_X$ has a density f (with respect to the Lebesgue measure λ). Show (using Theorem 4.15) that

$$\mathbf{E}[X] = \int_{\mathbb{R}} x f(x) \lambda(dx)$$

Exercise 5.1.2 Let $X \sim \beta_{r,s}$ be a Beta-distributed random variable with parameters $r, s > 0$ (see Example 1.107(ii)). Show that

$$\mathbf{E}[X^n] = \prod_{k=0}^{n-1} \frac{r+k}{r+s+k} \quad \text{for any } n \in \mathbb{N}$$

Exercise 5.1.3 Let $X_1, X_2, \dots$ be i.i.d. nonnegative random variables. By virtue of the Borel-Cantelli lemma, show that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} X_n = \begin{cases} 0 \text{ a.s.,} & \text{if } \mathbf{E}[X_1] < \infty, \\ \infty \text{ a.s.,} & \text{if } \mathbf{E}[X_1] = \infty. \end{cases}$$

Exercise 5.1.4 Let $X_1, X_2, \dots$ be i.i.d. nonnegative random variables. By virtue of the Borel-Cantelli lemma, show that for any $c \in (0, 1)$

$$\sum_{n=1}^{\infty} e^{X_n} c^n \begin{cases} < \infty \text{ a.s.,} & \text{if } \mathbf{E}[X_1] < \infty \\ = \infty \text{ a.s.,} & \text{if } \mathbf{E}[X_1] = \infty \end{cases}$$

5.5.2 5.2 Weak Law of Large Numbers

Theorem 5.11 (Markov inequality, Chebyshev inequality) Let X be a real random variable and let $f : [0, \infty) \rightarrow [0, \infty)$ be monotone increasing. Then for any $\varepsilon > 0$ with $f(\varepsilon) > 0$, the Markov inequality holds,

$$\mathbf{P}[|X| \geq \varepsilon] \leq \frac{\mathbf{E}[f(|X|)]}{f(\varepsilon)}$$

In the special case $f(x) = x^2$, we get $\mathbf{P}[|X| \geq \varepsilon] \leq \varepsilon^{-2} \mathbf{E}[X^2]$. In particular, if $X \in \mathcal{L}^2(\mathbf{P})$, the Chebyshev inequality holds:

$$\mathbf{P}[|X - \mathbf{E}[X]| \geq \varepsilon] \leq \varepsilon^{-2} \text{Var}[X].$$

Proof We have

$$\begin{aligned} \mathbf{E}[f(|X|)] &\geq \mathbf{E}[f(|X|)\mathbb{1}_{\{f(|X|) \geq f(\varepsilon)\}}] \\ &\geq \mathbf{E}[f(\varepsilon)\mathbb{1}_{\{f(|X|) \geq f(\varepsilon)\}}] \\ &\geq f(\varepsilon)\mathbf{P}[|X| \geq \varepsilon] \end{aligned}$$

Definition 5.12 Let $(X_n)_{n \in \mathbb{N}}$ be a sequence of real random variables in $\mathcal{L}^1(\mathbf{P})$ and let $\tilde{S}_n = \sum_{i=1}^n (X_i - \mathbf{E}[X_i])$.

(i) We say that $(X_n)_{n \in \mathbb{N}}$ fulfills the weak law of large numbers if

$$\lim_{n \rightarrow \infty} \mathbf{P}\left[\left|\frac{1}{n}\tilde{S}_n\right| > \varepsilon\right] = 0 \quad \text{for any } \varepsilon > 0$$

(ii) We say that $(X_n)_{n \in \mathbb{N}}$ fulfills the strong law of large numbers if

$$\mathbf{P}\left[\limsup_{n \rightarrow \infty} \left|\frac{1}{n}\tilde{S}_n\right| = 0\right] = 1$$

Remark 5.13 The strong law of large numbers implies the weak law. Indeed, if $A_n^\varepsilon := \left\{\left|\frac{1}{n}\tilde{S}_n\right| > \varepsilon\right\}$ and $A = \left\{\limsup_{n \rightarrow \infty} \left|\frac{1}{n}\tilde{S}_n\right| > 0\right\}$, then clearly

$$A = \bigcup_{m \in \mathbb{N}} \limsup_{n \rightarrow \infty} A_n^{1/m};$$

hence $\mathbf{P}[\limsup_{n \rightarrow \infty} A_n^\varepsilon] = 0$ for $\varepsilon > 0$. By Fatou's lemma (Theorem 4.21), we obtain

$$\begin{aligned} \limsup_{n \rightarrow \infty} \mathbf{P}[A_n^\varepsilon] &= 1 - \liminf_{n \rightarrow \infty} \mathbf{E}[\mathbb{1}_{(A_n^\varepsilon)^c}] \\ &\leq 1 - \mathbf{E}\left[\liminf_{n \rightarrow \infty} \mathbb{1}_{(A_n^\varepsilon)^c}\right] = \mathbf{E}\left[\limsup_{n \rightarrow \infty} \mathbb{1}_{A_n^\varepsilon}\right] = 0. \end{aligned}$$

Theorem 5.14 Let $X_1, X_2, \dots$ be uncorrelated random variables in $\mathcal{L}^2(\mathbf{P})$ with $V := \sup_{n \in \mathbb{N}} \text{Var}[X_n] < \infty$. Then $(X_n)_{n \in \mathbb{N}}$ fulfills the weak law of large numbers. More precisely, for any $\varepsilon > 0$, we have

$$\mathbf{P}\left[\left|\frac{1}{n}\tilde{S}_n\right| \geq \varepsilon\right] \leq \frac{V}{\varepsilon^2 n} \quad \text{for all } n \in \mathbb{N} \quad (5.5)$$

Proof Without loss of generality, assume $\mathbf{E}[X_i] = 0$ for all $i \in \mathbb{N}$ and thus $\tilde{S}_n = X_1 + \dots + X_n$. By Bienaymé's formula (Theorem 5.7), we obtain

$$\text{Var}\left[\frac{1}{n}\tilde{S}_n\right] = n^{-2} \sum_{i=1}^n \text{Var}[X_i] \leq \frac{V}{n}$$

By Chebyshev's inequality (Theorem 5.11), for any $\varepsilon > 0$,

$$\mathbf{P}\left[\left|\tilde{S}_n/n\right| \geq \varepsilon\right] \leq \frac{V}{\varepsilon^2 n} \xrightarrow{n \rightarrow \infty} 0$$

Example 5.15 (Weierstraß's approximation theorem) Let $f : [0, 1] \rightarrow \mathbb{R}$ be a continuous map. By Weierstraß's approximation theorem, there exist polynomials f_n of degree at most n such that

$$\|f_n - f\|_\infty \xrightarrow{n \rightarrow \infty} 0,$$

where $\|f\|_\infty := \sup\{|f(x)| : x \in [0, 1]\}$ denotes the supremum norm of $f \in C([0, 1])$ (the space of continuous functions $[0, 1] \rightarrow \mathbb{R}$).

We present a probabilistic proof of this theorem. For $n \in \mathbb{N}$, define the polynomial f_n by

$$f_n(x) := \sum_{k=0}^n f(k/n) \binom{n}{k} x^k (1-x)^{n-k} \quad \text{for } x \in [0, 1].$$

f_n is called the Bernstein polynomial of order n .

Fix $\varepsilon > 0$. As f is continuous on the compact interval $[0, 1]$, f is uniformly continuous. Hence there exists a $\delta > 0$ such that

$$|f(x) - f(y)| < \varepsilon \quad \text{for all } x, y \in [0, 1] \text{ with } |x - y| < \delta.$$

Now fix $p \in [0, 1]$ and let $X_1, X_2, \dots$ be independent random variables with $X_i \sim \text{Ber}_p$, $i \in \mathbb{N}$. Then $S_n := X_1 + \dots + X_n \sim b_{n,p}$ and thus

$$\mathbf{E}[f(S_n/n)] = \sum_{k=0}^n f(k/n) \mathbf{P}[S_n = k] = f_n(p)$$

We get

$$|f(S_n/n) - f(p)| \leq \varepsilon + 2\|f\|_\infty \mathbb{1}_{\{|(S_n/n) - p| \geq \delta\}}$$

and thus (by Theorem 5.14 with $V = p(1-p) \leq \frac{1}{4}$)

$$\begin{aligned} |f_n(p) - f(p)| &\leq \mathbf{E}[|f(S_n/n) - f(p)|] \\ &\leq \varepsilon + 2\|f\|_\infty \mathbf{P}\left[\left|\frac{S_n}{n} - p\right| \geq \delta\right] \\ &\leq \varepsilon + \frac{\|f\|_\infty}{2\delta^2 n} \end{aligned}$$

for any $p \in [0, 1]$. Hence $\|f_n - f\|_\infty \xrightarrow{n \rightarrow \infty} 0$.

Exercise 5.2.1 (Bernstein-Chernov bound) Let $n \in \mathbb{N}$ and $p_1, \dots, p_n \in [0, 1]$. Let $X_1, \dots, X_n$ be independent random variables with $X_i = \text{Ber}_{p_i}$ for any $i = 1, \dots, n$. Define $S_n = X_1 + \dots + X_n$ and $m := \mathbf{E}[S_n]$. Show that, for any $\delta > 0$, the following two estimates hold:

$$\mathbf{P}[S_n \geq (1 + \delta)m] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}}\right)^m$$

and

$$\mathbf{P}[S_n \leq (1 - \delta)m] \leq \exp\left(-\frac{\delta^2 m}{2}\right)$$

Hint: For S_n , use Markov's inequality with $f(x) = e^{\lambda x}$ for some $\lambda > 0$ and then find the λ that optimizes the bound.

5.5.3 5.3 Strong Law of Large Numbers

We show Etemadi's version [47] of the strong law of large numbers for identically distributed, pairwise independent random variables. There is a zoo of strong laws of large numbers, each of which varies in the exact assumptions it makes on the underlying sequence of random variables. For example, the assumption that the random variables be identically distributed can

be waived if other assumptions are introduced such as bounded variances. We do not strive for completeness but show only a few of the statements.

In order to illustrate the method of the proof of Etemadi's theorem, we first present (and prove) a strong law of large numbers under stronger assumptions.

Theorem 5.16 Let $X_1, X_2, \dots \in \mathcal{L}^2(\mathbf{P})$ be pairwise independent (that is, X_i and X_j are independent for all $i, j \in \mathbb{N}$ with $i \neq j$) and identically distributed. Then $(X_n)_{n \in \mathbb{N}}$ fulfills the strong law of large numbers.

Proof The random variables $(X_n^+)_{n \in \mathbb{N}}$ and $(X_n^-)_{n \in \mathbb{N}}$ again form pairwise independent families of square integrable random variables (compare Remark 2.15(ii)). Hence, it is enough to consider $(X_n^+)_{n \in \mathbb{N}}$. Thus we henceforth assume $X_n \geq 0$ almost surely for all $n \in \mathbb{N}$.

Let $S_n = X_1 + \dots + X_n$ for $n \in \mathbb{N}$. Fix $\varepsilon > 0$. For any $n \in \mathbb{N}$, define $k_n = \lfloor (1 + \varepsilon)^n \rfloor \geq \frac{1}{2}(1 + \varepsilon)^n$. Then, by Chebyshev's inequality (Theorem 5.11),

$$\begin{aligned} \sum_{n=1}^{\infty} \mathbf{P} \left[\left| \frac{S_{k_n}}{k_n} - \mathbf{E}[X_1] \right| \geq (1 + \varepsilon)^{-n/4} \right] \\ \leq \sum_{n=1}^{\infty} (1 + \varepsilon)^{n/2} \text{Var} [k_n^{-1} S_{k_n}] \\ = \sum_{n=1}^{\infty} (1 + \varepsilon)^{n/2} k_n^{-1} \text{Var} [X_1] \\ \leq 2 \text{Var} [X_1] \sum_{n=1}^{\infty} (1 + \varepsilon)^{-n/2} < \infty \end{aligned} \quad (5.6)$$

Thus, by the Borel-Cantelli lemma, for $\mathbf{P}$ -a.a. ω , there is an $n_0 = n_0(\omega)$ such that

$$\left| \frac{S_{k_n}}{k_n} - \mathbf{E}[X_1] \right| < (1 + \varepsilon)^{-n/4} \quad \text{for all } n \geq n_0$$

whence

$$\limsup_{n \rightarrow \infty} |k_n^{-1} S_{k_n} - \mathbf{E}[X_1]| = 0 \quad \text{almost surely.}$$

Note that $k_{n+1} \leq (1 + 2\varepsilon)k_n$ for sufficiently large $n \in \mathbb{N}$. For $l \in \{k_n, \dots, k_{n+1}\}$, we get

$$\frac{1}{1 + 2\varepsilon} k_n^{-1} S_{k_n} \leq k_{n+1}^{-1} S_{k_n} \leq l^{-1} S_l \leq k_n^{-1} S_{k_{n+1}} \leq (1 + 2\varepsilon) k_{n+1}^{-1} S_{k_{n+1}}.$$

Now $1 - (1 + 2\varepsilon)^{-1} \leq 2\varepsilon$ implies

$$\begin{aligned} \limsup_{l \rightarrow \infty} |l^{-1} S_l - \mathbf{E}[X_1]| &\leq \limsup_{n \rightarrow \infty} |k_n^{-1} S_{k_n} - \mathbf{E}[X_1]| + 2\varepsilon \limsup_{n \rightarrow \infty} k_n^{-1} S_{k_n} \\ &\leq 2\varepsilon \mathbf{E}[X_1] \quad \text{almost surely.} \end{aligned}$$

Hence the strong law of large numbers is in force.

The similarity of the variance estimates in the weak law of large numbers and in (5.6) suggests that in the preceding theorem the condition that the random variables $X_1, X_2, \dots$ be identically distributed could be replaced by the condition that the variances be bounded (see Exercise 5.3.1).

We can weaken the condition in Theorem 5.16 in a different direction by requiring integrability only instead of square integrability of the random variables.

Theorem 5.17 (Etemadi's strong law of large numbers (1981)) Let $X_1, X_2, \dots \in \mathcal{L}^1(\mathbf{P})$ be pairwise independent and identically distributed. Then $(X_n)_{n \in \mathbb{N}}$ fulfills the strong law of large numbers.

We follow the proof in [39, Section 2.4]. Define $\mu = \mathbf{E}[X_1]$ and $S_n = X_1 + \dots + X_n$. We start with some preparatory lemmas. (For the "a.s." notation see Definition 1.68.)

Lemma 5.18 For $n \in \mathbb{N}$, define $Y_n := X_n \mathbb{1}_{\{|X_n| \leq n\}}$ and $T_n = Y_1 + \dots + Y_n$. The sequence $(X_n)_{n \in \mathbb{N}}$ fulfills the strong law of large numbers if $T_n/n \xrightarrow{n \rightarrow \infty} \mu$ a.s.

Proof By Theorem 4.26, we have $\sum_{n=1}^{\infty} \mathbf{P}[|X_n| > n] \leq \mathbf{E}[|X_1|] < \infty$. Thus, by the Borel-Cantelli lemma,

$$\mathbf{P}[X_n \neq Y_n \text{ for infinitely many } n] = 0.$$

Hence there is an $n_0 = n_0(\omega)$ with $X_n = Y_n$ for all $n \geq n_0$, whence for $n \geq n_0$

$$\frac{T_n - S_n}{n} = \frac{T_{n_0} - S_{n_0}}{n} \xrightarrow{n \rightarrow \infty} 0.$$

Lemma 5.19 $\sum_{n \geq x} n^{-2} \leq 4$ for all $x \geq 0$.

Proof For $m \in \mathbb{N}$, by comparison with the corresponding integral, we get

$$\sum_{n=m}^{\infty} n^{-2} \leq m^{-2} + \int_m^{\infty} t^{-2} dt = m^{-2} + m^{-1} \leq \frac{2}{m}$$

Lemma 5.20 $\sum_{n=1}^{\infty} \frac{\mathbf{E}[Y_n^2]}{n^2} \leq 4\mathbf{E}[|X_1|]$.

Proof By Theorem 4.26,

$$\mathbf{E}[Y_n^2] = \int_0^{\infty} \mathbf{P}[Y_n^2 > t] dt$$

Substituting $x = \sqrt{t}$, we obtain

$$\mathbf{E}[Y_n^2] = \int_0^{\infty} 2x \mathbf{P}[|Y_n| > x] dx \leq \int_0^n 2x \mathbf{P}[|X_1| > x] dx$$

By Lemma 5.19, for $m \rightarrow \infty$,

$$f_m(x) = \left(\sum_{n=1}^m n^{-2} \mathbb{1}_{\{x < n\}} \right) 2x \mathbf{P}[|X_1| > x] \uparrow f(x) \leq 4\mathbf{P}[|X_1| > x]$$

Hence, by the monotone limit theorem, we can interchange the summation and the integral and obtain

$$\begin{aligned} \sum_{n=1}^{\infty} \frac{\mathbf{E}[Y_n^2]}{n^2} &\leq \sum_{n=1}^{\infty} n^{-2} \int_0^{\infty} \mathbb{1}_{\{x < n\}} 2x \mathbf{P}[|X_1| > x] dx \\ &= \int_0^{\infty} \left(\sum_{n=1}^{\infty} n^{-2} \mathbb{1}_{\{x < n\}} \right) 2x \mathbf{P}[|X_1| > x] dx \\ &\leq 4 \int_0^{\infty} \mathbf{P}[|X_1| > x] dx = 4\mathbf{E}[|X_1|] \end{aligned}$$

Proof of Theorem 5.17 As in the proof of Theorem 5.16, it is enough to consider the case $X_n \geq 0$. Fix $\varepsilon > 0$ and let $\alpha = 1 + \varepsilon$. For $n \in \mathbb{N}$, define $k_n = \lfloor \alpha^n \rfloor$. Note that $k_n \geq \alpha^n/2$. Hence, for all $m \in \mathbb{N}$ (with $n_0 = \lceil \log m / \log \alpha \rceil$),

$$\sum_{n:k_n \geq m} k_n^{-2} \leq 4 \sum_{n=n_0}^{\infty} \alpha^{-2n} = 4\alpha^{-2n_0} (1 - \alpha^{-2})^{-1} \leq 4(1 - \alpha^{-2})^{-1} m^{-2} \quad (5.7)$$

The aim is to employ Lemma 5.20 to refine the estimate (5.6) for $(Y_n)_{n \in \mathbb{N}}$ and $(T_n)_{n \in \mathbb{N}}$. For $\delta > 0$, Chebyshev's inequality yields (together with (5.7))

$$\begin{aligned} & \sum_{n=1}^{\infty} \mathbf{P} [|T_{k_n} - \mathbf{E}[T_{k_n}]| > \delta k_n] \\ & \leq \delta^{-2} \sum_{n=1}^{\infty} \frac{\mathbf{Var}[T_{k_n}]}{k_n^2} \\ & = \delta^{-2} \sum_{n=1}^{\infty} k_n^{-2} \sum_{m=1}^{k_n} \mathbf{Var}[Y_m] = \delta^{-2} \sum_{m=1}^{\infty} \mathbf{Var}[Y_m] \sum_{n:k_n \geq m} k_n^{-2} \\ & \leq 4(1 - \alpha^{-2})^{-1} \delta^{-2} \sum_{m=1}^{\infty} m^{-2} \mathbf{E}[Y_m^2] < \infty \quad \text{by Lemma 5.20.} \end{aligned}$$

(In the third step, we could change the order of summation since all summands are nonnegative.) Letting $\delta \downarrow 0$, we infer by the Borel-Cantelli lemma

$$\lim_{n \rightarrow \infty} \frac{T_{k_n} - \mathbf{E}[T_{k_n}]}{k_n} = 0 \quad \text{almost surely.} \quad (5.8)$$

By the monotone convergence theorem (Theorem 4.20), we have

$$\mathbf{E}[Y_n] = \mathbf{E}[X_1 \mathbb{1}_{\{X_1 \leq n\}}] \xrightarrow{n \rightarrow \infty} \mathbf{E}[X_1].$$

Hence $\mathbf{E}[T_{k_n}]/k_n \xrightarrow{n \rightarrow \infty} \mathbf{E}[X_1]$. By (5.8), we also have $T_{k_n}/k_n \xrightarrow{n \rightarrow \infty} \mathbf{E}[X_1]$ a.s. As in the proof of Theorem 5.16, we also get (since $Y_n \geq 0$)

$$\lim_{l \rightarrow \infty} \frac{T_l}{l} = \mathbf{E}[X_1] \quad \text{almost surely.}$$

By Lemma 5.18, this implies the claim of Theorem 5.17.

Example 5.21 (Monte Carlo integration) Let $f : [0, 1] \rightarrow \mathbb{R}$ be a function and assume we want to determine the value of its integral $I := \int_0^1 f(x) dx$ numerically. Assume that the computer generates numbers $X_1, X_2, \dots$ that can be considered as independent random numbers, uniformly distributed on $[0, 1]$. For $n \in \mathbb{N}$, define the estimated value

$$\hat{I}_n := \frac{1}{n} \sum_{i=1}^n f(X_i).$$

Assuming $f \in \mathcal{L}^1([0, 1])$, the strong law of large numbers yields $\hat{I}_n \xrightarrow{n \rightarrow \infty} I$ a.s.

Note that the last theorem made no statement on the speed of convergence. That is, we do not have control on the quantity $\mathbf{P} \left[\left| \hat{I}_n - I \right| > \varepsilon \right]$. In order to get more precise estimates for the integral, we need additional information; for example, the value $V_1 := \int f^2(x) dx - I^2$ if $f \in \mathcal{L}^2([0, 1])$. (For bounded f , V_1 can easily be bounded.) Indeed, in this case, $\mathbf{Var}[\hat{I}_n] = V_1/n$; hence, by Chebyshev's inequality,

$$\mathbf{P} \left[\left| \hat{I}_n - I \right| > \varepsilon n^{-1/2} \right] \leq V_1/\varepsilon^2.$$

Hence the error is at most of order $n^{-1/2}$. The central limit theorem will show that the error is indeed exactly of this order.

If f is smooth in some sense, then the usual numerical procedures yield better orders of convergence. Hence Monte Carlo simulation should be applied only if all other methods fail. This is the case in particular if $[0, 1]$ is replaced by $G \subset \mathbb{R}^d$ for very large d .

Definition 5.22 (Empirical distribution function) Let $X_1, X_2, \dots$ be real random variables. The map $F_n : \mathbb{R} \rightarrow [0, 1], x \mapsto \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{(-\infty, x]}(X_i)$ is called the empirical distribution function of $X_1, \dots, X_n$.

Theorem 5.23 (Glivenko-Cantelli) Let $X_1, X_2, \dots$ be i.i.d. real random variables with distribution function F , and let $F_n, n \in \mathbb{N}$, be the empirical distribution functions. Then

$$\limsup_{n \rightarrow \infty} \sup_{x \in \mathbb{R}} |F_n(x) - F(x)| = 0 \quad \text{almost surely.}$$

Proof Fix $x \in \mathbb{R}$ and let $Y_n(x) = \mathbb{1}_{(-\infty, x]}(X_n)$ and $Z_n(x) = \mathbb{1}_{(-\infty, x)}(X_n)$ for $n \in \mathbb{N}$. Additionally, define the left-sided limits $F(x-) = \lim_{y \uparrow x} F(y)$ and similarly for F_n . Then each of the families $(Y_n(x))_{n \in \mathbb{N}}$ and $(Z_n(x))_{n \in \mathbb{N}}$ is independent. Furthermore, $\mathbf{E}[Y_n(x)] = \mathbf{P}[X_n \leq x] = F(x)$ and $\mathbf{E}[Z_n(x)] = \mathbf{P}[X_n < x] = F(x-)$. By the strong law of large numbers, we thus have

$$F_n(x) = \frac{1}{n} \sum_{i=1}^n Y_i(x) \xrightarrow{n \rightarrow \infty} F(x) \quad \text{almost surely}$$

and

$$F_n(x-) = \frac{1}{n} \sum_{i=1}^n Z_i(x) \xrightarrow{n \rightarrow \infty} F(x-) \quad \text{almost surely.}$$

Formally, define $F(-\infty) = 0$ and $F(\infty) = 1$. Fix some $N \in \mathbb{N}$ and define

$$x_j := \inf\{x \in \overline{\mathbb{R}} : F(x) \geq j/N\}, \quad j = 0, \dots, N,$$

and

$$R_n := \max_{j=1, \dots, N-1} (|F_n(x_j) - F(x_j)| + |F_n(x_j-) - F(x_j-)|).$$

As shown above, $R_n \xrightarrow{n \rightarrow \infty} 0$ almost surely. For $x \in (x_{j-1}, x_j)$, we have (by definition of x_j)

$$F_n(x) \leq F_n(x_j-) \leq F(x_j-) + R_n \leq F(x) + R_n + \frac{1}{N}$$

and

$$F_n(x) \geq F_n(x_{j-1}) \geq F(x_{j-1}) - R_n \geq F(x) - R_n - \frac{1}{N}$$

Hence

$$\limsup_{n \rightarrow \infty} \sup_{x \in \mathbb{R}} |F_n(x) - F(x)| \leq \frac{1}{N} + \limsup_{n \rightarrow \infty} R_n = \frac{1}{N}.$$

Letting $N \rightarrow \infty$, the claim follows.

Example 5.24 (Shannon's theorem) Consider a source of information that sends a sequence of independent random symbols $X_1, X_2, \dots$ drawn from a finite alphabet E (that is, from an arbitrary finite set E). Let p_e be the probability of the symbol $e \in E$. Formally, the $X_1, X_2, \dots$ are i.i.d. E -valued random variables with $\mathbf{P}[X_i = e] = p_e$ for $e \in E$.

For any $\omega \in \Omega$ and $n \in \mathbb{N}$, let

$$\pi_n(\omega) := \prod_{i=1}^n p_{X_i(\omega)}$$

be the probability that the observed sequence $X_1(\omega), \dots, X_n(\omega)$ occurs. Define $Y_n(\omega) := -\log(p_{X_n(\omega)})$. Then $(Y_n)_{n \in \mathbb{N}}$ is i.i.d. and $\mathbf{E}[Y_n] = H(p)$, where

$$H(p) := - \sum_{e \in E} p_e \log(p_e)$$

is the entropy of the distribution $p = (p_e)_{e \in E}$ (compare Definition 5.25). By the strong law of large numbers, we infer Shannon's theorem:

$$-\frac{1}{n} \log \pi_n = \frac{1}{n} \sum_{i=1}^n Y_i \xrightarrow{n \rightarrow \infty} H(p) \quad \text{almost surely}$$

Entropy and Source Coding Theorem*

We briefly discuss the importance of π_n and the entropy. How can we quantify the information inherent in a message $X_1(\omega), \dots, X_n(\omega)$? This information can be measured by the length of the shortest sequence of zeros and ones by which the message can be encoded. Of course, you do not want to invent a new code for

every message but rather use one code that allows for the shortest average coding of the messages for the particular information source. To this end, associate with each symbol $e \in E$ a sequence of zeros and ones that when concatenated yield the message. The length $l(e)$ of the sequence that codes for e may depend on e . Hence, for efficiency, those symbols that appear more often get a shorter code than the more rare symbols. The Morse alphabet is constructed similarly (the letters "e" and "t", which are the most frequent letters in English, have the shortest codes ("dot" and "dash"), and the rare letter "q" has the code "dash-dash-dot-dash"). However, the Morse code also consists of gaps of different lengths that signal ends of letters and words. As we want to use only zeros and ones (and no gap-like symbols), we have to arrange the code in such a way that no code is the beginning of the code of a different symbol. For example, we could not encode one symbol with 0110 and a different one with 011011. A code that fulfills this condition is called a binary prefix code. Denote by $c(e) \in \{0, 1\}^{l(e)}$ the code of e , where $l(e)$ is its length. We can represent the codes of all letters in a tree.

Let us construct a code $C = (c(e), e \in E)$ that is efficient in the sense that it minimizes the expected length of the code (of a random symbol)

$$L_p(C) := \sum_{e \in E} p_e l(e)$$

We first define a specific code and then show that it is almost optimal. As a first step, we enumerate $E = \{e_1, \dots, e_N\}$ such that $p_{e_1} \geq p_{e_2} \geq \dots \geq p_{e_N}$. Define $\ell(e) \in \mathbb{N}$ for any $e \in E$ by

$$2^{-\ell(e)} \leq p_e < 2^{-\ell(e)+1}$$

Let $\tilde{p}_e = 2^{-\ell(e)}$ for any $e \in E$ and let $\tilde{q}_k = \sum_{l < k} \tilde{p}_{e_l}$ for $k = 1, \dots, N$.

By construction, $\ell(e_l) \leq \ell(e_k)$ for all $l \leq k$; hence the binary representation of $\tilde{q}_k$ has at most $\ell(e_k)$ digits:

$$\tilde{q}_k = \sum_{i=1}^{\ell(e_k)} c_i(e_k) 2^{-i}$$

Here the numbers $c_1(e_k), \dots, c_{\ell(e_k)}(e_k) \in \{0, 1\}$ are uniquely determined. Clearly, $\tilde{q}_l \geq \tilde{q}_k + 2^{-\ell(e_k)}$ for any $l > k$; hence

$$(c_1(e_k), \dots, c_{\ell(e_k)}(e_k)) \neq (c_1(e_l), \dots, c_{\ell(e_k)}(e_l)) \quad \text{for all } l > k.$$

Thus $C = (c(e), e \in E)$ is a prefix code.

For any $b > 0$ and $x > 0$, denote by $\log_b(x) := \frac{\log(x)}{\log(b)}$ the logarithm of x to base b . By construction, $-\log_2(p_e) \leq l(e) \leq 1 - \log_2(p_e)$. Hence the expected length is

$$-\sum_{e \in E} p_e \log_2(p_e) \leq L_p(C) \leq 1 - \sum_{e \in E} p_e \log_2(p_e).$$

The length of this code for the first n symbols of our random information source is thus approximately $-\sum_{k=1}^n \log_2(p_{X_k(\omega)}) = -\log_2 \pi_n(\omega)$. Here we have the connection to Shannon's theorem. That theorem thus makes a statement about the length of a binary prefix code needed to transmit a long message.

Now, is the code constructed above optimal, or are there codes with smaller mean length? The answer is given by the source coding theorem for which we prepare with a definition and a lemma.

Definition 5.25 (Entropy) Let $p = (p_e)_{e \in E}$ be a probability distribution on the countable set E . For $b > 0$, define

$$H_b(p) := -\sum_{e \in E} p_e \log_b(p_e)$$

with the convention $0 \log_b(0) := 0$. We call $H(p) := H_e(p)$ ($e = 2.71\dots$ Euler's number) the entropy and $H_2(p)$ the binary entropy of p .

Note that, for infinite E , the entropy need not be finite.

Lemma 5.26 (Entropy inequality) Let b and p be as above. Further, let q be a sub-probability distribution; that is, $q_e \geq 0$ for all $e \in E$ and $\sum_{e \in E} q_e \leq 1$. Then

$$H_b(p) \leq -\sum_{e \in E} p_e \log_b(q_e) \tag{5.9}$$

with equality if and only if $H_b(p) = \infty$ or $q = p$.

Proof Without loss of generality, we can do the computation with $b = e$; that is, with the natural logarithm. Note that $\log(1+x) \leq x$ for $x > -1$ with equality if and only if $x = 0$. If in (5.9) the left-hand side is finite, then we can subtract the right-hand side from the left-hand side and obtain

$$\begin{aligned} H(p) + \sum_{e \in E} p_e \log(q_e) &= \sum_{e: p_e > 0} p_e \log(q_e/p_e) \\ &= \sum_{e: p_e > 0} p_e \log\left(1 + \frac{q_e - p_e}{p_e}\right) \\ &\leq \sum_{e: p_e > 0} p_e \frac{q_e - p_e}{p_e} = \sum_{e \in E} (q_e - p_e) \leq 0 \end{aligned}$$

If $q \neq p$, then there is an $e \in E$ with $p_e > 0$ and $q_e \neq p_e$. If this is the case, then strict inequality holds if $H(p) < \infty$.

Theorem 5.27 (Source coding theorem) Let $p = (p_e)_{e \in E}$ be a probability distribution on the finite alphabet E . For any binary prefix code $C = (c(e), e \in E)$, we have $L_p(C) \geq H_2(p)$. Furthermore, there is a binary prefix code C with $L_p(C) \leq H_2(p) + 1$.

Proof The second part of the theorem was shown in the above construction. Now assume that a prefix code is given. Let $L = \max_{e \in E} l(e)$. For $e \in E$, let

$$C_L(e) = \{c \in \{0, 1\}^L : c_k = c_k(e) \text{ for } k \leq l(e)\}$$

the set of all dyadic sequences of length L that start like $c(e)$. Since we have a prefix code, the sets $C_L(e), e \in E$, are pairwise disjoint and $\bigcup_{e \in E} C_L(e) \subset \{0, 1\}^L$. Hence, if we define $q_e := 2^{-l(e)}$, then (note that $\#C_L(e) = 2^{L-l(e)}$)

$$\sum_{e \in E} q_e = 2^{-L} \sum_{e \in E} \#C_L(e) \leq 1.$$

By Lemma 5.26, we have $L_p(C) = \sum_{e \in E} p_e l(e) = -\sum_{e \in E} p_e \log_2(q_e) \geq H_2(p)$.

Exercise 5.3.1 Show the following improvement of Theorem 5.16: If $X_1, X_2, \dots \in \mathcal{L}^2(\mathbf{P})$ are pairwise independent with bounded variances, then $(X_n)_{n \in \mathbb{N}}$ fulfills the strong law of large numbers.

Exercise 5.3.2 Let $(X_n)_{n \in \mathbb{N}}$ be a sequence of independent identically distributed random variables with $\frac{1}{n}(X_1 + \dots + X_n) \xrightarrow{n \rightarrow \infty} Y$ almost surely for some random variable Y . Show that $X_1 \in \mathcal{L}^1(\mathbf{P})$ and $Y = \mathbf{E}[X_1]$ almost surely.

Hint: First show that

$$\mathbf{P}[|X_n| > n \text{ for infinitely many } n] = 0 \iff X_1 \in \mathcal{L}^1(\mathbf{P}).$$

Exercise 5.3.3 Let E be a finite set and let p be a probability vector on E . Show that the entropy $H(p)$ is minimal (in fact, zero) if $p = \delta_e$ for some $e \in E$. It is maximal (in fact, $\log(\#E)$) if p is the uniform distribution on E .

Exercise 5.3.4 (Subadditivity of Entropy) For $i = 1, 2$, let E^i be a finite set and p^i a probability vector on E^i . Let p be a probability vector on $E^1 \times E^2$ with marginals p^1 and p^2 . That is,

$$\sum_{e^2 \in E^2} p(e^1, e^2) = p_{e^1}^1 \quad \text{and} \quad \sum_{f^1 \in E^1} p(f^1, f^2) = p_{f^2}^2 \quad \text{for all } e^1 \in E^1, f^2 \in E^2.$$

Show that $H(p) \leq H(p^1) + H(p^2)$.

Exercise 5.3.5 Let $b \in \{2, 3, 4, \dots\}$. A b -adic prefix code is defined in a similar way as a binary prefix code; however, instead of 0 and 1, now all numbers $0, 1, \dots, b-1$ are admissible. Show that the statement of the source coding theorem holds for b -adic prefix codes with $H_2(p)$ replaced by $H_b(p)$.

Exercise 5.3.6 We want to check the efficiency of the Morse alphabet. To this end we need a table of the Morse code as well as the frequencies of the letters in a typical text. The following frequencies for letters in German texts are taken from [11, p. 10]. The frequencies for other languages can be found easily, e.g., at Wikipedia.

Letter	Morse code	Frequency
A	.-	0.0651
B	-...	0.0189
C	..-	0.0306
D	..-	0.0508
E	.	0.1740
F	...-	0.0166
G	.-	0.0301
H		0.0476
I	..	0.0755
J	.-	0.0027
K	.-	0.0121
L	...-	0.0344
M	-	0.0253

Letter	Morse code	Frequency
N	.-	0.0978
O	—	0.0251
P	..-	0.0079
Q	.-	0.0002
R	..-	0.07
S	...	0.0727
T	-	0.0615
U	...-	0.0435
V	...-	0.0067
W	.-	0.0189
X	..-	0.0003
Y	.-	0.0004
Z	..-	0.0113

Here '.' denotes a short signal while '-' denotes a long signal. Each letter is finished by a pause sign. Thus the Morse code can be interpreted as a ternary prefix code.

Determine the average code length of a letter and compare it with the entropy H_3 in order to check the efficiency of the Morse code.

5.5.4 5.4 Speed of Convergence in the Strong Lln

In the weak law of large numbers, we had a statement on the speed of convergence (Theorem 5.14). In the strong law of large numbers, however, we did not. As we required only first moments, in general, we cannot expect to get any useful statements. However, if we assume the existence of higher moments, we get reasonable estimates on the rate of convergence.

The core of the weak law of large numbers is Chebyshev's inequality. Here we present a stronger inequality that claims the same bound but now for the maximum over all partial sums until a fixed time.

Theorem 5.28 (Kolmogorov's inequality) Let $n \in \mathbb{N}$ and let $X_1, X_2, \dots, X_n$ be independent random variables with $\mathbf{E}[X_i] = 0$ and $\text{Var}[X_i] < \infty$ for $i = 1, \dots, n$. Further, let $S_k = X_1 + \dots + X_k$ for $k = 1, \dots, n$. Then, for any $t > 0$,

$$\mathbf{P}[\max\{S_k : k = 1, \dots, n\} \geq t] \leq \frac{\text{Var}[S_n]}{t^2 + \text{Var}[S_n]} \quad (5.10)$$

Furthermore, Kolmogorov's inequality holds:

$$\mathbf{P}[\max\{|S_k| : k = 1, \dots, n\} \geq t] \leq t^{-2} \text{Var}[S_n] \quad (5.11)$$

In Theorem 11.2 we will see Doob's inequality, which is a generalization of Kolmogorov's inequality.

Proof We decompose the probability space according to the first time τ at which the partial sums exceed the value t . Hence, let

$$\tau := \min\{k \in \{1, \dots, n\} : S_k \geq t\}$$

and $A_k = \{\tau = k\}$ for $k = 1, \dots, n$. Further, let

$$A = \bigcup_{k=1}^n A_k = \{\max\{S_k : k = 1, \dots, n\} \geq t\}$$

Let $c \geq 0$. The random variable $(S_k + c) \mathbb{1}_{A_k}$ is $\sigma(X_1, \dots, X_k)$ -measurable and $S_n - S_k$ is $\sigma(X_{k+1}, \dots, X_n)$ -measurable. By Theorem 2.26, the two random variables are independent, and

$$\mathbf{E}[(S_k + c) \mathbb{1}_{A_k} (S_n - S_k)] = \mathbf{E}[(S_k + c) \mathbb{1}_{A_k}] \mathbf{E}[S_n - S_k] = 0$$

Clearly, the events $A_1, \dots, A_n$ are pairwise disjoint; hence $\sum_{k=1}^n \mathbb{1}_{A_k} = \mathbb{1}_A \leq 1$. We thus obtain

$$\begin{aligned} \text{Var}[S_n] + c^2 &= \mathbf{E}[(S_n + c)^2] \\ &\geq \mathbf{E}\left[\sum_{k=1}^n (S_n + c)^2 \mathbb{1}_{A_k}\right] = \sum_{k=1}^n \mathbf{E}[(S_n + c)^2 \mathbb{1}_{A_k}] \\ &= \sum_{k=1}^n \mathbf{E}[(S_k + c)^2 + 2(S_k + c)(S_n - S_k) + (S_n - S_k)^2] \mathbb{1}_{A_k} \\ &= \sum_{k=1}^n \mathbf{E}[(S_k + c)^2 \mathbb{1}_{A_k}] + \sum_{k=1}^n \mathbf{E}[(S_n - S_k)^2 \mathbb{1}_{A_k}] \\ &\geq \sum_{k=1}^n \mathbf{E}[(S_k + c)^2 \mathbb{1}_{A_k}]. \end{aligned} \quad (5.12)$$

Since $c \geq 0$, we have $(S_k + c)^2 \mathbb{1}_{A_k} \geq (t + c)^2 \mathbb{1}_{A_k}$. Hence we can continue (5.12) to get

$$\text{Var}[S_n] + c^2 \geq \sum_{k=1}^n \mathbf{E}[(t + c)^2 \mathbb{1}_{A_k}] = (t + c)^2 \mathbf{P}[A]$$

For $c = \text{Var}[S_n]/t \geq 0$, we obtain

$$\mathbf{P}[A] \leq \frac{\text{Var}[S_n] + c^2}{(t + c)^2} = \frac{c(t + c)}{(t + c)^2} = \frac{tc}{t^2 + tc} = \frac{\text{Var}[S_n]}{t^2 + \text{Var}[S_n]}.$$

This shows (5.10). In order to show (5.11), choose

$$\bar{\tau} := \min\{k \in \{1, \dots, n\} : |S_k| \geq t\}.$$

Let $\bar{A}_k = \{\bar{\tau} = k\}$ and $\bar{A} = \{\bar{\tau} \leq n\}$. We cannot now continue (5.12) as above with $c > 0$. However, if we choose $c = 0$, then $S_k^2 \mathbb{1}_{\bar{A}_k} \geq t^2 \mathbb{1}_{\bar{A}_k}$. The same calculation as in (5.12) does then yield $\mathbf{P}[\bar{A}] \leq t^{-2} \text{Var}[S_n]$.

From Kolmogorov's inequality, we derive the following sharpening of the strong law of large numbers.

Theorem 5.29 Let $X_1, X_2, \dots$ be independent random variables with $\mathbf{E}[X_n] = 0$ for any $n \in \mathbb{N}$ and $V := \sup \{\text{Var}[X_n] : n \in \mathbb{N}\} < \infty$. Then, for any $\varepsilon > 0$,

$$\limsup_{n \rightarrow \infty} \frac{|S_n|}{n^{1/2}(\log(n))^{(1/2)+\varepsilon}} = 0 \quad \text{almost surely.}$$

Proof Let $k_n = 2^n$ and $l(n) = n^{1/2}(\log(n))^{(1/2)+\varepsilon}$ for $n \in \mathbb{N}$. Then we have $l(k_{n+1})/l(k_n) \xrightarrow{n \rightarrow \infty} \sqrt{2}$. Hence, for $n \in \mathbb{N}$ sufficiently large and $k \in \mathbb{N}$ with $k_{n-1} \leq k \leq k_n$, we have $|S_k|/l(k) \leq 2|S_{k_n}|/l(k_n)$. Hence, it is enough to show for every $\delta > 0$ that

$$\limsup_{n \rightarrow \infty} l(k_n)^{-1} \max \{|S_k| : k \leq k_n\} \leq \delta \quad \text{almost surely.} \quad (5.13)$$

For $\delta > 0$ and $n \in \mathbb{N}$, define $A_n^\delta := \{\max \{|S_k| : k \leq k_n\} > \delta l(k_n)\}$. Kolmogorov's inequality yields

$$\sum_{n=1}^{\infty} \mathbf{P}[A_n^\delta] \leq \sum_{n=1}^{\infty} \delta^{-2} (l(k_n))^{-2} V k_n = \frac{V}{\delta^2 (\log 2)^{1+2\varepsilon}} \sum_{n=1}^{\infty} n^{-1-2\varepsilon} < \infty.$$

The Borel-Cantelli lemma then gives $\mathbf{P}[\limsup_{n \rightarrow \infty} A_n^\delta] = 0$ and hence (5.13).

In Chapter 22, we will see that for independent identically distributed, square integrable, centered random variables $X_1, X_2, \dots$, the following strengthening holds,

$$\limsup_{n \rightarrow \infty} \frac{|S_n|}{\sqrt{2n \text{Var}[X_1] \log(\log(n))}} = 1 \quad \text{almost surely.}$$

Hence, in this case, the speed of convergence is known precisely. If the $X_1, X_2, \dots$ are not independent but only pairwise independent, then the rate of convergence deteriorates, although not drastically. Here we cite without proof a theorem that was found independently by Rademacher (1922) [141] and Menshov (1923) [113].

Theorem 5.30 (Rademacher-Menshov) Let $X_1, X_2, \dots$ be uncorrelated, square integrable, centered random variables and let $(a_n)_{n \in \mathbb{N}}$ be an increasing sequence of nonnegative numbers such that

$$\sum_{n=1}^{\infty} (\log n)^2 a_n^{-2} \text{Var}[X_n] < \infty \quad (5.14)$$

Then

$$\limsup_{n \rightarrow \infty} \left| a_n^{-1} \sum_{k=1}^n X_k \right| = 0 \quad \text{almost surely.}$$

Proof See, for example, [128].

Remark 5.31 Condition (5.14) is sharp in the sense that for any increasing sequence $(a_n)_{n \in \mathbb{N}}$ with $\sum_{n=1}^{\infty} a_n^{-2} (\log n)^2 = \infty$, there exists a sequence of pairwise independent, square integrable, centered random variables $X_1, X_2, \dots$ with $\text{Var}[X_n] = 1$ for all $n \in \mathbb{N}$ such that

$$\limsup_{n \rightarrow \infty} \left| a_n^{-1} \sum_{k=1}^n X_k \right| = \infty \quad \text{almost surely}$$

See [22]. There an example of [163] (see also [164, 165]) for orthogonal series is developed further. See also [117].

For random variables with infinite variance, the statements about the rate of convergence naturally get weaker. For example (see [8]), see the following theorem.

Theorem 5.32 (Baum and Katz (1965)) Let $\gamma > 1$ and let $X_1, X_2, \dots$ be i.i.d. Define $S_n = X_1 + \dots + X_n$ for $n \in \mathbb{N}$. Then

$$\sum_{n=1}^{\infty} n^{\gamma-2} \mathbf{P}[|S_n|/n > \varepsilon] < \infty \quad \text{for any } \varepsilon > 0$$

$$\iff \mathbf{E}[|X_1|^\gamma] < \infty \quad \text{and} \quad \mathbf{E}[X_1] = 0$$

Exercise 5.4.1 Let $X_1, \dots, X_n$ be independent real random variables and let $S_k = X_1 + \dots + X_k$ for $k = 1, \dots, n$. Show that for $t > 0$ Etemadi's inequality holds:

$$\mathbf{P}\left[\max_{k=1, \dots, n} |S_k| \geq t\right] \leq 3 \max_{k=1, \dots, n} \mathbf{P}[|S_k| \geq t/3].$$

5.5.5 5.5 The Poisson Process

We develop a model for the number of clicks of a Geiger counter in the (time) interval $I = (a, b]$. The number of clicks should obey the following rules. It should

- be random and independent for disjoint intervals,
- be homogeneous in time in the sense that the number of clicks in $I = (a, b]$ has the same distribution as the number of clicks in $c + I = (a + c, b + c]$,
- have finite expectation, and
- have no double points: At any point of time, the counter makes at most one click.

We formalize these requirements by introducing the following notation:

$$\mathcal{I} := \{(a, b] : a, b \in [0, \infty), a \leq b\}$$

$$\ell((a, b]) := b - a \quad (\text{the length of the interval } I = (a, b])$$

For $I \in \mathcal{I}$, let N_I be the number of clicks after time a but no later than b . In particular, we define $N_t := N_{(0, t]}$ as the total number of clicks until time t . The above requirements translate to: $(N_I, I \in \mathcal{I})$ being a family of random variables with values in $\mathbb{N}_0$ and with the following properties:

(P1) $N_{I \cup J} = N_I + N_J$ if $I \cap J = \emptyset$ and $I \cup J \in \mathcal{I}$.

(P2) The distribution of N_I depends only on the length of I : $\mathbf{P}_{N_I} = \mathbf{P}_{N_J}$ for all $I, J \in \mathcal{I}$ with $\ell(I) = \ell(J)$.

(P3) If $\mathcal{J} \subset \mathcal{I}$ with $I \cap J = \emptyset$ for all $I, J \in \mathcal{J}$ with $I \neq J$, then $(N_J, J \in \mathcal{J})$ is an independent family.

(P4) For any $I \in \mathcal{I}$, we have $\mathbf{E}[N_I] < \infty$.

(P5) $\limsup \varepsilon^{-1} \mathbf{P}[N_\varepsilon \geq 2] = 0$.

The meaning of (P5) is explained by the following calculation. Define


$$\lambda := \limsup_{\varepsilon \downarrow 0} \varepsilon^{-1} \mathbf{P}[N_\varepsilon \geq 2]$$

For any $n \in \mathbb{N}$ and $\varepsilon > 0$, we have

$$\mathbf{P}[N_{2^{-n}} \geq 2] \geq \lfloor 2^{-n}/\varepsilon \rfloor \mathbf{P}[N_\varepsilon \geq 2] - \lfloor 2^{-n}/\varepsilon \rfloor^2 \mathbf{P}[N_\varepsilon \geq 2]^2.$$

Hence

$$2^n \mathbf{P}[N_{2^{-n}} \geq 2] \geq \lambda - 2^{-n} \lambda^2 \xrightarrow{n \rightarrow \infty} \lambda.$$

Fig. 5.1 Simulation of a Poisson process with rate $\alpha = 0.5$

Then (because $(1 - a_k/k)^k \xrightarrow{k \rightarrow \infty} e^{-a}$ if $a_k \xrightarrow{k \rightarrow \infty} a$)
 $\mathbf{P}[\text{there is a double click in } (0, 1]]$

$$\begin{aligned} &= \lim_{n \rightarrow \infty} \mathbf{P} \left[\bigcup_{k=0}^{2^n-1} \{N_{(k2^{-n}, (k+1)2^{-n}]} \geq 2\} \right] \\ &= 1 - \lim_{n \rightarrow \infty} \mathbf{P} \left[\bigcap_{k=0}^{2^n-1} \{N_{(k2^{-n}, (k+1)2^{-n}]} \leq 1\} \right] \\ &= 1 - \lim_{n \rightarrow \infty} \prod_{k=0}^{2^n-1} \mathbf{P}[N_{(k2^{-n}, (k+1)2^{-n}]} \leq 1] \\ &= 1 - \lim_{n \rightarrow \infty} (1 - \mathbf{P}[N_{2^{-n}} \geq 2])^{2^n} \\ &= 1 - e^{-\lambda} \end{aligned}$$

Hence we have to postulate $\lambda = 0$. This, however, is exactly (P5).

The following theorem shows that properties (P1)-(P5) characterize the random variables $(N_I, I \in \mathcal{I})$ uniquely and that they form a Poisson process.

Definition 5.33 (Poisson process) A family $(N_t, t \geq 0)$ of $\mathbb{N}_0$ -valued random variables is called a Poisson process with intensity $\alpha \geq 0$ if $N_0 = 0$ and if:

- (i) For any $n \in \mathbb{N}$ and any choice of $n + 1$ numbers $0 = t_0 < t_1 < \dots < t_n$, the family $(N_{t_i} - N_{t_{i-1}}, i = 1, \dots, n)$ is independent.
- (ii) For $t > s \geq 0$, the difference $N_t - N_s$ is Poisson-distributed with parameter $\alpha(t - s)$; that is,

$$\mathbf{P}[N_t - N_s = k] = e^{-\alpha(t-s)} \frac{(\alpha(t-s))^k}{k!} \quad \text{for all } k \in \mathbb{N}_0.$$

See Fig. 5.1 for a computer simulation of a Poisson process.

The existence of the Poisson process has not yet been shown. We come back to this point in Theorem 5.36.

Theorem 5.34 If $(N_I, I \in \mathcal{I})$ has properties (P1)-(P5), then $(N_{(0,t]}, t \geq 0)$ is a Poisson process with intensity $\alpha := \mathbf{E}[N_{(0,1]}]$. If, on the other hand, $(N_t, t \geq 0)$ is a Poisson process, then $(N_t - N_s, (s, t] \in \mathcal{I})$ has properties (P1)-(P5).

Proof First assume that $(N_t, t \geq 0)$ is a Poisson process with intensity $\alpha \geq 0$. Then, for $I = (a, b]$, clearly $\mathbf{P}_{N_I} = \text{Poi}_{\alpha(b-a)} = \text{Poi}_{\alpha \ell(I)}$. Hence (P2) holds. By (i), we have (P3). Clearly, $\mathbf{E}[N_I] = \alpha \ell(I) < \infty$; thus we have (P4). Finally, $\mathbf{P}[N_\varepsilon \geq 2] = 1 - e^{-\alpha\varepsilon} - \alpha\varepsilon e^{-\alpha\varepsilon} = f(0) - f(\alpha\varepsilon)$, where $f(x) := e^{-x} + xe^{-x}$. The derivative is $f'(x) = -xe^{-x}$, whence

$$\lim_{\varepsilon \downarrow 0} \varepsilon^{-1} \mathbf{P} [N_\varepsilon \geq 2] = -\alpha f'(0) = 0.$$

This implies (P5).

Now assume that $(N_I, I \in \mathcal{I})$ fulfills (P1)-(P5). Define $\alpha(t) := \mathbf{E} [N_t]$. Then (owing to (P2))

$$\alpha(s+t) = \mathbf{E} [N_{(0,s]} + N_{(s,s+t]}] = \mathbf{E} [N_{(0,s]}] + \mathbf{E} [N_{(0,t]}] = \alpha(s) + \alpha(t).$$

As $t \mapsto \alpha(t)$ is monotone increasing, this implies linearity: $\alpha(t) = t\alpha(1)$ for any $t \geq 0$. Letting $\alpha := \alpha(1)$, we obtain $\mathbf{E} [N_I] = \alpha \ell(I)$. It remains to show that $\mathbf{P}_{N_t} = \text{Poi}_{\alpha t}$. In order to apply the Poisson approximation theorem (Theorem 3.7), for fixed $n \in \mathbb{N}$, we decompose the interval $(0, t]$ into 2^n disjoint intervals of equal length,

$$I^n(k) := ((k-1)2^{-n}t, k2^{-n}t], \quad k = 1, \dots, 2^n.$$

Now define $X^n(k) := N_{I^n(k)}$ and

$$\bar{X}^n(k) := \begin{cases} 1, & \text{if } X^n(k) \geq 1 \\ 0, & \text{else} \end{cases}$$

By properties (P2) and (P3), the random variables $(X^n(k), k = 1, \dots, 2^n)$ are independent and identically distributed. Hence also $(\bar{X}^n(k), k = 1, \dots, 2^n)$ are i.i.d., namely $\bar{X}^n(k) \sim \text{Ber}_{p_n}$, where $p_n = \mathbf{P} [N_{2^{-n}t} \geq 1]$.

Finally, let $N_t^n := \sum_{k=1}^{2^n} \bar{X}^n(k)$. Then $N_t^n \sim b_{2^n, p_n}$. Clearly, $N_t^{n+1} - N_t^n \geq 0$. Now, by (P5),

$$\mathbf{P} [N_t \neq N_t^n] \leq \sum_{k=1}^{2^n} \mathbf{P} [X^n(k) \geq 2] = 2^n \mathbf{P} [N_{2^{-n}t} \geq 2] \xrightarrow{n \rightarrow \infty} 0. \quad (5.15)$$

Hence $\mathbf{P} [N_t = \lim_{n \rightarrow \infty} N_t^n] = 1$. By the monotone convergence theorem, we get

$$\alpha t = \mathbf{E} [N_t] = \lim_{n \rightarrow \infty} \mathbf{E} [N_t^n] = \lim_{n \rightarrow \infty} p_n 2^n.$$

Using the Poisson approximation theorem (Theorem 3.7), we infer that, for any $l \in \mathbb{N}_0$,

$$\mathbf{P} [N_t = l] = \lim_{n \rightarrow \infty} \mathbf{P} [N_t^n = l] = \text{Poi}_{\alpha t}(\{l\}).$$

Hence $\mathbf{P}_{N_t} = \text{Poi}_{\alpha t}$.

At this point, we still have to show that there are Poisson processes at all. We present a general two-step construction principle that will be used in a similar form later in Chapter 24 in a more general setting. In the first step, we determine the (random) number of jumps in $(0, 1]$. In the second step, we distribute these jumps uniformly and independently on $(0, 1]$. Strictly speaking, this gives the Poisson process only on the time interval $(0, 1]$, but it is clear how to move on: We perform the same procedure independently for each of the intervals $(1, 2]$, $(2, 3]$ and so on and then collect the jumps (see also Exercise 5.5.1).

Let $\alpha > 0$ and let L be a Poi_α random variable. Further, let $X_1, X_2, \dots$ be independent random variables, that are uniformly distributed on $(0, 1]$, i.e., $X_k \sim \mathcal{U}_{(0,1]}$ for each k . We assume that $\{L, X_1, X_2, \dots\}$ is an independent family of random variables. We now define $N = (N_t)_{t \in [0,1]}$ by

$$N_t := \sum_{l=1}^L \mathbb{1}_{(0,t]}(X_l) \quad \text{for } t \in [0, 1]. \quad (5.16)$$

Theorem 5.35 The family N of random variables defined in (5.16) is a Poisson process with intensity α (and time set $[0, 1]$).

Proof We have to show that the increments of N in finitely many pairwise disjoint intervals are independent and Poisson distributed. Hence let $m \in \mathbb{N}$ and $0 = t_0 < t_1 < \dots < t_m = 1$. We use the abbreviations $p_i := t_i - t_{i-1}$ and $\lambda_i = \alpha \cdot (t_i - t_{i-1})$ and show that

$$(N_{t_i} - N_{t_{i-1}})_{i=1, \dots, m} \quad \text{is independent} \quad (5.17)$$

and

$$N_{t_i} - N_{t_{i-1}} \sim \text{Poi}_{\lambda_i} \quad \text{for all } i = 1, \dots, m. \quad (5.18)$$

This is equivalent to showing that for each choice of $k_1, \dots, k_m \in \mathbb{N}_0$, we have

$$\mathbf{P} [N_{t_i} - N_{t_{i-1}} = k_i \quad \text{for any } i = 1, \dots, m] = \prod_{i=1}^m \left(e^{-\lambda_i} \frac{\lambda_i^{k_i}}{k_i!} \right) \quad (5.19)$$

Write

$$M_{n,i} := \# \{l \leq n : t_{i-1} < X_l \leq t_i\} = \sum_{l=1}^n \mathbb{1}_{(t_{i-1}, t_i]}(X_l).$$

By Exercise 2.2.3, the vector $(M_{n,1}, \dots, M_{n,m})$ is multinomially distributed with parameters n and $p = (p_1, \dots, p_m)$. That is, if we assume $n := k_1 + \dots + k_m$, then

$$\mathbf{P} [M_{n,1} = k_1, \dots, M_{n,m} = k_m] = \frac{n!}{k_1! \dots k_m!} p_1^{k_1} \dots p_m^{k_m}$$

In order to show (5.19), note that the event in (5.19) implies $L = n$ and that L and $(M_{n,1}, \dots, M_{n,m})$ are independent. Hence we have

$$\begin{aligned} & \mathbf{P} [N_{t_i} - N_{t_{i-1}} = k_i \text{ for } i = 1, \dots, m] \\ &= \mathbf{P} [\{N_{t_i} - N_{t_{i-1}} = k_i \text{ for } i = 1, \dots, m\} \cap \{L = n\}] \\ &= \mathbf{P} [\{M_{n,1} = k_1, \dots, M_{n,m} = k_m\} \cap \{L = n\}] \\ &= \mathbf{P} [M_{n,1} = k_1, \dots, M_{n,m} = k_m] \cdot \mathbf{P} [L = n] \\ &= \frac{n!}{k_1! \dots k_m!} p_1^{k_1} \dots p_m^{k_m} e^{-\alpha} \frac{\alpha^n}{n!} \\ &= \prod_{i=1}^m \left(e^{-\lambda_i} \frac{\lambda_i^{k_i}}{k_i!} \right). \end{aligned}$$

We close this section by presenting a further, rather elementary and instructive construction of the Poisson process based on specifying the waiting times between the clicks of the Geiger counter, or, more formally, between the points of discontinuity of the map $t \mapsto N_t(\omega)$. At time s , what is the probability that we have to wait another t time units (or longer) for the next click? Since we modeled the clicks as a Poisson process with intensity α , this probability can easily be computed:

$$\mathbf{P} [N_{(s, s+t]} = 0] = e^{-\alpha t}$$

Hence the waiting time for the next click is exponentially distributed with parameter α . Furthermore, the waiting times should be independent. We now take the waiting times as the starting point and, based on them, construct the Poisson process.

Let $W_1, W_2, \dots$ be an independent family of exponentially distributed random variables with parameter $\alpha > 0$; hence $\mathbf{P}[W_n > x] = e^{-\alpha x}$. We define

$$T_n := \sum_{k=1}^n W_k$$

and interpret W_n as the waiting time between the $(n-1)$ th click and the n th click. T_n is the time of the n th click. Appealing to this intuition we define the number of clicks until time t by

$$N_t := \# \{n \in \mathbb{N}_0 : T_n \leq t\}$$

Hence

$$\{N_t = k\} = \{T_k \leq t < T_{k+1}\}.$$

In particular, N_t is a random variable; that is, measurable.

Theorem 5.36 The family $(N_t, t \geq 0)$ is a Poisson process with intensity α .

Proof (We follow the proof in [59, Theorem 3.34].) We must show that for any $n \in \mathbb{N}$ and any sequence $0 = t_0 < t_1 < \dots < t_n$, we have that $(N_{t_i} - N_{t_{i-1}}, i = 1, \dots, n)$ is independent and $N_{t_i} - N_{t_{i-1}} \sim \text{Poi}_{\alpha(t_i - t_{i-1})}$. We are well aware that it is not enough to show this for the case $n = 2$ only. However, the notational complications become overwhelming for $n \geq 3$, and the idea for general $n \in \mathbb{N}$ becomes clear in the case $n = 2$. Hence we restrict ourselves to the case $n = 2$.

Hence we show for $0 < s < t$ and $l, k \in \mathbb{N}_0$ that

$$\mathbf{P}[N_s = k, N_t - N_s = l] = \left(e^{-\alpha s} \frac{(\alpha s)^k}{k!} \right) \left(e^{-\alpha(t-s)} \frac{(\alpha(t-s))^l}{l!} \right). \quad (5.20)$$

This implies that N_s and $(N_t - N_s)$ are independent. Furthermore, by summing over $k \in \mathbb{N}_0$, this yields $N_t - N_s \sim \text{Poi}_{\alpha(t-s)}$.

By Corollary 2.22, the distribution $\mathbf{P}_{(W_1, \dots, W_{k+l+1})}$ has the density

$$x \mapsto \alpha^{k+l+1} e^{-\alpha S_{k+l+1}(x)},$$

where $S_n(x) := x_1 + \dots + x_n$. It is sufficient to consider $l \geq 1$ since we get the $l = 0$ term from the fact that the probability measure has total mass one. Hence, let $l \geq 1$. We compute

$$\begin{aligned} & \mathbf{P}[N_s = k, N_t - N_s = l] \\ &= \mathbf{P}[T_k \leq s < T_{k+1}, T_{k+l} \leq t < T_{k+l+1}] \\ &= \int_0^\infty \dots \int_0^\infty dx_1 \dots dx_{k+l+1} \\ &\quad \times \alpha^{k+l+1} e^{-\alpha S_{k+l+1}(x)} \mathbb{1}_{\{S_k(x) \leq s < S_{k+1}(x)\}} \mathbb{1}_{\{S_{k+l}(x) \leq t < S_{k+l+1}(x)\}} \end{aligned}$$

Starting with x_{k+l+1} , we integrate successively. In the first step, substitute $z = S_{k+l+1}(x)$ to obtain

$$\int_0^\infty dx_{k+l+1} \alpha e^{-\alpha S_{k+l+1}(x)} \mathbb{1}_{\{S_{k+l+1}(x) > t\}} = \int_t^\infty dz \alpha e^{-\alpha z} = e^{-\alpha t}$$

Now keep $x_1, \dots, x_k$ fixed and substitute for the remaining variables by letting $y_1 = S_{k+1}(x) - s, y_2 = x_{k+2}, \dots, y_l = x_{k+l}$ to obtain

$$\begin{aligned} & \int_0^\infty \dots \int_0^\infty dx_{k+1} \dots dx_{k+l} \mathbb{1}_{\{s < S_{k+1}(x) \leq S_{k+l}(x) \leq t\}} \\ &= \int_0^\infty \dots \int_0^\infty dy_1 \dots dy_l \mathbb{1}_{\{y_1 + \dots + y_l \leq t-s\}} = \frac{(t-s)^l}{l!} \end{aligned}$$

(The last identity can be obtained, for example, by induction on l .) Now integrate the remaining variables $x_1, \dots, x_k$ to get

$$\int_0^\infty \dots \int_0^\infty dx_1 \dots dx_k \mathbb{1}_{\{S_k(x) \leq s\}} = \frac{s^k}{k!}$$

In total, we have

$$\mathbf{P}[N_s = k, N_t - N_s = l] = e^{-\alpha t} \alpha^{k+l} \frac{s^k}{k!} \frac{(t-s)^l}{l!}$$

hence (5.20) holds.

Exercise 5.5.1 Let $L_n, X_k^n, k, n \in \mathbb{N}$ be independent random variables with $L_n \sim \text{Poi}_\alpha$ and $X_k^n \sim \mathcal{U}_{(n-1, n]}$ (the uniform distribution on $(n-1, n]$) for all $k, n \in \mathbb{N}$. Define

$$N_t := \# \{(k, n) \in \mathbb{N}^2 : k \leq L_n \text{ and } X_k^n \leq t\}.$$

Show that $(N_t)_{t \geq 0}$ is a Poisson process with intensity α .

Exercise 5.5.2 Let $T > 0$ and let $X_1, X_2, \dots$ be i.i.d. random variables that are uniformly distributed on $[0, 1]$. Let

$$N := \max \{n \in \mathbb{N}_0 : X_1 + \dots + X_n \leq T\}$$

and compute $\mathbf{E}[N]$.

5.6 6 Convergence Theorems

In the strong and the weak laws of large numbers, we implicitly introduced the notions of almost sure convergence and convergence in probability of random variables. We saw that almost sure convergence implies convergence in measure/probability. This chapter is devoted to a systematic treatment of almost sure convergence, convergence in measure and convergence of integrals. The key role for connecting convergence in measure and convergence of integrals is played by the concept of uniform integrability.

5.6.1 6.1 Almost Sure and Measure Convergence

In the following, $(\Omega, \mathcal{A}, \mu)$ will be a σ -finite measure space. We first define in metric spaces almost sure convergence and convergence in measure and then compare both concepts. To this end, we need two lemmas that ensure that the distance function associated with two measurable maps is again measurable. In the following, let (E, d) be a separable metric space with Borel σ -algebra $\mathcal{B}(E)$. "Separable" means that there exists a countable dense subset. For $x \in E$ and $r > 0$, denote by $B_r(x) = \{y \in E : d(x, y) < r\}$ the ball with radius r centered at x .

Lemma 6.1 Let $f, g : \Omega \rightarrow E$ be measurable with respect to $\mathcal{A} - \mathcal{B}(E)$. Then the map $H : \Omega \rightarrow [0, \infty), \omega \mapsto d(f(\omega), g(\omega))$ is $\mathcal{A} - \mathcal{B}([0, \infty))$ -measurable.

Proof Let $F \subset E$ be countable and dense. By the triangle inequality, $d(x, z) + d(z, y) \geq d(x, y)$ for all $x, y \in E$ and $z \in F$. Let $(z_n)_{n \in \mathbb{N}}$ be a sequence in F with $z_n \xrightarrow{n \rightarrow \infty} x$. Since d is continuous, we have $d(x, z_n) + d(z_n, y) \xrightarrow{n \rightarrow \infty} d(x, y)$. Putting things together, we infer $\inf_{z \in F} (d(x, z) + d(z, y)) = d(x, y)$. Since $x \mapsto d(x, z)$ is continuous and hence measurable, the

maps $f_z, g_z : \Omega \rightarrow [0, \infty)$ with $f_z(\omega) = d(f(\omega), z)$ and $g_z(\omega) = d(g(\omega), z)$ are also measurable. Thus $f_z + g_z$ and $H = \inf_{z \in F} (f_z + g_z)$ are measurable.

(A somewhat more systematic proof is based on the fact that (f, g) is $\mathcal{A} - \mathcal{B}(E \times E)$ -measurable (this will follow from Theorem 14.8) and that $d : E \times E \rightarrow [0, \infty)$ is continuous and hence $\mathcal{B}(E \times E) - \mathcal{B}([0, \infty))$ -measurable. As a composition of measurable maps, $\omega \mapsto d(f(\omega), g(\omega))$ is measurable.)

Let $f, f_1, f_2, \dots : \Omega \rightarrow E$ be measurable with respect to $\mathcal{A} - \mathcal{B}(E)$.

Definition 6.2 We say that $(f_n)_{n \in \mathbb{N}}$ converges to f

(i) in μ -measure (or, briefly, in measure), symbolically $f_n \xrightarrow{\text{meas}} f$, if

$$\mu(\{d(f, f_n) > \varepsilon\} \cap A) \xrightarrow{n \rightarrow \infty} 0$$

for all $\varepsilon > 0$ and all $A \in \mathcal{A}$ with $\mu(A) < \infty$, and

(ii) μ -almost everywhere (a.e.), symbolically $f_n \xrightarrow{\text{a.e.}} f$, if there exists a μ -null set $N \in \mathcal{A}$ such that

$$d(f(\omega), f_n(\omega)) \xrightarrow{n \rightarrow \infty} 0 \quad \text{for any } \omega \in \Omega \setminus N.$$

If μ is a probability measure, then convergence in μ -measure is also called convergence in probability. If $(f_n)_{n \in \mathbb{N}}$ converges a.e., then we also say that $(f_n)_{n \in \mathbb{N}}$ converges almost surely (a.s.) and write $f_n \xrightarrow{\text{a.s.}} f$. Sometimes we will drop the qualifications "almost everywhere" and "almost surely".

Remark 6.3 Let $A_1, A_2, \dots \in \mathcal{A}$ with $A_n \uparrow \Omega$ and $\mu(A_n) < \infty$ for any $n \in \mathbb{N}$. Then a.e. convergence is equivalent to a.e. convergence on each A_n .

Remark 6.4 Almost everywhere convergence implies convergence in measure: For $\varepsilon > 0$, define

$$D_n(\varepsilon) = \{d(f, f_m) > \varepsilon \text{ for some } m \geq n\}$$

Then $D(\varepsilon) := \bigcap_{n=1}^{\infty} D_n(\varepsilon) \subset N$, where N is the null set from the definition of almost everywhere convergence. Upper semicontinuity of μ implies

$$\mu(D_n(\varepsilon) \cap A) \xrightarrow{n \rightarrow \infty} \mu(D(\varepsilon) \cap A) = 0$$

for any $A \in \mathcal{A}$ with $\mu(A) < \infty$.

Remark 6.5 Almost everywhere convergence and convergence in measure determine the limit up to equality almost everywhere. Indeed, let $f_n \xrightarrow{\text{meas}} f$ and $f_n \xrightarrow{\text{meas}} g$. Let $A_1, A_2, \dots \in \mathcal{A}$ with $A_n \uparrow \Omega$ and $\mu(A_n) < \infty$ for any $n \in \mathbb{N}$. Then (since $d(f, g) \leq d(f, f_n) + d(g, f_n)$), for any $m \in \mathbb{N}$ and $\varepsilon > 0$,

$$\begin{aligned} \mu(A_m \cap \{d(f, g) > \varepsilon\}) \\ \leq \mu(A_m \cap \{d(f, f_n) > \varepsilon/2\}) + \mu(A_m \cap \{d(g, f_n) > \varepsilon/2\}) \xrightarrow{n \rightarrow \infty} 0 \end{aligned}$$

Hence $\mu(\{d(f, g) > 0\}) = 0$.

Remark 6.6 In general, convergence in measure does not imply almost everywhere convergence. Indeed, let $(X_n)_{n \in \mathbb{N}}$ be an independent family of random variables with $X_n \sim \text{Ber}_{1/n}$. Then $X_n \xrightarrow{n \rightarrow \infty} 0$ in probability but the Borel-Cantelli lemma implies $\limsup_{n \rightarrow \infty} X_n = 1$ almost surely.

Theorem 6.7 Let $A_1, A_2, \dots \in \mathcal{A}$ with $A_N \uparrow \Omega$ and $\mu(A_N) < \infty$ for all $N \in \mathbb{N}$. For measurable $f, g : \Omega \rightarrow E$, let

$$\tilde{d}(f, g) := \sum_{N=1}^{\infty} \frac{2^{-N}}{1 + \mu(A_N)} \int_{A_N} (1 \wedge d(f(\omega), g(\omega))) \mu(d\omega) \quad (6.1)$$

Then $\tilde{d}$ is a metric that induces convergence in measure: If $f, f_1, f_2, \dots$ are measurable, then

$$f_n \xrightarrow{\text{meas}} f \iff \tilde{d}(f, f_n) \xrightarrow{n \rightarrow \infty} 0.$$

Proof For $N \in \mathbb{N}$, define

$$\tilde{d}_N(f, g) := \int_{A_N} (1 \wedge d(f(\omega), g(\omega))) \mu(d\omega)$$

Then $\tilde{d}(f, f_n) \xrightarrow{n \rightarrow \infty} 0$ if and only if $\tilde{d}_N(f, f_n) \xrightarrow{n \rightarrow \infty} 0$ for all $N \in \mathbb{N}$.

" $\implies$ " Assume $f_n \xrightarrow{\text{meas}} f$. Then, for any $\varepsilon \in (0, 1)$,

$$\tilde{d}_N(f, f_n) \leq \mu(A_N \cap \{d(f, f_n) > \varepsilon\}) + \varepsilon \mu(A_N) \xrightarrow{n \rightarrow \infty} \varepsilon \mu(A_N).$$

Letting $\varepsilon \downarrow 0$ yields $\tilde{d}_N(f, f_n) \xrightarrow{n \rightarrow \infty} 0$.

" $\impliedby$ " Assume $\tilde{d}(f, f_n) \xrightarrow{n \rightarrow \infty} 0$. Let $B \in \mathcal{A}$ with $\mu(B) < \infty$. Fix $\delta > 0$ and choose $N \in \mathbb{N}$ large enough that $\mu(B \setminus A_N) < \delta$. Then, for $\varepsilon \in (0, 1)$,

$$\begin{aligned} \mu(B \cap \{d(f, f_n) > \varepsilon\}) &\leq \delta + \mu(A_N \cap \{d(f, f_n) > \varepsilon\}) \\ &\leq \delta + \varepsilon^{-1} \tilde{d}_N(f, f_n) \xrightarrow{n \rightarrow \infty} \delta. \end{aligned}$$

Letting $\delta \downarrow 0$ yields $\mu(B \cap \{d(f, f_n) > \varepsilon\}) \xrightarrow{n \rightarrow \infty} 0$; hence $f_n \xrightarrow{\text{meas}} f$.

Consider the most prominent case $E = \mathbb{R}$ equipped with the Euclidean metric. Here the integral is the basis for another concept of convergence.

Definition 6.8 (Mean convergence) Let $f, f_1, f_2, \dots \in \mathcal{L}^1(\mu)$. We say that the sequence $(f_n)_{n \in \mathbb{N}}$ converges in mean to f , symbolically

$$f_n \xrightarrow{L^1} f,$$

if $\|f_n - f\|_1 \xrightarrow{n \rightarrow \infty} 0$.

Remark 6.9 If $f_n \xrightarrow{L^1} f$, then in particular $\int f_n d\mu \xrightarrow{n \rightarrow \infty} \int f d\mu$.

Remark 6.10 If $f_n \xrightarrow{L^1} f$ and $f_n \xrightarrow{L^1} g$, then $f = g$ almost everywhere. Indeed, by the triangle inequality, $\|f - g\|_1 \leq \|f_n - f\|_1 + \|f_n - g\|_1 \xrightarrow{n \rightarrow \infty} 0$.

Remark 6.11 Both L^1 -convergence and almost everywhere convergence imply convergence in measure. All other implications are incorrect in general.

Theorem 6.12 (Fast convergence) Let (E, d) be a separable metric space. In order for the sequence $(f_n)_{n \in \mathbb{N}}$ of measurable maps $\Omega \rightarrow E$ to converge almost everywhere, it is sufficient that one of the following conditions holds.

(i) $E = \mathbb{R}$ and there is a $p \in [1, \infty)$ with $f_n \in \mathcal{L}^p(\mu)$ for all $n \in \mathbb{N}$ and there is an $f \in \mathcal{L}^p(\mu)$ with $\sum_{n=1}^{\infty} \|f_n - f\|_p < \infty$.

(ii) There is a measurable f with $\sum_{n=1}^{\infty} \mu(A \cap \{d(f, f_n) > \varepsilon\}) < \infty$ for all $\varepsilon > 0$ and for all $A \in \mathcal{A}$ with $\mu(A) < \infty$.

In both cases, we have $f_n \xrightarrow{n \rightarrow \infty} f$ almost everywhere.

(iii) E is complete and there is a summable sequence $(\varepsilon_n)_{n \in \mathbb{N}}$ such that

$$\sum_{n=1}^{\infty} \mu(A \cap \{d(f_n, f_{n+1}) > \varepsilon_n\}) < \infty \quad \text{for all } A \in \mathcal{A} \text{ with } \mu(A) < \infty$$

Proof Clearly, condition (i) implies (ii) since Markov's inequality yields that

$$\mu(\{|f - f_n| > \varepsilon\}) \leq \varepsilon^{-p} \|f - f_n\|_p^p$$

By Remark 6.3, it is enough to consider the case $\mu(\Omega) < \infty$.

Assume (ii). Let $B_n(\varepsilon) = \{d(f, f_n) > \varepsilon\}$ and $B(\varepsilon) = \limsup_{n \rightarrow \infty} B_n(\varepsilon)$. By the Borel-Cantelli

lemma, $\mu(B(\varepsilon)) = 0$. Let $N = \bigcup_{n=1}^{\infty} B(1/n)$. Then $\mu(N) = 0$ and $f_n(\omega) \xrightarrow{n \rightarrow \infty} f(\omega)$ for any $\omega \in \Omega \setminus N$.

Assume (iii). Let $B_n = \{d(f_n, f_{n+1}) > \varepsilon_n\}$ and $B = \limsup_{n \rightarrow \infty} B_n$. Then $\mu(B) = 0$ and $(f_n(\omega))_{n \in \mathbb{N}}$ is a Cauchy sequence in E for any $\omega \in \Omega \setminus B$. Since E is complete, the limit $f(\omega) := \lim_{n \rightarrow \infty} f_n(\omega)$ exists. For $\omega \in B$, define $f(\omega) = 0$.

Corollary 6.13 Let (E, d) be a separable metric space. Let $f, f_1, f_2, \dots$ be measurable maps $\Omega \rightarrow E$. Then the following statements are equivalent.

- (i) $f_n \xrightarrow{n \rightarrow \infty} f$ in measure.
- (ii) For any subsequence of $(f_n)_{n \in \mathbb{N}}$, there exists a sub-subsequence that converges to f almost everywhere.

Proof "(ii) $\implies$ (i)" Assume that (i) does not hold. Let $\tilde{d}$ be a metric that induces convergence in measure (see Theorem 6.7). Then there exists an $\varepsilon > 0$ and a subsequence $(f_{n_k})_{k \in \mathbb{N}}$ with $\tilde{d}(f_{n_k}, f) > \varepsilon$ for all $k \in \mathbb{N}$. Clearly, no subsequence of $(f_{n_k})_{k \in \mathbb{N}}$ converges to f in measure; hence neither converges almost everywhere.

"(i) $\implies$ (ii)" Now assume (i). Let $A_1, A_2, \dots \in \mathcal{A}$ with $A_N \uparrow \Omega$ and $\mu(A_N) < \infty$ for any $N \in \mathbb{N}$. Since $f_{n_k} \xrightarrow{\text{meas}} f$ for $k \rightarrow \infty$, we can choose a subsequence $(f_{n_{k_l}})_{l \in \mathbb{N}}$ such that $\mu\left(A_l \cap \left(d(f, f_{n_{k_l}}) > 1/l\right)\right) < 2^{-l}$ for any $l \in \mathbb{N}$. Hence, for each $N \in \mathbb{N}$, we have

$$\sum_{l=1}^{\infty} \mu\left(A_N \cap \left(d(f, f_{n_{k_l}}) > \frac{1}{l}\right)\right) \leq N\mu(A_N) + \sum_{l=N+1}^{\infty} 2^{-l} < \infty.$$

By Theorem 6.12(ii), $(f_{n_{k_l}})_{l \in \mathbb{N}}$ converges to f almost everywhere on A_N . By Remark 6.3, $(f_{n_{k_l}})_{l \in \mathbb{N}}$ converges to f almost everywhere.

Corollary 6.14 Let $(\Omega, \mathcal{A}, \mu)$ be a measure space in which almost everywhere convergence and convergence in measure do not coincide. Then there does not exist a topology on the set of measurable maps $\Omega \rightarrow E$ that induces almost everywhere convergence.

Proof Assume that there does exist a topology that induces almost everywhere convergence. Let $f, f_1, f_2, \dots$ be measurable maps with the property that $f_n \xrightarrow{\text{meas}} f$, but not $f_n \xrightarrow{n \rightarrow \infty} f$ almost everywhere. Now let U be an open set that contains f , but with $f_n \notin U$ for infinitely many $n \in \mathbb{N}$. Hence, let $(f_{n_k})_{k \in \mathbb{N}}$ be a subsequence with $f_{n_k} \notin U$ for all $k \in \mathbb{N}$. Since $f_{n_k} \xrightarrow{k \rightarrow \infty} f$ in measure, by Corollary 6.13, there exists a further subsequence $(f_{n_{k_l}})_{l \in \mathbb{N}}$ of $(f_{n_k})_{k \in \mathbb{N}}$ with $f_{n_{k_l}} \xrightarrow{l \rightarrow \infty} f$ almost everywhere. However, then $f_{n_{k_l}} \in U$ for all but finitely many l , which yields a contradiction!

Corollary 6.15 Let (E, d) be a separable complete metric space. Let $(f_n)_{n \in \mathbb{N}}$ be a Cauchy sequence in measure in E ; that is, for any $A \in \mathcal{A}$ with $\mu(A) < \infty$ and any $\varepsilon > 0$, we have

$$\mu(A \cap \{d(f_m, f_n) > \varepsilon\}) \longrightarrow 0 \quad \text{for } m, n \rightarrow \infty$$

Then $(f_n)_{n \in \mathbb{N}}$ converges in measure.

Proof Without loss of generality, we may assume $\mu(\Omega) < \infty$. Choose a subsequence $(f_{n_k})_{k \in \mathbb{N}}$ such that

$$\mu(\{d(f_n, f_{n_k}) > 2^{-k}\}) < 2^{-k} \quad \text{for all } n \geq n_k.$$

By Theorem 6.12(iii), there is an f with $f_{n_k} \xrightarrow{k \rightarrow \infty} f$ almost everywhere; hence, in particular, $\mu(\{d(f_{n_k}, f) > \varepsilon/2\}) \xrightarrow{k \rightarrow \infty} 0$ for all $\varepsilon > 0$. Now

$$\mu(\{d(f_n, f) > \varepsilon\}) \leq \mu(\{d(f_{n_k}, f) > \varepsilon/2\}) + \mu(\{d(f_{n_k}, f) > \varepsilon/2\}).$$

If k is large enough that $2^{-k} < \varepsilon/2$ and if $n \geq n_k$, then the first summand is smaller than 2^{-k} . Hence we have $\mu(\{d(f_n, f) > \varepsilon\}) \xrightarrow{n \rightarrow \infty} 0$; that is, $f_n \xrightarrow{\text{meas}} f$.

Exercise 6.1.1 Let Ω be countable. Show that convergence in probability implies almost everywhere convergence.

Exercise 6.1.2 Give an example of a sequence that

- (i) converges in L^1 but not almost everywhere,
- (ii) converges almost everywhere but not in L^1 .

Exercise 6.1.3 (Egorov's theorem (1911), [41]) Let $(\Omega, \mathcal{A}, \mu)$ be a finite measure space and let $f_1, f_2, \dots$ be measurable functions that converge to some f almost everywhere. Show that, for every $\varepsilon > 0$, there is a set $A \in \mathcal{A}$ with $\mu(\Omega \setminus A) < \varepsilon$ and $\sup_{\omega \in A} |f_n(\omega) - f(\omega)| \xrightarrow{n \rightarrow \infty} 0$.

Exercise 6.1.4 Let $X_1, X_2, \dots$ be independent, square integrable, centered random variables with $\sum_{i=1}^{\infty} \text{Var}[X_i] < \infty$. Show that there exists a square integrable X with $X = \lim_{n \rightarrow \infty} \sum_{i=1}^n X_i$ almost surely.

5.6.2 6.2 Uniform Integrability

From the preceding section, we can conclude that convergence in measure plus existence of L^1 limit points implies L^1 -convergence. Hence convergence in measure plus relative sequential compactness in L^1 yields convergence in L^1 . In this section, we study a criterion for relative sequential compactness in L^1 , the so-called uniform integrability.

Definition 6.16 A family $\mathcal{F} \subset \mathcal{L}^1(\mu)$ is called uniformly integrable if

$$\inf_{0 \leq g \in \mathcal{L}^1(\mu)} \sup_{f \in \mathcal{F}} \int (|f| - g)^+ d\mu = 0 \quad (6.2)$$

Theorem 6.17 The family $\mathcal{F} \subset \mathcal{L}^1(\mu)$ is uniformly integrable if and only if

$$\inf_{0 \leq \tilde{g} \in \mathcal{L}^1(\mu)} \sup_{f \in \mathcal{F}} \int_{\{|f| > \tilde{g}\}} |f| d\mu = 0. \quad (6.3)$$

If $\mu(\Omega) < \infty$, then uniform integrability is equivalent to either of the following two conditions:

- (i) $\inf_{a \in [0, \infty)} \sup_{f \in \mathcal{F}} \int (|f| - a)^+ d\mu = 0$,
- (ii) $\inf_{a \in [0, \infty)} \sup_{f \in \mathcal{F}} \int_{\{|f| > a\}} |f| d\mu = 0$.

Proof Clearly, $(|f| - g)^+ \leq |f| \cdot \mathbb{1}_{\{|f| > g\}}$; hence (6.3) implies uniform integrability.

Now assume (6.2). For $\varepsilon > 0$, choose $g_\varepsilon \in \mathcal{L}^1(\mu)$ such that

$$\sup_{f \in \mathcal{F}} \int (|f| - g_\varepsilon)^+ d\mu \leq \varepsilon \quad (6.4)$$

Define $\tilde{g}_\varepsilon = 2g_{\varepsilon/2}$. Then, for $f \in \mathcal{F}$,

$$\int_{\{|f| > \tilde{g}_\varepsilon\}} |f| d\mu \leq \int_{\{|f| > \tilde{g}_\varepsilon\}} (|f| - g_{\varepsilon/2})^+ d\mu + \int_{\{|f| > \tilde{g}_\varepsilon\}} g_{\varepsilon/2} d\mu.$$

By construction, $\int_{\{|f| > \tilde{g}_\varepsilon\}} (|f| - g_{\varepsilon/2})^+ d\mu \leq \varepsilon/2$ and

$$g_{\varepsilon/2} \mathbb{1}_{\{|f| > \tilde{g}_\varepsilon\}} \leq (|f| - g_{\varepsilon/2})^+ \mathbb{1}_{\{|f| > \tilde{g}_\varepsilon\}};$$

hence also

$$\int_{\{|f| > \tilde{g}_\varepsilon\}} g_{\varepsilon/2} d\mu \leq \int_{\{|f| > \tilde{g}_\varepsilon\}} (|f| - g_{\varepsilon/2})^+ d\mu \leq \frac{\varepsilon}{2}$$

Summing up, we have

$$\sup_{f \in \mathcal{F}} \int_{\{|f| > \tilde{g}_\varepsilon\}} |f| d\mu \leq \varepsilon \quad (6.5)$$

Clearly, (ii) implies (i). If $\mu(\Omega) < \infty$, then (i) implies uniform integrability of $\mathcal{F}$ since the infimum is taken over the smaller set of constant functions. We still have to show that uniform integrability implies (ii). Accordingly, assume $\mathcal{F}$ is uniformly integrable (but not necessarily $\mu(\Omega) < \infty$). For any $\varepsilon > 0$ (and g_ε and $\tilde{g}_\varepsilon$ as above), choose a_ε such that $\int_{\{\tilde{g}_{\varepsilon/2} > a_\varepsilon\}} \tilde{g}_{\varepsilon/2} d\mu < \frac{\varepsilon}{2}$. Then

$$\int_{\{|f| > a_\varepsilon\}} |f| d\mu \leq \int_{\{|f| > \tilde{g}_{\varepsilon/2}\}} |f| d\mu + \int_{\{\tilde{g}_{\varepsilon/2} > a_\varepsilon\}} \tilde{g}_{\varepsilon/2} d\mu < \varepsilon.$$

Theorem 6.18

- (i) If $\mathcal{F} \subset \mathcal{L}^1(\mu)$ is a finite set, then $\mathcal{F}$ is uniformly integrable.
- (ii) If $\mathcal{F}, \mathcal{G} \subset \mathcal{L}^1(\mu)$ are uniformly integrable, then $(f + g : f \in \mathcal{F}, g \in \mathcal{G})$, $(f - g : f \in \mathcal{F}, g \in \mathcal{G})$ and $\{|f| : f \in \mathcal{F}\}$ are also uniformly integrable.
- (iii) If $\mathcal{F}$ is uniformly integrable and if, for any $g \in \mathcal{G}$, there exists an $f \in \mathcal{F}$ with $|g| \leq |f|$, then $\mathcal{G}$ is also uniformly integrable.

Proof The proof is simple and is left as an exercise.

The following theorem describes a very useful criterion for uniform integrability. We will use it in many places.

Theorem 6.19 For finite μ , $\mathcal{F} \subset \mathcal{L}^1(\mu)$ is uniformly integrable if and only if there is a function $H : [0, \infty) \rightarrow [0, \infty)$ with $\lim_{x \rightarrow \infty} H(x)/x = \infty$ and

$$\sup_{f \in \mathcal{F}} \int H(|f|) d\mu < \infty$$

H can be chosen to be monotone increasing and convex.

Proof " $\Leftarrow$ " Assume there is an H with the advertised properties. Then $K_a := \inf_{x \geq a} \frac{H(x)}{x} \uparrow \infty$ if $a \uparrow \infty$. Hence, for $a > 0$,

$$\begin{aligned} \sup_{f \in \mathcal{F}} \int_{\{|f| \geq a\}} |f| d\mu &\leq \frac{1}{K_a} \sup_{f \in \mathcal{F}} \int_{\{|f| \geq a\}} H(|f|) d\mu \\ &\leq \frac{1}{K_a} \sup_{f \in \mathcal{F}} \int H(|f|) d\mu \xrightarrow{a \rightarrow \infty} 0 \end{aligned}$$

" $\Rightarrow$ " Assume $\mathcal{F}$ is uniformly integrable. As we have $\mu(\Omega) < \infty$, by Theorem 6.17, there exists a sequence $a_n \uparrow \infty$ with

$$\sup_{f \in \mathcal{F}} \int (|f| - a_n)^+ d\mu < 2^{-n}$$

Define

$$H(x) = \sum_{n=1}^{\infty} (x - a_n)^+ \quad \text{for any } x \geq 0$$

As a sum of convex functions, H is convex. Further, for any $n \in \mathbb{N}$ and $x \geq 2a_n$, $H(x)/x \geq \sum_{k=1}^n (1 - a_k/x)^+ \geq n/2$; hence we have $H(x)/x \uparrow \infty$. Finally, by monotone convergence, for any $f \in \mathcal{F}$,

$$\int H(|f(\omega)|) \mu(d\omega) = \sum_{n=1}^{\infty} \int (|f| - a_n)^+ d\mu \leq \sum_{n=1}^{\infty} 2^{-n} = 1$$

Recall the notation $\|\cdot\|_p$ from Definition 4.16.

Definition 6.20 Let $p \in [1, \infty]$. A family $\mathcal{F} \subset \mathcal{L}^p(\mu)$ is called bounded in $\mathcal{L}^p(\mu)$ if $\sup \{\|f\|_p : f \in \mathcal{F}\} < \infty$.

Corollary 6.21 Let $\mu(\Omega) < \infty$ and $p > 1$. If $\mathcal{F}$ is bounded in $\mathcal{L}^p(\mu)$, then $\mathcal{F}$ is uniformly integrable.

Proof Apply Theorem 6.19 with the convex map $H(x) = x^p$.

Corollary 6.22 If $(X_i)_{i \in I}$ is a family of random variables with

$$\sup \{\mathbf{E}[X_i] : i \in I\} < \infty \quad \text{and} \quad \sup \{\text{Var}[X_i] : i \in I\} < \infty,$$

then $(X_i)_{i \in I}$ is uniformly integrable.

Proof Since $\mathbf{E}[X_i^2] = \mathbf{E}[X_i]^2 + \text{Var}[X_i]$, $i \in I$, is bounded, this follows from Corollary 6.21 with $p = 2$.

Lemma 6.23 There is a map $h \in \mathcal{L}^1(\mu)$ with $h > 0$ almost everywhere.

Proof Let $A_1, A_2, \dots \in \mathcal{A}$ with $A_n \uparrow \Omega$ and $\mu(A_n) < \infty$ for all $n \in \mathbb{N}$. Define

$$h = \sum_{n=1}^{\infty} 2^{-n} (1 + \mu(A_n))^{-1} \mathbb{1}_{A_n}$$

Then $h > 0$ almost everywhere and $\int h d\mu = \sum_{n=1}^{\infty} 2^{-n} \frac{\mu(A_n)}{1 + \mu(A_n)} \leq 1$.

Theorem 6.24 A family $\mathcal{F} \subset \mathcal{L}^1(\mu)$ is uniformly integrable if and only if the following two conditions are fulfilled.

(i) $C := \sup_{f \in \mathcal{F}} \int |f| d\mu < \infty$.

(ii) There is a function $0 \leq h \in \mathcal{L}^1(\mu)$ such that for any $\varepsilon > 0$, there is a $\delta(\varepsilon) > 0$ with

$$\sup_{f \in \mathcal{F}} \int_A |f| d\mu \leq \varepsilon \quad \text{for all } A \in \mathcal{A} \quad \text{such that} \quad \int_A h d\mu < \delta(\varepsilon)$$

If $\mu(\Omega) < \infty$, then (ii) is equivalent to (iii):

(iii) For all $\varepsilon > 0$, there is a $\delta(\varepsilon) > 0$ such that

$$\sup_{f \in \mathcal{F}} \int_A |f| d\mu \leq \varepsilon \quad \text{for all } A \in \mathcal{A} \text{ with } \mu(A) < \delta(\varepsilon)$$

Proof " $\implies$ " Let $\mathcal{F}$ be uniformly integrable. Let $h \in \mathcal{L}^1(\mu)$ with $h > 0$ a.e. Let $\varepsilon > 0$ and let $\tilde{g}_{\varepsilon/3}$ be an $\varepsilon/3$ -bound for $\mathcal{F}$ (as in (6.5)). Since $\{\tilde{g}_{\varepsilon/3} \geq \alpha h\} \downarrow \emptyset$ for $\alpha \rightarrow \infty$, for sufficiently large $\alpha = \alpha(\varepsilon)$, we have

$$\int_{\{\tilde{g}_{\varepsilon/3} \geq \alpha h\}} \tilde{g}_{\varepsilon/3} d\mu < \frac{\varepsilon}{3}$$

Letting $\delta(\varepsilon) := \frac{\varepsilon}{3\alpha(\varepsilon)}$, we get for any $A \in \mathcal{A}$ with $\int_A h d\mu < \delta(\varepsilon)$ and any $f \in \mathcal{F}$,

$$\begin{aligned} \int_A |f| d\mu &\leq \int_{\{|f| > \tilde{g}_{\varepsilon/3}\}} |f| d\mu + \int_A \tilde{g}_{\varepsilon/3} d\mu \\ &\leq \frac{\varepsilon}{3} + \alpha \int_A h d\mu + \int_{\{\tilde{g}_{\varepsilon/3} \geq \alpha h\}} \tilde{g}_{\varepsilon/3} d\mu \leq \varepsilon \end{aligned}$$

Hence we have shown (ii). In the above computation, let $A = \Omega$ to obtain

$$\int |f| d\mu \leq \frac{2\varepsilon}{3} + \alpha \int h d\mu < \infty$$

Hence we have also shown (i).

"%" Assume (i) and (ii). Let $\varepsilon > 0$. Choose h and $\delta(\varepsilon) > 0$ as in (ii) and C as in (i). Define $\tilde{h} = \frac{C}{\delta(\varepsilon)}h$. Then

$$\int_{\{|f| > \tilde{h}\}} h d\mu = \frac{\delta(\varepsilon)}{C} \int_{\{|f| > \tilde{h}\}} \tilde{h} d\mu \leq \frac{\delta(\varepsilon)}{C} \int |f| d\mu \leq \delta(\varepsilon);$$

hence, by assumption, $\int_{\{|f| > \tilde{h}\}} |f| d\mu < \varepsilon$.

"(ii) $\implies$ (iii)" Assume (ii). Let $\varepsilon > 0$ and choose $\delta = \delta(\varepsilon)$ as in (ii). Choose $K < \infty$ large enough that $\int_{\{h \geq K\}} h d\mu < \delta/2$. For all $A \in \mathcal{A}$ with $\mu(A) < \delta/(2K)$, we obtain

$$\int_A h d\mu \leq K\mu(A) + \int_{\{h \geq K\}} h d\mu < \delta$$

hence $\int_A |f| d\mu \leq \varepsilon$ for all $f \in \mathcal{F}$.

"(iii) $\implies$ (ii)" Assume (iii) and $\mu(\Omega) < \infty$. Then $h \equiv 1$ serves the purpose.

We come to the main theorem of this section.

Theorem 6.25 Let $\{f_n : n \in \mathbb{N}\} \subset \mathcal{L}^1(\mu)$. The following statements are equivalent.

- (i) There is an $f \in \mathcal{L}^1(\mu)$ with $f_n \xrightarrow{n \rightarrow \infty} f$ in L^1 .
 - (ii) $(f_n)_{n \in \mathbb{N}}$ is an $\mathcal{L}^1(\mu)$ -Cauchy sequence; that is, $\|f_n - f_m\|_1 \rightarrow 0$ for $m, n \rightarrow \infty$.
 - (iii) $(f_n)_{n \in \mathbb{N}}$ is uniformly integrable and there is a measurable map f such that $f_n \xrightarrow{\text{meas}} f$.
- The limits in (i) and (iii) coincide.

Proof "(i) $\implies$ (ii)" This is evident.

"(ii) $\implies$ (iii)" For any $\varepsilon > 0$, there is an $n_\varepsilon \in \mathbb{N}$ such that $\|f_n - f_{n_\varepsilon}\|_1 < \varepsilon$ for all $n \geq n_\varepsilon$. Hence $\|(|f_n| - |f_{n_\varepsilon}|)^+\|_1 < \varepsilon$ for all $n \geq n_\varepsilon$. Thus $g_\varepsilon = \max\{|f_1|, \dots, |f_{n_\varepsilon}|\}$ is an ε -bound for $(f_n)_{n \in \mathbb{N}}$ (as in (6.4)). For $\varepsilon > 0$, let

$$\mu(\{|f_m - f_n| > \varepsilon\}) \leq \varepsilon^{-1} \|f_m - f_n\|_1 \rightarrow 0 \quad \text{for } m, n \rightarrow \infty.$$

Thus $(f_n)_{n \in \mathbb{N}}$ is also a Cauchy sequence in measure; hence it converges in measure by Corollary 6.15.

"(iii) $\implies$ (i)" Let f be the limit in measure of the sequence $(f_n)_{n \in \mathbb{N}}$. Assume that $(f_n)_{n \in \mathbb{N}}$ does not converge to f in L^1 . Then there is an $\varepsilon > 0$ and a subsequence $(f_{n_k})_{k \in \mathbb{N}}$ with

$$\|f - f_{n_k}\|_1 > 2\varepsilon \quad \text{for all } k \in \mathbb{N} \tag{6.6}$$

Here we define $\|f - f_{n_k}\|_1 = \infty$ if $f \notin \mathcal{L}^1(\mu)$. By Corollary 6.13, there is a subsequence $(f_{n'_k})_{k \in \mathbb{N}}$ of $(f_{n_k})_{k \in \mathbb{N}}$ with $f_{n'_k} \xrightarrow{k \rightarrow \infty} f$ almost everywhere. By Fatou's lemma (Theorem 4.21) with 0 as a minorant, we thus get

$$\int |f| d\mu \leq \liminf_{k \rightarrow \infty} \int |f_{n'_k}| d\mu < \infty$$

Hence $f \in \mathcal{L}^1(\mu)$. By Theorem 6.18(ii) (with $\mathcal{G} = \{f\}$), we obtain that the family $(f - f_{n'_k})_{k \in \mathbb{N}}$ is uniformly integrable; hence there is a $0 \leq g \in \mathcal{L}^1(\mu)$ such that $\int (|f - f_{n'_k}| - g)^+ d\mu < \varepsilon$. Define

$$g_k = |f_{n'_k} - f| \wedge g \quad \text{for } k \in \mathbb{N}$$

Then $g_k \xrightarrow{k \rightarrow \infty} 0$ almost everywhere and $g - g_k \geq 0$. By Fatou's lemma,

$$\begin{aligned} \limsup_{k \rightarrow \infty} \int g_k d\mu &= \int g d\mu - \liminf_{k \rightarrow \infty} \int (g - g_k) d\mu \\ &\leq \int g d\mu - \int \left(\lim_{k \rightarrow \infty} (g - g_k) \right) d\mu = 0 \end{aligned}$$

Since $|f - f_{n'_k}| = (|f - f_{n'_k}| - g)^+ + g_k$, this implies that

$$\limsup_{k \rightarrow \infty} \|f - f_{n'_k}\|_1 \leq \limsup_{k \rightarrow \infty} \int (|f - f_{n'_k}| - g)^+ d\mu + \limsup_{k \rightarrow \infty} \int g_k d\mu \leq \varepsilon$$

contradicting (6.6).

Corollary 6.26 (Lebesgue's convergence theorem, dominated convergence) Let f be measurable and let $(f_n)_{n \in \mathbb{N}}$ be a sequence in $\mathcal{L}^1(\mu)$ with $f_n \xrightarrow{n \rightarrow \infty} f$ in measure. Assume that there is an integrable dominating function $0 \leq g \in \mathcal{L}^1(\mu)$ with $|f_n| \leq g$ almost everywhere for all $n \in \mathbb{N}$. Then $f \in \mathcal{L}^1(\mu)$ and $f_n \xrightarrow{n \rightarrow \infty} f$ in L^1 ; hence in particular $\int f_n d\mu \xrightarrow{n \rightarrow \infty} \int f d\mu$.

Proof This is a consequence of Theorem 6.25, as the dominating function ensures uniform integrability of the sequence $(f_n)_{n \in \mathbb{N}}$.

Exercise 6.2.1 Let $H \in \mathcal{L}^1(\mu)$ with $H > 0$ μ -a.e. (see Lemma 6.23) and let (E, d) be a separable metric space. For measurable $f, g : \Omega \rightarrow E$, define

$$d_H(f, g) := \int (1 \wedge d(f(\omega), g(\omega))) H(\omega) \mu(d\omega)$$

- (i) Show that d_H is a metric that induces convergence in measure.
- (ii) Show that d_H is complete if (E, d) is complete.

5.6.3 6.3 Exchanging Integral and Differentiation

We study how properties such as continuity and differentiability of functions of two variables behave under integration with respect to one of the variables.

Theorem 6.27 (Continuity lemma) Let (E, d) be a metric space, $x_0 \in E$ and let $f : \Omega \times E \rightarrow \mathbb{R}$ be a map with the following properties.

- (i) For any $x \in E$, the map $\omega \mapsto f(\omega, x)$ is in $\mathcal{L}^1(\mu)$.
- (ii) For almost all $\omega \in \Omega$, the map $x \mapsto f(\omega, x)$ is continuous at the point x_0 .
- (iii) There is a map $h \in \mathcal{L}^1(\mu)$, $h \geq 0$, such that $|f(\cdot, x)| \leq h \mu$ -a.e. for all $x \in E$.

Then the map $F : E \rightarrow \mathbb{R}$, $x \mapsto \int f(\omega, x) \mu(d\omega)$ is continuous at x_0 .

Proof Let $(x_n)_{n \in \mathbb{N}}$ be a sequence in E with $\lim_{n \rightarrow \infty} x_n = x_0$. Define $f_n = f(\cdot, x_n)$. By assumption, $|f_n| \leq h$ and $f_n \xrightarrow{n \rightarrow \infty} f(\cdot, x_0)$ almost everywhere. By the dominated convergence theorem (Corollary 6.26), we get

$$F(x_n) = \int f_n d\mu \xrightarrow{n \rightarrow \infty} \int f(\cdot, x_0) d\mu = F(x_0)$$

Hence F is continuous at x_0 .

Theorem 6.28 (Differentiation lemma) Let $I \subset \mathbb{R}$ be a nontrivial open interval and let $f : \Omega \times I \rightarrow \mathbb{R}$ be a map with the following properties.

- (i) For any $x \in I$, the map $\omega \mapsto f(\omega, x)$ is in $\mathcal{L}^1(\mu)$.
- (ii) For almost all $\omega \in \Omega$, the map $I \rightarrow \mathbb{R}$, $x \mapsto f(\omega, x)$ is differentiable with derivative f' .
- (iii) There is a map $h \in \mathcal{L}^1(\mu)$, $h \geq 0$, such that $|f'(\cdot, x)| \leq h \mu$ -a.e. for all $x \in I$.

Then, for any $x \in I$, $f'(\cdot, x) \in \mathcal{L}^1(\mu)$ and the function $F : I \rightarrow \mathbb{R}$, $x \mapsto \int f(\omega, x) \mu(d\omega)$ is differentiable with derivative

$$F'(x) = \int f'(\omega, x) \mu(d\omega)$$

Proof Let $x_0 \in I$ and let $(x_n)_{n \in \mathbb{N}}$ be a sequence in I with $x_n \neq x_0$ for all $n \in \mathbb{N}$ and such that $\lim_{n \rightarrow \infty} x_n = x_0$. We show that, along the sequence $(x_n)_{n \in \mathbb{N}}$, the difference quotients converge. Define

$$g_n(\omega) = \frac{f(\omega, x_n) - f(\omega, x_0)}{x_n - x_0} \quad \text{for all } \omega \in \Omega$$

By assumption (ii), we have

$$g_n \xrightarrow{n \rightarrow \infty} f'(\cdot, x_0) \quad \mu\text{-almost everywhere.}$$

By the mean value theorem of calculus, for all $n \in \mathbb{N}$ and for almost all $\omega \in \Omega$, there exists a $y_n(\omega) \in I$ with $g_n(\omega) = f'(\omega, y_n(\omega))$. In particular, $|g_n| \leq h$ almost everywhere for all $n \in \mathbb{N}$. By the dominated convergence theorem (Corollary 6.26), the limiting function $f'(\cdot, x_0)$ is in $\mathcal{L}^1(\mu)$ and

$$\lim_{n \rightarrow \infty} \frac{F(x_n) - F(x_0)}{x_n - x_0} = \lim_{n \rightarrow \infty} \int g_n(\omega) \mu(d\omega) = \int f'(\omega, x_0) \mu(d\omega)$$

Example 6.29 (Laplace transform) Let X be a nonnegative random variable on $(\Omega, \mathcal{A}, \mathbf{P})$. Using the notation of Theorem 6.28, let $I = [0, \infty)$ and $f(x, \lambda) = e^{-\lambda x}$ for $\lambda \in I$. Then

$$F(\lambda) = \mathbf{E}[e^{-\lambda X}]$$

is infinitely often differentiable in $(0, \infty)$. The first two derivatives of F are $F'(\lambda) = -\mathbf{E}[Xe^{-\lambda X}]$ and $F''(\lambda) = \mathbf{E}[X^2 e^{-\lambda X}]$. Successively, we get the n th derivative $F^{(n)}(\lambda) = \mathbf{E}[(-X)^n e^{-\lambda X}]$. By monotone convergence, we get

$$\mathbf{E}[X] = -\lim_{\lambda \downarrow 0} F'(\lambda) \tag{6.7}$$

and

$$\mathbf{E}[X^n] = (-1)^n \lim_{\lambda \downarrow 0} F^{(n)}(\lambda) \quad \text{for all } n \in \mathbb{N} \tag{6.8}$$

Indeed, for $\varepsilon > 0$ and $I = (\varepsilon, \infty)$, we have

$$\sup_{x \geq 0, \lambda \in I} \left| \frac{d}{d\lambda} f(x, \lambda) \right| = \sup_{x \geq 0, \lambda \in I} x e^{-\lambda x} = \varepsilon^{-1} e^{-1} < \infty.$$

Thus F fulfills the assumptions of Theorem 6.28. Inductively, we get the statement for $F^{(n)}$ since

$$\left| \frac{d^n}{d\lambda^n} f(x, \lambda) \right| \leq (n/\varepsilon)^n e^{-n} < \infty \quad \text{for } x \geq 0 \text{ and } \lambda \geq \varepsilon.$$

Exercise 6.3.1 Let X be a random variable on $(\Omega, \mathcal{A}, \mathbf{P})$ and let

$$\Lambda(t) := \log(\mathbf{E}[e^{tX}]) \quad \text{for all } t \in \mathbb{R}$$

Show that $D := \{t \in \mathbb{R} : \Lambda(t) < \infty\}$ is a nonempty interval and that Λ is infinitely often differentiable in the interior of D .

5.7 7 L^p -Spaces and the Radon-Nikodym Theorem

In this chapter, we study the spaces of functions whose p th power is integrable. In Section 7.2, we first derive some of the important inequalities (Hölder, Minkowski, Jensen) and then in Section 7.3 investigate the case $p = 2$ in more detail. Apart from the inequalities, the important results for probability theory are Lebesgue's decomposition theorem and the Radon-Nikodym theorem in Section 7.4. At first reading, some readers might wish to skip some of the more analytic parts of this chapter.

5.7.1 7.1 Definitions

We always assume that $(\Omega, \mathcal{A}, \mu)$ is a σ -finite measure space. In Definition 4.16, for measurable $f : \Omega \rightarrow \overline{\mathbb{R}}$, we defined

$$\|f\|_p := \left(\int |f|^p d\mu \right)^{1/p} \quad \text{for } p \in [1, \infty)$$

and

$$\|f\|_\infty := \inf\{K \geq 0 : \mu(|f| > K) = 0\}.$$

Further, we defined the spaces of functions where these expressions are finite:

$$\mathcal{L}^p(\Omega, \mathcal{A}, \mu) = \mathcal{L}^p(\mathcal{A}, \mu) = \mathcal{L}^p(\mu) = \{f : \Omega \rightarrow \overline{\mathbb{R}} \text{ measurable and } \|f\|_p < \infty\}.$$

We saw that $\|\cdot\|_1$ is a seminorm on $\mathcal{L}^1(\mu)$. Here our first goal is to change $\|\cdot\|_p$ into a proper norm for all $p \in [1, \infty]$. Apart from the fact that we still have to show the triangle inequality, to this end, we have to change the space a little bit since we only have

$$\|f - g\|_p = 0 \iff f = g \quad \mu\text{-a.e.}$$

For a proper norm (that is, not only a seminorm), the left-hand side has to imply equality (not only a.e.) of f and g . Hence we now consider f and g as equivalent if $f = g$ almost everywhere. Thus let

$$\mathcal{N} = \{f \text{ is measurable and } f = 0 \mu\text{-a.e.}\}$$

For any $p \in [1, \infty]$, $\mathcal{N}$ is a subvector space of $\mathcal{L}^p(\mu)$. Thus formally we can build the factor space. This is the standard procedure in order to change a seminorm into a proper norm.

Definition 7.1 (Factor space) For any $p \in [1, \infty]$, define

$$L^p(\Omega, \mathcal{A}, \mu) := \mathcal{L}^p(\Omega, \mathcal{A}, \mu) / \mathcal{N} = \{\bar{f} := f + \mathcal{N} : f \in \mathcal{L}^p(\mu)\}.$$

For $\bar{f} \in L^p(\mu)$, define $\|\bar{f}\|_p = \|f\|_p$ for any $f \in \bar{f}$. Also let $\int \bar{f} d\mu = \int f d\mu$ if this expression is defined for f .

Note that $\|\bar{f}\|_p$ and $\int \bar{f} d\mu$ do not depend on the choice of the representative $f \in \bar{f}$. Recall from Theorem 4.19 that $\int \bar{f} d\mu$ is well-defined if $f \in \mathcal{L}^p(\mu)$ and if μ is finite but it need not be if μ is infinite.

We first investigate convergence with respect to $\|\cdot\|_p$. To this end, we extend the corresponding theorem (Theorem 6.25) on convergence with respect to $\|\cdot\|_1$.

Definition 7.2 Let $p \in [1, \infty]$ and $f, f_1, f_2, \dots \in \mathcal{L}^p(\mu)$. If $\|f_n - f\|_p \xrightarrow{n \rightarrow \infty} 0$, then we say that $(f_n)_{n \in \mathbb{N}}$ converges to f in $L^p(\mu)$ and we write $f_n \xrightarrow{L^p} f$.

Theorem 7.3 Let $p \in [1, \infty]$ and $f_1, f_2, \dots \in \mathcal{L}^p(\mu)$. Then the following statements are equivalent:

(i) There is an $f \in \mathcal{L}^p(\mu)$ with $f_n \xrightarrow{L^p} f$.

(ii) $(f_n)_{n \in \mathbb{N}}$ is a Cauchy sequence in $\mathcal{L}^p(\mu)$.

If $p < \infty$, then, in addition, (i) and (ii) are equivalent to:

(iii) $(|f_n|^p)_{n \in \mathbb{N}}$ is uniformly integrable and there exists a measurable f with $f_n \xrightarrow{\text{meas}} f$.

The limits in (i) and (ii) coincide.

Proof For $p = \infty$, the equivalence of (i) and (ii) is a simple consequence of the triangle inequality.

Now let $p \in [1, \infty)$. The proof is similar to the proof of Theorem 6.25.

"(i) $\implies$ (ii)" Note that $|x + y|^p \leq 2^p (|x|^p + |y|^p)$ for all $x, y \in \mathbb{R}$. Hence

$$\|f_m - f_n\|_p^p \leq 2^p \left(\|f_m - f\|_p^p + \|f_n - f\|_p^p \right) \xrightarrow{n \rightarrow \infty} 0 \quad \text{for } m, n \rightarrow \infty.$$

"(ii) $\implies$ (iii)" This works as in the proof of Theorem 6.25.

"(iii) $\implies$ (i)" Since $|f_n|^p \xrightarrow{n \rightarrow \infty} |f|^p$ in measure, by Theorem 6.25, we have $|f|^p \in \mathcal{L}^1(\mu)$ and hence $f \in \mathcal{L}^p(\mu)$. For $n \in \mathbb{N}$, define $g_n = |f_n - f|^p$. Then

$g_n \xrightarrow{n \rightarrow \infty} 0$ in measure, and $(g_n)_{n \in \mathbb{N}}$ is uniformly integrable since $g_n \leq 2^p (|f_n|^p + |f|^p)$. Hence we get (by Theorem 6.25) $\|f_n - f\|_p^p = \|g_n\|_1 \xrightarrow{n \rightarrow \infty} 0$.

Exercise 7.1.1 Let $(X_i)_{i \in \mathbb{N}}$ be independent, square integrable random variables with $\mathbf{E}[X_i] = 0$ for all $i \in \mathbb{N}$.

(i) Show that $\sum_{i=1}^{\infty} \text{Var}[X_i] < \infty$ implies that there exists a real random variable X with $\sum_{i=1}^n X_i \xrightarrow{n \rightarrow \infty} X$ almost surely.

(ii) Does the converse implication hold in (i)?

Exercise 7.1.2 Let $f : \Omega \rightarrow \mathbb{R}$ be measurable. Show that the following hold.

(i) If $\int |f|^p d\mu < \infty$ for some $p \in (0, \infty)$, then $\|f\|_p \xrightarrow{p \rightarrow \infty} \|f\|_{\infty}$.

(ii) The integrability condition in (i) cannot be waived.

Exercise 7.1.3 Let $p \in (1, \infty)$, $f \in \mathcal{L}^p(\lambda)$, where λ is the Lebesgue measure on $\mathbb{R}$. Let $T : \mathbb{R} \rightarrow \mathbb{R}, x \mapsto x + 1$. Show that

$$\frac{1}{n} \sum_{k=0}^{n-1} f \circ T^k \xrightarrow{n \rightarrow \infty} 0 \quad \text{in } L^p(\lambda)$$

5.7.2 7.2 Inequalities and the Fischer-riesz Theorem

We present one of the most important inequalities of probability theory, Jensen's inequality for convex functions, and indicate how to derive from it Hölder's inequality and Minkowski's inequality. They in turn yield the triangle inequality for $\|\cdot\|_p$ and help in determining the dual space of $L^p(\mu)$. However, for the formal proofs of the latter inequalities, we will follow a different route.

Before stating Jensen's inequality, we give a primer on the basics of convexity of sets and functions.

Definition 7.4 A subset G of a vector space (or of an affine linear space) is called convex if, for any two points $x, y \in G$ and any $\lambda \in [0, 1]$, we have $\lambda x + (1 - \lambda)y \in G$.

Example 7.5

(i) The convex subsets of $\mathbb{R}$ are the intervals.

(ii) A linear subspace of a vector space is convex.

(iii) The set of all probability measures on a measurable space is a convex set.

Definition 7.6 Let G be a convex set. A map $\varphi : G \rightarrow \mathbb{R}$ is called convex if for any two points $x, y \in G$ and any $\lambda \in [0, 1]$, we have

$$\varphi(\lambda x + (1 - \lambda)y) \leq \lambda \varphi(x) + (1 - \lambda)\varphi(y)$$

φ is called concave if $(-\varphi)$ is convex.

Let $I \subset \mathbb{R}$ be an interval. Let $\varphi : I \rightarrow \mathbb{R}$ be continuous and in the interior I° twice continuously differentiable with second derivative φ'' . Then φ is convex if and only if $\varphi''(x) \geq 0$ for all $x \in I^\circ$. To put it differently, the first derivative φ' of a convex function is a monotone increasing function. In the next theorem, we will see that this is still true even if φ is not twice continuously differentiable when we pass to the right-sided derivative $D^+\varphi$ (or to the left-sided derivative), which we show always exists.

Theorem 7.7 Let $I \subset \mathbb{R}$ be an interval with interior I° and let $\varphi : I \rightarrow \mathbb{R}$ be a convex map. Then:

- (i) φ is continuous on I° and hence measurable with respect to $\mathcal{B}(I)$.
- (ii) For $x \in I^\circ$, define the function of difference quotients

$$g_x(y) := \frac{\varphi(y) - \varphi(x)}{y - x} \quad \text{for } y \in I \setminus \{x\}.$$

Then g_x is monotone increasing and there exist the left-sided and right-sided derivatives

$$D^-\varphi(x) := \lim_{y \uparrow x} g_x(y) = \sup \{g_x(y) : y < x\}$$

and

$$D^+\varphi(x) := \lim_{y \downarrow x} g_x(y) = \inf \{g_x(y) : y > x\}$$

- (iii) For $x \in I^\circ$, we have $D^-\varphi(x) \leq D^+\varphi(x)$ and

$$\varphi(x) + (y - x)t \leq \varphi(y) \quad \text{for any } y \in I \iff t \in [D^-\varphi(x), D^+\varphi(x)].$$

Hence $D^-\varphi(x)$ and $D^+\varphi(x)$ are the minimal and maximal slopes of a tangent at x .

- (iv) The maps $x \mapsto D^-\varphi(x)$ and $x \mapsto D^+\varphi(x)$ are monotone increasing. $x \mapsto D^-\varphi(x)$ is left continuous and $x \mapsto D^+\varphi(x)$ is right continuous. We have $D^-\varphi(x) = D^+\varphi(x)$ at all points of continuity of $D^-\varphi$ and $D^+\varphi$.

- (v) φ is differentiable at x if and only if $D^-\varphi(x) = D^+\varphi(x)$. In this case, the derivative is $\varphi'(x) = D^+\varphi(x)$.

- (vi) φ is almost everywhere differentiable and $\varphi(b) - \varphi(a) = \int_a^b D^+\varphi(x) dx$ for $a, b \in I^\circ$.

Proof (i) Let $x \in I^\circ$. Assume that $\liminf_{n \rightarrow \infty} \varphi(x - 1/n) \leq \varphi(x) - \varepsilon$ for some $\varepsilon > 0$. Since φ is convex, we have

$$\varphi(y) \geq \varphi(x) + n(y - x)(\varphi(x) - \varphi(x - 1/n)) \quad \text{for all } y > x \text{ and } n \in \mathbb{N}.$$

Combining this with the assumption, we get $\varphi(y) = \infty$ for all $y > x$. Hence the assumption was false. A similar argument for the right-hand side yields continuity of φ at x .

- (ii) Monotonicity is implied by convexity. The other claims are evident.

- (iii) By monotonicity of g_x , we have $D^-\varphi(x) \leq D^+\varphi(x)$. By construction, $\varphi(x) + (y - x)t \leq \varphi(y)$ for all $y < x$ if and only if $t \geq D^-\varphi(x)$. The inequality holds for all $y > x$ if and only if $t \leq D^+\varphi(x)$.

- (iv) For $\varepsilon > 0$, by the convexity, the map $x \mapsto g_x(x + \varepsilon)$ is monotone increasing and is continuous by (i). Being an infimum of monotone increasing and continuous functions the map $x \mapsto D^+\varphi(x)$ is monotone increasing and right continuous. The statement for $D^-\varphi$ follows similarly. As $x \mapsto g_x(y)$ is monotone, we get $D^+\varphi(x') \geq D^-\varphi(x') \geq D^+\varphi(x)$ for $x' > x$. If $D^+\varphi$ is continuous at x , then $D^-\varphi(x) = D^+\varphi(x)$.

- (v) This is obvious since $D^-\varphi$ and $D^+\varphi$ are the limits of the sequences of slopes of the left-sided and right-sided secant lines, respectively.

(vi) For $\varepsilon > 0$, let $A_\varepsilon = \{x \in I : D^+\varphi(x) \geq \varepsilon + \lim_{y \uparrow x} D^+\varphi(y)\}$ be the set of points of discontinuity of size at least ε . For any two points $a, b \in I$ with $a < b$, we have $\#(A_\varepsilon \cap (a, b)) \leq \varepsilon^{-1}(D^+\varphi(b) - D^+\varphi(a))$; hence $A_\varepsilon \cap (a, b)$ is a finite set. Thus A_ε is countable. Hence also $A = \bigcup_{n=1}^\infty A_{1/n}$ is countable and thus a null set. By (iv) and (v), φ is differentiable in $I^\circ \setminus A$ with derivative $D^+\varphi$.

If I is an interval, then a map $g : I \rightarrow \mathbb{R}$ is called affine linear if there are numbers $a, b \in \mathbb{R}$ such that $g(x) = ax + b$ for all $x \in I$. If $\varphi : I \rightarrow \mathbb{R}$ is a map, then we write

$$L(\varphi) := \{g : I \rightarrow \mathbb{R} \text{ is affine linear and } g \leq \varphi\}.$$

As a shorthand, we write $\sup L(\varphi)$ for the map $x \mapsto \sup\{f(x) : f \in L(\varphi)\}$.

Corollary 7.8 Let $I \subset \mathbb{R}$ be an open interval and let $\varphi : I \rightarrow \mathbb{R}$ be a map. Then the following are equivalent.

- (i) φ is convex.
- (ii) For any $x_0 \in I$, there exists a $g \in L(\varphi)$ with $g(x_0) = \varphi(x_0)$.
- (iii) $L(\varphi)$ is nonempty and $\varphi = \sup L(\varphi)$.
- (iv) There is a sequence $(g_n)_{n \in \mathbb{N}}$ in $L(\varphi)$ with $\varphi = \lim_{n \rightarrow \infty} \max\{g_1, \dots, g_n\}$.

Proof "(ii) $\implies$ (iii) $\implies$ (iv)" This is obvious.

"(iii) $\implies$ (i)" The supremum of convex functions is convex and any affine linear map is convex. Hence $\sup L(\varphi)$ is convex if $L(\varphi) \neq \emptyset$.

"(i) $\implies$ (ii)" By Theorem 7.7(iii), for any $x_0 \in I$, the map

$$x \mapsto \varphi(x_0) + (x - x_0) D^+\varphi(x_0)$$

is in $L(\varphi)$.

Theorem 7.9 (Jensen's inequality) Let $I \subset \mathbb{R}$ be an interval and let X be an I valued random variable with $\mathbf{E}[|X|] < \infty$. If φ is convex, then $\mathbf{E}[\varphi(X)^-] < \infty$ and

$$\mathbf{E}[\varphi(X)] \geq \varphi(\mathbf{E}[X])$$

Proof As $L(\varphi) \neq \emptyset$ by Corollary 7.8(iii), we can choose numbers $a, b \in \mathbb{R}$ such that $ax + b \leq \varphi(x)$ for all $x \in I$. Hence

$$\mathbf{E}[\varphi(X)^-] \leq \mathbf{E}[(aX + b)^-] \leq |b| + |a| \cdot \mathbf{E}[|X|] < \infty.$$

We distinguish the cases where $\mathbf{E}[X]$ is in the interior I° or at the boundary ∂I .

Case 1. If $\mathbf{E}[X] \in I^\circ$, then let $t^+ := D^+\varphi(\mathbf{E}[X])$ be the maximal slope of a tangent of φ at $\mathbf{E}[X]$. Then $\varphi(x) \geq t^+ \cdot (x - \mathbf{E}[X]) + \varphi(\mathbf{E}[X])$ for all $x \in I$; hence

$$\mathbf{E}[\varphi(X)] \geq t^+ \mathbf{E}[X - \mathbf{E}[X]] + \mathbf{E}[\varphi(\mathbf{E}[X])] = \varphi(\mathbf{E}[X]).$$

Case 2. If $\mathbf{E}[X] \in \partial I$, then $X = \mathbf{E}[X]$ a.s.; hence $\mathbf{E}[\varphi(X)] = \mathbf{E}[\varphi(\mathbf{E}[X])] = \varphi(\mathbf{E}[X])$.

Jensen's inequality can be extended to $\mathbb{R}^n$. To this end, we need a representation of convex functions of many variables as a supremum of affine linear functions. Recall that a function $g : \mathbb{R}^n \rightarrow \mathbb{R}$ is called affine linear if there is an $a \in \mathbb{R}^n$ and a $b \in \mathbb{R}$ such that $g(x) = \langle a, x \rangle + b$ for all x . Here $\langle \cdot, \cdot \rangle$ denotes the usual inner product on $\mathbb{R}^n$.

Theorem 7.10 Let $G \subset \mathbb{R}^n$ be open and convex and let $\varphi : G \rightarrow \mathbb{R}$ be a map. Then Corollary 7.8 holds with I replaced by G . If φ is convex, then φ is continuous and hence measurable. If φ is twice continuously differentiable, then φ is convex if and only if the Hessian matrix is positive semidefinite.

Proof As we need these statements only in the proof of the multidimensional Jensen inequality, which will not play a central role in the following, we only give references for the proofs. In Rockafellar's book [145], continuity follows from Theorem 10.1, and the statements of Corollary

7.8 follow from Theorem 12.1 and Theorem 18.8. The claim about the Hessian matrix can be found in Theorem 4.5.

Theorem 7.11 (Jensen's inequality in $\mathbb{R}^n$) Let $G \subset \mathbb{R}^n$ be a convex set and let $X_1, \dots, X_n$ be integrable real random variables with $\mathbf{P}[(X_1, \dots, X_n) \in G] = 1$. Further, let $\varphi : G \rightarrow \mathbb{R}$ be convex. Then $\mathbf{E}[\varphi(X_1, \dots, X_n)] < \infty$ and

$$\mathbf{E}[\varphi(X_1, \dots, X_n)] \geq \varphi(\mathbf{E}[X_1], \dots, \mathbf{E}[X_n]).$$

Proof First consider the case where G is open. Here, the argument is similar to the proof of Theorem 7.9. Let $g \in L(\varphi)$ with

$$g(\mathbf{E}[X_1], \dots, \mathbf{E}[X_n]) = \varphi(\mathbf{E}[X_1], \dots, \mathbf{E}[X_n]).$$

As $g \leq \varphi$ is linear, we get

$$\mathbf{E}[\varphi(X_1, \dots, X_n)] \geq \mathbf{E}[g(X_1, \dots, X_n)] = g(\mathbf{E}[X_1], \dots, \mathbf{E}[X_n]).$$

Integrability of $\varphi(X_1, \dots, X_n)$ can be derived in a similar way to the onedimensional case.

Now consider the general case where G is not necessarily open. Here the problem that arises when $(\mathbf{E}[X_1], \dots, \mathbf{E}[X_n]) \in \partial G$ is a bit more tricky than in the onedimensional case since ∂G can have flat pieces that in turn, however, are convex. Hence one cannot infer that $(X_1, \dots, X_n)$ equals its expectation almost surely. We only sketch the argument. First infer that $(X_1, \dots, X_n)$ is almost surely in one of those flat pieces. This piece is necessarily of dimension smaller than n . Now restrict φ to that flat piece and inductively reduce its dimension until reaching a point, the case that has already been treated above. Details can be found, e.g., in [37, Theorem 10.2.6].

Example 7.12 Let X be a real random variable with $\mathbf{E}[X^2] < \infty$, $I = \mathbb{R}$ and $\varphi(x) = x^2$. By Jensen's inequality, we get

$$\text{Var}[X] = \mathbf{E}[X^2] - (\mathbf{E}[X])^2 \geq 0.$$

Example 7.13 Let $G = [0, \infty) \times [0, \infty)$, $\alpha \in (0, 1)$ and $\varphi(x, y) = x^\alpha y^{1-\alpha}$. Then φ is concave (exercise!); hence, for nonnegative random variables X and Y with finite expectation (by Theorem 7.11),

$$\mathbf{E}[X^\alpha Y^{1-\alpha}] \leq (\mathbf{E}[X])^\alpha (\mathbf{E}[Y])^{1-\alpha}$$

Example 7.14 Let G, X and Y be as in Example 7.13. Let $p \in (1, \infty)$. Then $\psi(x, y) = (x^{1/p} + y^{1/p})^p$ is concave. Hence (by Theorem 7.11)

$$(\mathbf{E}[X]^{1/p} + \mathbf{E}[Y]^{1/p})^p \geq \mathbf{E}[(X^{1/p} + Y^{1/p})^p].$$

Before we present Hölder's inequality and Minkowski's inequality, we need a preparatory lemma.

Lemma 7.15 (Young's inequality) For $p, q \in (1, \infty)$ with $\frac{1}{p} + \frac{1}{q} = 1$ and for $x, y \in [0, \infty)$,

$$xy \leq \frac{x^p}{p} + \frac{y^q}{q}. \quad (7.1)$$

Proof Fix $y \in [0, \infty)$ and define $f(x) := \frac{x^p}{p} + \frac{y^q}{q} - xy$ for $x \in [0, \infty)$. f is twice continuously differentiable in $(0, \infty)$ with derivatives $f'(x) = x^{p-1} - y$ and $f''(x) = (p-1)x^{p-2}$. In particular, f is strictly convex and hence assumes its (unique) minimum at $x_0 = y^{1/(p-1)}$. By assumption, $q = \frac{p}{p-1}$; hence $x_0^p = y^q$ and thus

$$f(x_0) = \left(\frac{1}{p} + \frac{1}{q}\right) y^q - y^{1/(p-1)} y = 0.$$

Theorem 7.16 (Hölder's inequality) Let $p, q \in [1, \infty]$ with $\frac{1}{p} + \frac{1}{q} = 1$ and $f \in \mathcal{L}^p(\mu)$, $g \in \mathcal{L}^q(\mu)$. Then $(fg) \in \mathcal{L}^1(\mu)$ and

$$\|fg\|_1 \leq \|f\|_p \cdot \|g\|_q.$$

Proof The cases $p = 1$ and $p = \infty$ are trivial. Hence, let $p \in (1, \infty)$. Let $f \in \mathcal{L}^p(\mu)$ and $g \in \mathcal{L}^q(\mu)$ be nontrivial. By passing to $f/\|f\|_p$ and $g/\|g\|_q$, we may assume that $\|f\|_p = \|g\|_q = 1$. By Lemma 7.15, we have

$$\begin{aligned} \|fg\|_1 &= \int |f| \cdot |g| d\mu \leq \frac{1}{p} \int |f|^p d\mu + \frac{1}{q} \int |g|^q d\mu \\ &= \frac{1}{p} + \frac{1}{q} = 1 = \|f\|_p \cdot \|g\|_q \end{aligned}$$

Theorem 7.17 (Minkowski's inequality) For $p \in [1, \infty]$ and $f, g \in \mathcal{L}^p(\mu)$,

$$\|f + g\|_p \leq \|f\|_p + \|g\|_p \quad (7.2)$$

Proof The case $p = \infty$ is trivial. Hence, let $p \in [1, \infty)$. The left-hand side in (7.2) does not decrease if we replace f and g by $|f|$ and $|g|$. Hence we may assume $f \geq 0$ and $g \geq 0$ and (to avoid trivialities) $\|f + g\|_p > 0$.

Now $(f + g)^p \leq 2^p (f^p \vee g^p) \leq 2^p (f^p + g^p)$; hence $f + g \in \mathcal{L}^p(\mu)$. Apply Hölder's inequality to $f \cdot (f + g)^{p-1}$ and to $g \cdot (f + g)^{p-1}$ to get

$$\begin{aligned} \|f + g\|_p^p &= \int (f + g)^p d\mu = \int f(f + g)^{p-1} d\mu + \int g(f + g)^{p-1} d\mu \\ &\leq \|f\|_p \cdot \|(f + g)^{p-1}\|_q + \|g\|_p \cdot \|(f + g)^{p-1}\|_q \\ &= (\|f\|_p + \|g\|_p) \cdot \|f + g\|_p^{p-1} \end{aligned}$$

Note that in the last step, we used the fact that $p - p/q = 1$. Dividing both sides by $\|f + g\|_p^{p-1}$ yields (7.2).

In Theorem 7.17, we verified the triangle inequality and hence that $\|\cdot\|_p$ is a norm. Theorem 7.3 says that this norm is complete (i.e., every Cauchy sequence converges). A complete normed vector space is called a Banach space. Summing up, we have shown the following theorem.

Theorem 7.18 (Fischer-Riesz) $(\mathcal{L}^p(\mu), \|\cdot\|_p)$ is a Banach space for every $p \in [1, \infty]$.

Exercise 7.2.1 Show Hölder's inequality by applying Jensen's inequality to the function of Example 7.13.

Exercise 7.2.2 Show Minkowski's inequality by applying Jensen's inequality to the function of Example 7.14.

Exercise 7.2.3 Let X be a real random variable and let $p, q \in (1, \infty)$ with $\frac{1}{p} + \frac{1}{q} = 1$. Show that X is in $\mathcal{L}^p(\mathbf{P})$ if and only if there exists a $C < \infty$ such that $|\mathbf{E}[XY]| \leq C\|Y\|_q$ for any bounded random variable Y .

5.7.3 7.3 Hilbert Spaces

In this section, we study the case $p = 2$ in more detail. The main goal is the representation theorem for continuous linear functionals on Hilbert spaces due to Riesz and Fréchet. This theorem is a cornerstone for a functional analytic proof of the Radon-Nikodym theorem in Section 7.4.

Definition 7.19 Let V be a real vector space. A map $\langle \cdot, \cdot \rangle : V \times V \rightarrow \mathbb{R}$ is called an inner product if:

- (i) (Linearity) $\langle x, \alpha y + z \rangle = \alpha \langle x, y \rangle + \langle x, z \rangle$ for all $x, y, z \in V$ and $\alpha \in \mathbb{R}$.
- (ii) (Symmetry) $\langle x, y \rangle = \langle y, x \rangle$ for all $x, y \in V$.
- (iii) (Positive definiteness) $\langle x, x \rangle > 0$ for all $x \in V \setminus \{0\}$.

If only (i) and (ii) hold and $\langle x, x \rangle \geq 0$ for all x , then $\langle \cdot, \cdot \rangle$ is called a positive semidefinite symmetric bilinear form, or a semi-inner product.

If $\langle \cdot, \cdot \rangle$ is an inner product, then $(V, \langle \cdot, \cdot \rangle)$ is called a (real) Hilbert space if the norm defined by $\|x\| := \langle x, x \rangle^{1/2}$ is complete; that is, if $(V, \|\cdot\|)$ is a Banach space.

Definition 7.20 For $f, g \in \mathcal{L}^2(\mu)$, define

$$\langle f, g \rangle := \int f g d\mu$$

For $\bar{f}, \bar{g} \in L^2(\mu)$, define $\langle \bar{f}, \bar{g} \rangle := \langle f, g \rangle$, where $f \in \bar{f}$ and $g \in \bar{g}$.

Note that this definition is independent of the particular choices of the representatives of f and g .

Theorem 7.21 $\langle \cdot, \cdot \rangle$ is an inner product on $L^2(\mu)$ and a semi-inner product on $\mathcal{L}^2(\mu)$. In addition, $\|f\|_2 = \langle f, f \rangle^{1/2}$.

Proof This is left as an exercise.

As a corollary to Theorem 7.18, we get the following.

Corollary 7.22 $(L^2(\mu), \langle \cdot, \cdot \rangle)$ is a real Hilbert space.

Lemma 7.23 If $\langle \cdot, \cdot \rangle$ is a semi-inner product on the real vector space V , then $\langle \cdot, \cdot \rangle : V \times V \rightarrow \mathbb{R}$ is continuous (with respect to the product topology of the topology on V that is generated by the pseudo-metric $d(x, y) = \langle x - y, x - y \rangle^{1/2}$).

Proof This is obvious.

Definition 7.24 (Orthogonal complement) Let V be a real vector space with inner product $\langle \cdot, \cdot \rangle$. If $W \subset V$, then the orthogonal complement of W is the following linear subspace of V :

$$W^\perp := \{v \in V : \langle v, w \rangle = 0 \text{ for all } w \in W\}.$$

Theorem 7.25 (Orthogonal decomposition) Let $(V, \langle \cdot, \cdot \rangle)$ be a Hilbert space and let $W \subset V$ be a closed linear subspace. For any $x \in V$, there is a unique representation $x = y + z$ where $y \in W$ and $z \in W^\perp$.

Proof Let $x \in V$ and $c := \inf\{\|x - w\| : w \in W\}$. Further, let $(w_n)_{n \in \mathbb{N}}$ be a sequence in W with $\|x - w_n\| \xrightarrow{n \rightarrow \infty} c$. The parallelogram law yields

$$\|w_m - w_n\|^2 = 2\|w_m - x\|^2 + 2\|w_n - x\|^2 - 4\left\|\frac{1}{2}(w_m + w_n) - x\right\|^2.$$

As W is linear, we have $(w_m + w_n)/2 \in W$; hence $\left\|\frac{1}{2}(w_m + w_n) - x\right\| \geq c$. Thus $(w_n)_{n \in \mathbb{N}}$ is a Cauchy sequence: $\|w_m - w_n\| \rightarrow 0$ if $m, n \rightarrow \infty$.

Since V is complete and W is closed, W is also complete; hence there is a $y \in W$ with $w_n \xrightarrow{n \rightarrow \infty} y$. Now let $z := x - y$. Then $\|z\| = \lim_{n \rightarrow \infty} \|w_n - x\| = c$ by continuity of the norm (Lemma 7.23).

Consider an arbitrary $w \in W \setminus \{0\}$. We define $\varrho := -\langle z, w \rangle / \|w\|^2$ and get $y + \varrho w \in W$; hence

$$c^2 \leq \|x - (y + \varrho w)\|^2 = \|z\|^2 + \varrho^2 \|w\|^2 + 2\varrho \langle z, w \rangle = c^2 - \varrho^2 \|w\|^2.$$

Concluding, we have $\langle z, w \rangle = 0$ for all $w \in W$ and thus $z \in W^\perp$.

Uniqueness of the decomposition is easy: If $x = y' + z'$ is an orthogonal decomposition, then $y - y' \in W$ and $z - z' \in W^\perp$ as well as $y - y' + z - z' = 0$; hence

$$\begin{aligned}
0 &= \|y - y' + z - z'\|^2 \\
&= \|y - y'\|^2 + \|z - z'\|^2 + 2\langle y - y', z - z' \rangle \\
&= \|y - y'\|^2 + \|z - z'\|^2
\end{aligned}$$

whence $y = y'$ and $z = z'$.

Theorem 7.26 (Riesz-Fréchet representation theorem) Let $(V, \langle \cdot, \cdot \rangle)$ be a Hilbert space and let $F : V \rightarrow \mathbb{R}$ be a map. Then the following are equivalent.

- (i) F is continuous and linear.
- (ii) There is an $f \in V$ with $F(x) = \langle x, f \rangle$ for all $x \in V$.

The element $f \in V$ in (ii) is uniquely determined.

Proof "(ii) $\implies$ (i)" For any $f \in V$, by definition of the inner product, the map $x \mapsto \langle x, f \rangle$ is linear. By Lemma 7.23, this map is also continuous.

"(i) $\implies$ (ii)" If $F \equiv 0$, then choose $f = 0$. Now assume F is not identically zero. As F is continuous, the kernel $W := F^{-1}(\{0\})$ is a closed (proper) linear subspace of V . Let $v \in V \setminus W$ and let $v = y + z$ for $y \in W$ and $z \in W^\perp$ be the orthogonal decomposition of v . Then $z \neq 0$ and $F(z) = F(v) - F(y) = F(v) \neq 0$. Hence we can define $u := z/F(z) \in W^\perp$. Clearly, $F(u) = 1$ and for any $x \in V$, we have $F(x - F(x)u) = F(x) - F(x)F(u) = 0$; hence $x - F(x)u \in W$ and thus $\langle x - F(x)u, u \rangle = 0$. Consequently, $F(x) = \langle x, u \rangle / \|u\|^2$. Now define $f := u / \|u\|^2$. Then $F(x) = \langle x, f \rangle$ for all $x \in V$.

"Uniqueness" Let $\langle x, f \rangle = \langle x, g \rangle$ for all $x \in V$. Letting $x = f - g$, we get $0 = \langle f - g, f - g \rangle$; hence $f = g$.

In the following section, we will need the representation theorem for the space $\mathcal{L}^2(\mu)$, which, unlike $L^2(\mu)$, is not a Hilbert space. However, with a little bit of abstract nonsense, one can apply the preceding theorem to $\mathcal{L}^2(\mu)$. Recall that $\mathcal{N} = \{f \in \mathcal{L}^2(\mu) : \langle f, f \rangle = 0\}$ is the subspace of functions that equal zero almost everywhere. Let $L^2(\mu) = \mathcal{L}^2(\mu) / \mathcal{N}$ be the factor space. This is a special case of the situation where $(V, \langle \cdot, \cdot \rangle)$ is a linear space with complete semi-inner product. In this case, $\mathcal{N} := \{v \in V : \langle v, v \rangle = 0\}$ and $V_0 = V / \mathcal{N} := \{f + \mathcal{N} : f \in V\}$. Denote $\langle v + \mathcal{N}, w + \mathcal{N} \rangle_0 := \langle v, w \rangle$ to obtain a Hilbert space $(V_0, \langle \cdot, \cdot \rangle_0)$.

Corollary 7.27 Let $(V, \langle \cdot, \cdot \rangle)$ be a linear vector space with complete semi-inner product. The map $F : V \rightarrow \mathbb{R}$ is continuous and linear if and only if there is an $f \in V$ with $F(x) = \langle x, f \rangle$ for all $x \in V$.

Proof One implication is trivial. Hence, let F be continuous and linear. Then $F(0) = 0$ since F is linear. Note that $F(v) = F(0) = 0$ for all $v \in \mathcal{N}$ since F is continuous. Indeed, v lies in every open neighborhood of 0; hence F assumes at v the same value as at 0. Thus F induces a continuous linear map $F_0 : V_0 \rightarrow \mathbb{R}$ by $F_0(x + \mathcal{N}) = F(x)$. By Theorem 7.26, there is an $f + \mathcal{N} \in V_0$ with $F_0(x + \mathcal{N}) = \langle x + \mathcal{N}, f + \mathcal{N} \rangle_0$ for all $x + \mathcal{N} \in V_0$. However, $F(x) = \langle x, f \rangle$ for all $x \in V$ by the definition of F_0 and $\langle \cdot, \cdot \rangle_0$.

Corollary 7.28 The map $F : \mathcal{L}^2(\mu) \rightarrow \mathbb{R}$ is continuous and linear if and only if there is an $f \in \mathcal{L}^2(\mu)$ with $F(g) = \int g f d\mu$ for all $g \in \mathcal{L}^2(\mu)$.

Proof The space $\mathcal{L}^2(\mu)$ fulfills the conditions of Corollary 7.27.

Exercise 7.3.1 (Fourier series) For $n \in \mathbb{N}_0$, define $S_n, C_n : [0, 1] \rightarrow [0, 1]$ by $S_n(x) = \sqrt{2} \sin(2\pi n x)$, $C_n(x) = \sqrt{2} \cos(2\pi n x)$. For two square summable sequences $(a_n)_{n \in \mathbb{N}}$ and $(b_n)_{n \in \mathbb{N}_0}$, let $h_{a,b} := b_0 + \sum_{n=1}^{\infty} (a_n S_n + b_n C_n)$. Further, let W be the vector space of such $h_{a,b}$.

Show the following:

- (i) The functions $C_0, S_n, C_n, n \in \mathbb{N}$ form an orthogonal system in $L^2([0, 1], \lambda)$.
- (ii) The series defining $h_{a,b}$ converges in $L^2([0, 1], \lambda)$.
- (iii) W is a closed linear subspace of $L^2([0, 1], \lambda)$.
- (iv) $W = L^2([0, 1], \lambda)$. More precisely, for any $f \in L^2([0, 1], \lambda)$, there exist uniquely defined

square summable sequences $(a_n)_{n \in \mathbb{N}}$ and $(b_n)_{n \in \mathbb{N}_0}$ such that $f = h_{a,b}$. Furthermore, $\|f\|_2^2 = b_0^2 + \sum_{n=1}^{\infty} (a_n^2 + b_n^2)$.

Hint: Show (iv) first for step functions (see Exercise 4.2.6).

5.7.4 7.4 Lebesgue's Decomposition Theorem

In this section, we employ the properties of Hilbert spaces that we derived in the last section in order to decompose a measure into a singular part and a part that is absolutely continuous, both with respect to a second given measure. Furthermore, we show that the absolutely continuous part has a density. Let μ and ν be measures on $(\Omega, \mathcal{A})$. By Definition 4.13, a measurable function $f : \Omega \rightarrow [0, \infty)$ is called a density of ν with respect to μ if

$$v(A) := \int f \mathbb{1}_A d\mu \quad \text{for all } A \in \mathcal{A} \quad (7.3)$$

On the other hand, for any measurable $f : \Omega \rightarrow [0, \infty)$, Eq. (7.3) defines a measure ν on $(\Omega, \mathcal{A})$. In this case, we also write

$$v = f\mu \quad \text{and} \quad f = \frac{dv}{d\mu} \quad (7.4)$$

For example, the normal distribution $\nu = \mathcal{N}_{0,1}$ has the density $f(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$ with respect to the Lebesgue measure $\mu = \lambda$ on $\mathbb{R}$.

If $g : \Omega \rightarrow [0, \infty]$ is measurable, then (by Theorem 4.15)

$$\int g d\nu = \int g f d\mu \quad (7.5)$$

Hence $g \in \mathcal{L}^1(\nu)$ if and only if $gf \in \mathcal{L}^1(\mu)$, and in this case (7.5) holds.

If $\nu = f\mu$, then $\nu(A) = 0$ for all $A \in \mathcal{A}$ with $\mu(A) = 0$. The situation is quite the opposite for, e.g., the Poisson distribution $\mu = \text{Poi}_\varrho$ with parameter $\varrho > 0$ and $\nu = \mathcal{N}_{0,1}$. Here $\mathbb{N}_0 \subset \mathbb{R}$ is a ν -null set with $\mu(\mathbb{R} \setminus \mathbb{N}_0) = 0$. We say that ν is singular to μ .

The main goal of this chapter is to show that an arbitrary σ -finite measure ν on a measurable space $(\Omega, \mathcal{A})$ can be decomposed into a part that is singular to the σ -finite measure μ and a part that has a density with respect to μ (Lebesgue's decomposition theorem, Theorem 7.33).

Theorem 7.29 (Uniqueness of the density) Let ν be σ -finite. If f_1 and f_2 are densities of ν with respect to μ , then $f_1 = f_2\mu$ -almost everywhere. In particular, the density $\frac{d\nu}{d\mu}$ is unique up to equality μ -almost everywhere.

Proof Let $E_n \uparrow \Omega$ with $\nu(E_n) < \infty, n \in \mathbb{N}$. Let $A_n = E_n \cap \{f_1 > f_2\}$ for $n \in \mathbb{N}$. Then $\nu(A_n) < \infty$; hence

$$0 = \nu(A_n) - \nu(A_n) = \int_{A_n} (f_1 - f_2) d\mu$$

By Theorem 4.8(i), $f_2 \mathbb{1}_{A_n} = f_1 \mathbb{1}_{A_n} \mu$ -a.e. As $f_1 > f_2$ on A_n , we infer $\mu(A_n) = 0$ and

$$\mu(\{f_1 > f_2\}) = \mu\left(\bigcup_{n \in \mathbb{N}} A_n\right) = 0.$$

Similarly, we get $\mu(\{f_1 < f_2\}) = 0$; hence $f_1 = f_2\mu$ -a.e. $\square$

Definition 7.30 Let μ and ν be two measures on $(\Omega, \mathcal{A})$.

(i) ν is called absolutely continuous with respect to μ (symbolically $\nu \ll \mu$) if

$$v(A) = 0 \quad \text{for all } A \in \mathcal{A} \text{ with } \mu(A) = 0. \quad (7.6)$$

The measures μ and ν are called equivalent (symbolically $\mu \approx \nu$) if $\nu \ll \mu$ and $\mu \ll \nu$.

(ii) μ is called singular to ν (symbolically $\mu \perp \nu$) if there exists an $A \in \mathcal{A}$ such that $\mu(A) = 0$ and $\nu(\Omega \setminus A) = 0$.

Remark 7.31 Clearly, $\mu \perp \nu \iff \nu \perp \mu$. $\square$

Example 7.32

(i) Let μ be a measure on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ with density f with respect to the Lebesgue measure λ . Then $\mu(A) = \int_A f d\lambda = 0$ for every $A \in \mathcal{A}$ with $\lambda(A) = 0$; hence $\mu \ll \lambda$. If λ -almost everywhere $f > 0$, then $\mu(A) = \int_A f d\lambda > 0$ if $\lambda(A) > 0$; hence $\mu \approx \lambda$. If $\lambda(\{f = 0\}) > 0$, then (since $\mu(\{f = 0\}) = 0$) $\lambda \not\ll \mu$.

(ii) Consider the Bernoulli distributions Ber_p and Ber_q for $p, q \in [0, 1]$. If $p \in (0, 1)$, then $\text{Ber}_q \ll \text{Ber}_p$. If $p \in \{0, 1\}$, then $\text{Ber}_q \ll \text{Ber}_p$ if and only if $p = q$, and $\text{Ber}_q \perp \text{Ber}_p$ if and only if $q = 1 - p$.

(iii) Consider the Poisson distributions Poi_α and Poi_β for $\alpha, \beta \geq 0$. We have $\text{Poi}_\alpha \ll \text{Poi}_\beta$ if and only if $\beta > 0$ or $\alpha = 0$.

(iv) Consider the infinite product measures (see Theorem 1.64) $(\text{Ber}_p)^{\otimes \mathbb{N}}$ and $(\text{Ber}_q)^{\otimes \mathbb{N}}$ on $\Omega = \{0, 1\}^{\mathbb{N}}$. Then $(\text{Ber}_p)^{\otimes \mathbb{N}} \perp (\text{Ber}_q)^{\otimes \mathbb{N}}$ if $p \neq q$. Indeed, for $n \in \mathbb{N}$, let $X_n((\omega_1, \omega_2, \dots)) = \omega_n$ be the projection of Ω to the n th coordinate. Then under $(\text{Ber}_r)^{\otimes \mathbb{N}}$ the family $(X_n)_{n \in \mathbb{N}}$ is independent and Bernoulli distributed with parameter r (see Example 2.18). By the strong law of large

numbers, for any $r \in \{p, q\}$, there exists a measurable set $A_r \subset \Omega$ with $(\text{Ber}_r)^{\otimes \mathbb{N}}(\Omega \setminus A_r) = 0$ and

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n X_i(\omega) = r \quad \text{for all } \omega \in A_r.$$

In particular, $A_p \cap A_q = \emptyset$ if $p \neq q$ and thus $(\text{Ber}_p)^{\otimes \mathbb{N}} \perp (\text{Ber}_q)^{\otimes \mathbb{N}}$.

Theorem 7.33 (Lebesgue's decomposition theorem) Let μ and ν be σ -finite measures on $(\Omega, \mathcal{A})$. Then ν can be uniquely decomposed into an absolutely continuous part ν_a and a singular part ν_s (with respect to μ):

$$\nu = \nu_a + \nu_s, \quad \text{where } \nu_a \ll \mu \text{ and } \nu_s \perp \mu.$$

ν_a has a density with respect to μ , and $\frac{d\nu_a}{d\mu}$ is $\mathcal{A}$ -measurable and finite μ -a.e.

Corollary 7.34 (Radon-Nikodym theorem) Let μ and ν be σ -finite measures on $(\Omega, \mathcal{A})$. Then

$$\nu \text{ has a density w.r.t. } \mu \iff \nu \ll \mu.$$

In this case, $\frac{d\nu}{d\mu}$ is $\mathcal{A}$ -measurable and finite μ -a.e. $\frac{d\nu}{d\mu}$ is called the Radon-Nikodym derivative of ν with respect to μ .

Proof One direction is trivial. Hence, let $\nu \ll \mu$. By Theorem 7.33, we get that $\nu = \nu_a$ has a density with respect to μ .

Proof of Theorem 7.33 The idea goes back to von Neumann. We follow the exposition in [37].

By the usual exhaustion arguments, we can restrict ourselves to the case where μ and ν are finite. By Theorem 4.19, the canonical inclusion $i : \mathcal{L}^2(\Omega, \mathcal{A}, \mu + \nu) \hookrightarrow \mathcal{L}^1(\Omega, \mathcal{A}, \mu + \nu)$ is continuous. Since $\nu \leq \mu + \nu$, the linear functional $\mathcal{L}^2(\Omega, \mathcal{A}, \mu + \nu) \rightarrow \mathbb{R}, h \mapsto \int h d\nu$ is continuous. By the Riesz-Fréchet theorem (here Corollary 7.28), there exists a $g \in \mathcal{L}^2(\Omega, \mathcal{A}, \mu + \nu)$ such that

$$\int h dv = \int h g d(\mu + v) \quad \text{for all } h \in \mathcal{L}^2(\Omega, \mathcal{A}, \mu + v) \quad (7.7)$$

or equivalently

$$\int f(1 - g) d(\mu + v) = \int f d\mu \quad \text{for all } f \in \mathcal{L}^2(\Omega, \mathcal{A}, \mu + v) \quad (7.8)$$

If in (7.7) we choose $h = \mathbb{1}_{\{g < 0\}}$, then we get that $(\mu + \nu)$ -almost everywhere $g \geq 0$. Similarly, with $f = \mathbb{1}_{\{g > 1\}}$ in (7.8), we obtain that $(\mu + \nu)$ -almost everywhere $g \leq 1$. Hence $0 \leq g \leq 1$.

Now let $f \geq 0$ be measurable and let $(f_n)_{n \in \mathbb{N}}$ be a sequence of nonnegative functions in $\mathcal{L}^2(\Omega, \mathcal{A}, \mu + \nu)$ with $f_n \uparrow f$. By the monotone convergence theorem (applied to the measure $(1 - g)(\mu + \nu)$; that is, the measure with density $(1 - g)$ with respect to $\mu + \nu$), we obtain that (7.8) holds for all measurable $f \geq 0$. Similarly, we get (7.7) for all measurable $h \geq 0$.

Let $E := g^{-1}(\{1\})$. If we let $f = \mathbb{1}_E$ in (7.8), then we get $\mu(E) = 0$. Define the measures ν_a and ν_s for $A \in \mathcal{A}$ by

$$\nu_a(A) := \nu(A \setminus E) \quad \text{and} \quad \nu_s(A) := \nu(A \cap E).$$

Clearly, $\nu = \nu_a + \nu_s$ and $\nu_s(\Omega \setminus E) = 0$; hence $\nu_s \perp \mu$. If now $A \cap E = \emptyset$ and $\mu(A) = 0$, then $\int \mathbb{1}_A d\mu = 0$. Hence, by (7.8), also $\int_A (1 - g) d(\mu + \nu) = 0$. On the other hand, we have $1 - g > 0$ on A ; hence $\mu(A) + \nu(A) = 0$ and thus $\nu_a(A) = \nu(A) = 0$. If, more generally, B is measurable with $\mu(B) = 0$, then $\mu(B \setminus E) = 0$; hence, as shown above, $\nu_a(B) = \nu_a(B \setminus E) = 0$. Consequently, $\nu_a \ll \mu$ and $\nu = \nu_a + \nu_s$ is the decomposition we wanted to construct.

In order to obtain the density of ν_a with respect to μ , we define $f := \frac{g}{1-g} \mathbb{1}_{\Omega \setminus E}$. For any $A \in \mathcal{A}$, by (7.8) and (7.7) with $h = \mathbb{1}_{A \setminus E}$,

$$\int_A f d\mu = \int_{A \cap E^c} g d(\mu + \nu) = \nu(A \setminus E) = \nu_a(A)$$

Hence $f = \frac{d\nu_a}{d\mu}$.

Exercise 7.4.1 For every $x \in (0, 1]$, let $x = (0, x_1 x_2 x_3 \dots) := \sum_{n=1}^{\infty} x_n 2^{-n}$ be the dyadic expansion (with $\limsup_{n \rightarrow \infty} x_n = 1$ for definiteness). Define a map $F : (0, 1] \rightarrow (0, 1]$ by

$$F(x) = (0, x_1 x_1 x_2 x_2 x_3 x_3 \dots) = \sum_{n=1}^{\infty} 3x_n 4^{-n}$$

Let U be a random variable that is uniformly distributed on $(0, 1]$ and denote by $\mu := \mathbf{P}_{U \circ F^{-1}}$ the distribution of $F(U)$.

Show that the probability measure μ has a continuous distribution function and that μ is singular to the Lebesgue measure $\lambda|_{(0,1]}$.

Exercise 7.4.2 Let $n \in \mathbb{N}$ and $p, q \in [0, 1]$. For which values of p and q do we have $b_{n,p} \ll b_{n,q}$? Compute the Radon-Nikodym derivative $\frac{db_{n,p}}{db_{n,q}}$.

5.7.5 7.5 Supplement: Signed Measures

In this section, we show the decomposition theorems for signed measures (Hahn, Jordan) and deliver an alternative proof for Lebesgue's decomposition theorem. We owe some of the proofs to [89].

Definition 7.35 Let μ and ν be two measures on $(\Omega, \mathcal{A})$. ν is called totally continuous with respect to μ if, for any $\varepsilon > 0$, there exists a $\delta > 0$ such that for all $A \in \mathcal{A}$

$$\mu(A) < \delta \quad \text{implies} \quad \nu(A) < \varepsilon. \quad (7.9)$$

Remark 7.36 The definition of total continuity is similar to that of uniform integrability (see Theorem 6.24(iii)), at least for finite μ . We will come back to this connection in the framework of the martingale convergence theorem that will provide an alternative proof of the Radon-Nikodym theorem (Corollary 7.34).

Theorem 7.37 Let μ and ν be measures on $(\Omega, \mathcal{A})$. If ν is totally continuous with respect to μ , then $\nu \ll \mu$. If $\nu(\Omega) < \infty$, then the converse also holds.

Proof " " Let ν be totally continuous with respect to μ . Let $A \in \mathcal{A}$ with $\mu(A) = 0$. For all $\varepsilon > 0$, by assumption, $\nu(A) < \varepsilon$; hence $\nu(A) = 0$ and thus $\nu \ll \mu$.

" " Let ν be finite but not totally continuous with respect to μ . Then there exist an $\varepsilon > 0$ and sets $A_n \in \mathcal{A}$ with $\mu(A_n) < 2^{-n}$ but $\nu(A_n) \geq \varepsilon$ for all $n \in \mathbb{N}$. Define

$$A := \limsup_{n \rightarrow \infty} A_n = \bigcap_{n=1}^{\infty} \bigcup_{k=n}^{\infty} A_k$$

Then

$$\mu(A) = \lim_{n \rightarrow \infty} \mu\left(\bigcup_{k=n}^{\infty} A_k\right) \leq \lim_{n \rightarrow \infty} \sum_{k=n}^{\infty} \mu(A_k) \leq \lim_{n \rightarrow \infty} \sum_{k=n}^{\infty} 2^{-k} = 0.$$

Since ν is finite and upper semicontinuous (Theorem 1.36), we have

$$\nu(A) = \lim_{n \rightarrow \infty} \nu\left(\bigcup_{k=n}^{\infty} A_k\right) \geq \inf_{n \in \mathbb{N}} \nu(A_n) \geq \varepsilon > 0$$

Thus $\nu \not\ll \mu$.

Example 7.38 In the converse implication of the theorem, the assumption of finiteness is essential. For example, let $\mu = \mathcal{N}_{0,1}$ be the standard normal distribution on $\mathbb{R}$ and let ν be the Lebesgue measure on $\mathbb{R}$. Then ν has the density $f(x) = \sqrt{2\pi}e^{x^2/2}$ with respect to μ . In particular, we have $\nu \ll \mu$. On the other hand, $\mu([n, \infty)) \xrightarrow{n \rightarrow \infty} 0$ and $\nu([n, \infty)) = \infty$ for any $n \in \mathbb{N}$. Hence ν is not totally continuous with respect to μ .

Example 7.39 Let $(\Omega, \mathcal{A})$ be a measurable space and let μ and ν be finite measures on $(\Omega, \mathcal{A})$. Denote by $\mathcal{Z}$ the set of finite partitions of Ω into pairwise disjoint measurable sets. That is, $Z \in \mathcal{Z}$ is a finite subset of $\mathcal{A}$ such that the sets $C \in Z$ are pairwise disjoint and $\cup_{C \in Z} C = \Omega$ for all $Z \in \mathcal{Z}$. For $Z \in \mathcal{Z}$, define a function $f_Z : \Omega \rightarrow \mathbb{R}$ by

$$f_Z(\omega) = \sum_{C \in Z: \mu(C) > 0} \frac{\nu(C)}{\mu(C)} \mathbb{1}_C(\omega)$$

We show that the following three statements are equivalent.

- (i) The family $(f_Z : Z \in \mathcal{Z})$ is uniformly integrable in $\mathcal{L}^1(\mu)$ and $\int f_Z d\mu = \nu(\Omega)$ for any $Z \in \mathcal{Z}$.
- (ii) $\nu \ll \mu$.
- (iii) ν is totally continuous with respect to μ .

The equivalence of (ii) and (iii) was established in the preceding theorem. If (ii) holds, then, for all $Z \in \mathcal{Z}$,

$$\int f_Z d\mu = \sum_{C \in Z: \mu(C) > 0} \nu(C) = \nu(\Omega)$$

since $\nu(C) = 0$ for those C that do not appear in the sum. Now fix $\varepsilon > 0$. Since (ii) implies (iii), there is a $\delta' > 0$ such that $\nu(A) < \varepsilon/2$ for all $A \in \mathcal{A}$ with $\mu(A) \leq \delta'$. Let $K := \nu(\Omega)/\delta'$ and $\delta < \varepsilon/(2K)$. Then

$$\mu \left(\bigcup_{C \in \mathcal{Z}: K\mu(C) \leq \nu(C)} C \right) = \sum_{C \in \mathcal{Z}: K\mu(C) \leq \nu(C)} \mu(C) \leq \frac{1}{K} \nu(\Omega) = \delta';$$

hence

$$\sum_{C \in \mathcal{Z}: K\mu(C) \leq \nu(C)} \nu(C) = \nu \left(\bigcup_{C \in \mathcal{Z}: K\mu(C) \leq \nu(C)} C \right) < \frac{\varepsilon}{2}.$$

We conclude that for all $A \in \mathcal{A}$ with $\mu(A) < \delta$,

$$\begin{aligned} \int_A f_Z d\mu &= \sum_{C \in \mathcal{Z}: \mu(C) > 0} \mu(A \cap C) \frac{\nu(C)}{\mu(C)} \\ &= \sum_{0 < K\mu(C) \leq \nu(C)} \mu(A \cap C) \frac{\nu(C)}{\mu(C)} + \sum_{K\mu(C) > \nu(C)} \mu(A \cap C) \frac{\nu(C)}{\mu(C)} \\ &\leq \frac{\varepsilon}{2} + \sum_{K\mu(C) > \nu(C)} K\mu(A \cap C) \leq \frac{\varepsilon}{2} + K\mu(A) < \varepsilon. \end{aligned}$$

Hence $(f_Z, Z \in \mathcal{Z})$ is uniformly integrable by Theorem 6.24(iii).

Now assume (i). If $\mu = 0$, then $\int f d\mu = 0$ for all f ; hence $\nu(\Omega) = 0$ and thus $\nu \ll \mu$. Hence, let $\mu \neq 0$. Let $A \in \mathcal{A}$ with $\mu(A) = 0$. Then $Z = \{A, A^c\} \in \mathcal{Z}$ and $f_Z = \mathbb{1}_{A^c} \nu(A^c) / \mu(A^c)$. By assumption, $\nu(\Omega) = \int f_Z d\mu = \nu(A^c)$; hence $\nu(A) = 0$ and thus $\nu \ll \mu$.

Definition 7.40 (Signed measure) A set function $\varphi : \mathcal{A} \rightarrow \mathbb{R}$ is called a signed measure on $(\Omega, \mathcal{A})$ if it is σ -additive; that is, if for any sequence of pairwise disjoint sets $A_1, A_2, \dots \in \mathcal{A}$,

$$\varphi \left(\biguplus_{n=1}^{\infty} A_n \right) = \sum_{n=1}^{\infty} \varphi(A_n). \quad (7.10)$$

The set of all signed measures will be denoted by $\mathcal{M}^{\pm} = \mathcal{M}^{\pm}(\Omega, \mathcal{A})$.

Remark 7.41

(i) If φ is a signed measure, then in (7.10) we automatically have absolute convergence. Indeed, the value of the left-hand side does not change if we change the order of the sets $A_1, A_2, \dots$. In order for this to hold for the right-hand side, by Weierstraß's theorem on rearrangements of series, the series has to converge absolutely. In particular, for any sequence $(A_n)_{n \in \mathbb{N}}$ of pairwise disjoint sets, we have $\lim_{n \rightarrow \infty} \sum_{k=n}^{\infty} |\varphi(A_k)| = 0$.

(ii) If $\varphi \in \mathcal{M}^{\pm}$, then $\varphi(\emptyset) = 0$ since $\mathbb{R} \ni \nu(\emptyset) = \sum_{n \in \mathbb{N}} \nu(\emptyset)$.

(iii) In general, $\varphi \in \mathcal{M}^{\pm}$ is not σ -subadditive.

Example 7.42 If μ^+, μ^- are finite measures, then $\varphi := \mu^+ - \mu^- \in \mathcal{M}^{\pm}$. We will see that every signed measure has such a representation.

Theorem 7.43 (Hahn's decomposition theorem) Let φ be a signed measure. Then there is a set $\Omega^+ \in \mathcal{A}$ with $\varphi(A) \geq 0$ for all $A \in \mathcal{A}, A \subset \Omega^+$ and $\varphi(A) \leq 0$ for all $A \in \mathcal{A}, A \subset \Omega^- := \Omega \setminus \Omega^+$. Such a decomposition $\Omega = \Omega^- \uplus \Omega^+$ is called a Hahn decomposition of Ω (with respect to φ).

Proof Let $\alpha := \sup\{\varphi(A) : A \in \mathcal{A}\}$. We have to show that φ attains the maximum α ; that is, there exists an $\Omega^+ \in \mathcal{A}$ with $\varphi(\Omega^+) = \alpha$. If this is the case, then $\alpha \in \mathbb{R}$ and for $A \subset \Omega^+, A \in \mathcal{A}$, we would have

$$\alpha \geq \varphi(\Omega^+ \setminus A) = \varphi(\Omega^+) - \varphi(A) = \alpha - \varphi(A);$$

hence $\varphi(A) \geq 0$. For $A \subset \Omega^-$, $A \in \mathcal{A}$, we would have $\varphi(A) \leq 0$ since

$$\alpha \geq \varphi(\Omega^+ \cup A) = \varphi(\Omega^+) + \varphi(A) = \alpha + \varphi(A).$$

We now construct Ω^+ with $\varphi(\Omega^+) = \alpha$. Let $(A_n)_{n \in \mathbb{N}}$ be a sequence in $\mathcal{A}$ with $\alpha = \lim_{n \rightarrow \infty} \varphi(A_n)$. Let $A := \bigcup_{n=1}^{\infty} A_n$. As each A_n could still contain "portions with negative mass", we cannot simply choose $\Omega^+ = A$. Rather, we have to peel off the negative portions layer by layer.

Define $A_n^0 := A_n$, $A_n^1 := A \setminus A_n$, and let

$$\mathcal{P}_n := \left\{ \bigcap_{i=1}^n A_i^{s(i)} : s \in \{0, 1\}^n \right\}$$

be the partition of A that is generated by $A_1, \dots, A_n$. Clearly, for any $B, C \in \mathcal{P}_n$, either $B = C$ or $B \cap C = \emptyset$ holds. In addition, we have

$$A_n = \bigsqcup_{\substack{B \in \mathcal{P}_n \\ B \subset A_n}} B$$

Define

$$\mathcal{P}_n^- := \{B \in \mathcal{P}_n : \varphi(B) < 0\}, \quad \mathcal{P}_n^+ := \mathcal{P}_n \setminus \mathcal{P}_n^-$$

and

$$C_n := \bigcup_{B \in \mathcal{P}_n^+} B$$

Due to the finite additivity of φ , we have

$$\varphi(A_n) = \sum_{\substack{B \in \mathcal{P}_n \\ B \subset A_n}} \varphi(B) \leq \sum_{\substack{B \in \mathcal{P}_n^+ \\ B \subset A_n}} \varphi(B) \leq \sum_{B \in \mathcal{P}_n^+} \varphi(B) = \varphi(C_n).$$

For $m \leq n$, let $E_m^n = C_m \cup \dots \cup C_n$. Hence, for $m < n$, we have $E_m^n \setminus E_m^{n-1} \subset C_n$ and thus

$$E_m^n \setminus E_m^{n-1} = \bigsqcup_{\substack{B \in \mathcal{P}_n^+ \\ B \subset E_m^n \setminus E_m^{n-1}}} B$$

In particular, this implies $\varphi(E_m^n \setminus E_m^{n-1}) \geq 0$. For $E_m := \bigcup_{n \geq m} C_n$, we also have $E_m^n \uparrow E_m(n \rightarrow \infty)$ and

$$\begin{aligned} \varphi(A_m) &\leq \varphi(C_m) = \varphi(E_m^m) \leq \varphi(E_m^m) + \sum_{n=m+1}^{\infty} \varphi(E_m^n \setminus E_m^{n-1}) \\ &= \varphi\left(E_m^m \cup \bigcup_{n=m+1}^{\infty} (E_m^n \setminus E_m^{n-1})\right) = \varphi\left(\bigcup_{n=m}^{\infty} E_m^n\right) = \varphi(E_m) \end{aligned}$$

Now define $\Omega^+ = \bigcap_{m=1}^{\infty} E_m$; hence $E_m \downarrow \Omega^+$. Then

$$\begin{aligned}
\varphi(E_m) &= \varphi\left(\Omega^+ \uplus \biguplus_{n \geq m} (E_n \setminus E_{n+1})\right) \\
&= \varphi(\Omega^+) + \sum_{n=m}^{\infty} \varphi(E_n \setminus E_{n+1}) \xrightarrow{m \rightarrow \infty} \varphi(\Omega^+)
\end{aligned}$$

In the last step, we used Remark 7.41(i). Summing up, we have

$$\alpha = \lim_{m \rightarrow \infty} \varphi(A_m) \leq \lim_{m \rightarrow \infty} \varphi(E_m) = \varphi(\Omega^+)$$

However, by definition, $\alpha \geq \varphi(\Omega^+)$; hence $\alpha = \varphi(\Omega^+)$. This finishes the proof.

Corollary 7.44 (Jordan's decomposition theorem) Assume $\varphi \in \mathcal{M}^\pm(\Omega, \mathcal{A})$ is a signed measure. Then there exist uniquely determined finite measures φ^+, φ^- with $\varphi = \varphi^+ - \varphi^-$ and $\varphi^+ \perp \varphi^-$.

Proof Let $\Omega = \Omega^+ \uplus \Omega^-$ be a Hahn decomposition. Define $\varphi^+(A) := \varphi(A \cap \Omega^+)$ and $\varphi^-(A) := -\varphi(A \cap \Omega^-)$.

The uniqueness of the decomposition is trivial.

Corollary 7.45 Let $\varphi \in \mathcal{M}^\pm(\Omega, \mathcal{A})$ and let $\varphi = \varphi^+ - \varphi^-$ be the Jordan decomposition of φ . Let $\Omega = \Omega^+ \uplus \Omega^-$ be a Hahn decomposition of Ω . Then

$$\begin{aligned}
\|\varphi\|_{TV} &:= \sup\{\varphi(A) - \varphi(\Omega \setminus A) : A \in \mathcal{A}\} \\
&= \varphi(\Omega^+) - \varphi(\Omega^-) \\
&= \varphi^+(\Omega) + \varphi^-(\Omega)
\end{aligned}$$

defines a norm on $\mathcal{M}^\pm(\Omega, \mathcal{A})$, the so-called total variation norm.

Proof We only have to show the triangle inequality. Let $\varphi_1, \varphi_2 \in \mathcal{M}^\pm$. Let $\Omega = \Omega^+ \uplus \Omega^-$ be a Hahn decomposition with respect to $\varphi := \varphi_1 + \varphi_2$ and let $\Omega = \Omega_i^+ \uplus \Omega_i^-$ be a Hahn decomposition with respect to $\varphi_i, i = 1, 2$. Then

$$\begin{aligned}
\|\varphi_1 + \varphi_2\|_{TV} &= \varphi_1(\Omega^+) - \varphi_1(\Omega^-) + \varphi_2(\Omega^+) - \varphi_2(\Omega^-) \\
&\leq \varphi_1(\Omega_1^+) - \varphi_1(\Omega_1^-) + \varphi_2(\Omega_2^+) - \varphi_2(\Omega_2^-) \\
&= \|\varphi_1\|_{TV} + \|\varphi_2\|_{TV}
\end{aligned}$$

With a lemma, we prepare for an alternative proof of Lebesgue's decomposition theorem (Theorem 7.33).

Lemma 7.46 Let μ, ν be finite measures on $(\Omega, \mathcal{A})$ that are not mutually singular; in short, $\mu \not\perp \nu$. Then there is an $A \in \mathcal{A}$ with $\mu(A) > 0$ and an $\varepsilon > 0$ with

$$\varepsilon\mu(E) \leq \nu(E) \quad \text{for all } E \in \mathcal{A} \text{ with } E \subset A.$$

Proof For $n \in \mathbb{N}$, let $\Omega = \Omega_n^+ \uplus \Omega_n^-$ be a Hahn decomposition for $(\nu - \frac{1}{n}\mu) \in \mathcal{M}^\pm$. Define $M := \bigcap_{n \in \mathbb{N}} \Omega_n^-$. Clearly, $(\nu - \frac{1}{n}\mu)(M) \leq 0$; hence $\nu(M) \leq \frac{1}{n}\mu(M)$ for all $n \in \mathbb{N}$ and thus $\nu(M) = 0$. Since $\mu \not\perp \nu$, we get $\mu(\Omega \setminus M) = \mu(\bigcup_{n \in \mathbb{N}} \Omega_n^+) > 0$. Thus $\mu(\Omega_{n_0}^+) > 0$ for some $n_0 \in \mathbb{N}$. Define $A := \Omega_{n_0}^+$ and $\varepsilon := \frac{1}{n_0}$. Then $\mu(A) > 0$ and $(\nu - \varepsilon\mu)(E) \geq 0$ for all $E \subset A, E \in \mathcal{A}$.

Alternative proof of Theorem 7.33 We show only the existence of a decomposition. By choosing a suitable sequence $\Omega_n \uparrow \Omega$, we can assume that ν is finite. Consider the set of functions

$$\mathcal{G} := \left\{ g : \Omega \rightarrow [0, \infty] : g \text{ is measurable and } \int_A g d\mu \leq \nu(A) \text{ for all } A \in \mathcal{A} \right\},$$

and define

$$\gamma := \sup \left\{ \int g d\mu : g \in \mathcal{G} \right\}$$

Our aim is to find a maximal element f in $\mathcal{G}$ (i.e., an f for which $\int f d\mu = \gamma$). This f will be the density of ν_a .

Clearly, $0 \in \mathcal{G}$; hence $\mathcal{G} \neq \emptyset$. Furthermore,

$$f, g \in \mathcal{G} \quad \text{implies} \quad f \vee g \in \mathcal{G}. \quad (7.11)$$

Indeed, letting $E := \{f \geq g\}$, for all $A \in \mathcal{A}$, we have

$$\int_A (f \vee g) d\mu = \int_{A \cap E} f d\mu + \int_{A \setminus E} g d\mu \leq v(A \cap E) + v(A \setminus E) = v(A)$$

Choose a sequence $(g_n)_{n \in \mathbb{N}}$ in $\mathcal{G}$ such that $\int g_n d\mu \xrightarrow{n \rightarrow \infty} \gamma$, and define the function $f_n = g_1 \vee \dots \vee g_n$. Now (7.11) implies $f_n \in \mathcal{G}$. Letting $f := \sup \{f_n : n \in \mathbb{N}\}$, the monotone convergence theorem yields

$$\int_A f d\mu = \sup_{n \in \mathbb{N}} \int_A f_n d\mu \leq v(A) \quad \text{for all } A \in \mathcal{A}$$

(that is, $f \in \mathcal{G}$), and

$$\int f d\mu = \sup_{n \in \mathbb{N}} \int f_n d\mu \geq \sup_{n \in \mathbb{N}} \int g_n d\mu = \gamma.$$

Hence $\int f d\mu = \gamma \leq \nu(\Omega)$. Now define, for any $A \in \mathcal{A}$,

$$\nu_a(A) := \int_A f d\mu \quad \text{and} \quad \nu_s(A) := v(A) - \nu_a(A)$$

By construction, $\nu_a \ll \mu$ is a finite measure with density f with respect to μ . Since $f \in \mathcal{G}$, we have $\nu_s(A) = v(A) - \int_A f d\mu \geq 0$ for all $A \in \mathcal{A}$, and thus also ν_s is a finite measure. It remains to show $\nu_s \perp \mu$.

At this point we use Lemma 7.46. Assume that we had $\nu_s \neq \mu$. Then there would be an $\varepsilon > 0$ and an $A \in \mathcal{A}$ with $\mu(A) > 0$ such that $\varepsilon\mu(E) \leq \nu_s(E)$ for all $E \subset A$, $E \in \mathcal{A}$. Then, for $B \in \mathcal{A}$, we would have

$$\begin{aligned} \int_B (f + \varepsilon \mathbb{1}_A) d\mu &= \int_B f d\mu + \varepsilon \mu(A \cap B) \\ &\leq \nu_a(B) + \nu_s(A \cap B) \leq \nu_a(B) + \nu_s(B) = v(B) \end{aligned}$$

In other words, $(f + \varepsilon \mathbb{1}_A) \in \mathcal{G}$ and thus $\int (f + \varepsilon \mathbb{1}_A) d\mu = \gamma + \varepsilon \mu(A) > \gamma$, contradicting the definition of γ . Hence in fact $\nu_s \perp \mu$.

Exercise 7.5.1 Let μ be a σ -finite measure on $(\Omega, \mathcal{A})$ and let φ be a signed measure on $(\Omega, \mathcal{A})$. Show that, analogously to the Radon-Nikodym theorem, the following two statements are equivalent:

- (i) $\varphi(A) = 0$ for all $A \in \mathcal{A}$ with $\mu(A) = 0$.
- (ii) There is an $f \in \mathcal{L}^1(\mu)$ with $\varphi = f\mu$; hence $\int_A f d\mu = \varphi(A)$ for all $A \in \mathcal{A}$.

Exercise 7.5.2 Let μ, ν, α be finite measures on $(\Omega, \mathcal{A})$ with $\nu \ll \mu \ll \alpha$.

- (i) Show the chain rule for the Radon-Nikodym derivative:

$$\frac{d\nu}{d\alpha} = \frac{d\nu}{d\mu} \frac{d\mu}{d\alpha} \quad \alpha\text{-a.e.}$$

- (ii) Show that $f := \frac{d\nu}{d(\mu+\nu)}$ exists and that $\frac{d\nu}{d\mu} = \frac{f}{1-f}$ holds μ -a.e.

5.7.6 7.6 Supplement: Dual Spaces

By the Riesz-Fréchet theorem (Theorem 7.26), every continuous linear functional $F : L^2(\mu) \rightarrow \mathbb{R}$ has a representation $F(g) = \langle f, g \rangle$ for some $f \in L^2(\mu)$. On the other hand, for any $f \in L^2(\mu)$, the map $L^2(\mu) \rightarrow \mathbb{R}, g \mapsto \langle f, g \rangle$ is continuous and linear. Hence $L^2(\mu)$ is canonically isomorphic to its topological dual space $(L^2(\mu))'$. This dual space is defined as follows.

Definition 7.47 (Dual space) Let $(V, \|\cdot\|)$ be a Banach space. The dual space V' of V is defined by

$$V' := \{F : V \rightarrow \mathbb{R} \text{ is continuous and linear} \}.$$

For $F \in V'$, we define $\|F\|' := \sup\{|F(f)| : \|f\| = 1\}$.

Remark 7.48 As F is continuous, for any $\delta > 0$, there exists an $\varepsilon > 0$ such that $|F(f)| < \delta$ for all $f \in V$ with $\|f\| < \varepsilon$. Hence $\|F\|' \leq \delta/\varepsilon < \infty$.

We are interested in the case $V = L^p(\mu)$ for $p \in [1, \infty]$ and write $\|F\|'_p$ for the norm of $F \in V'$. In the particular case $V = L^2(\mu)$, by the Cauchy-Schwarz inequality, we have $\|F\|'_2 = \|f\|_2$. This can be generalized:

Lemma 7.49 Let $p, q \in [1, \infty]$ with $\frac{1}{p} + \frac{1}{q} = 1$. The canonical map

$$\begin{aligned} \kappa : L^q(\mu) &\rightarrow (L^p(\mu))' \\ \kappa(f)(g) &= \int fg d\mu \quad \text{for } f \in L^q(\mu), g \in L^p(\mu) \end{aligned}$$

is an isometry; that is, $\|\kappa(f)\|'_p = \|f\|_q$.

Proof We show equality by showing the two inequalities separately.

" $\leq$ " This follows from Hölder's inequality.

" $\geq$ " For any admissible pair p, q and all $f \in L^q(\mu), g \in L^p(\mu)$, by the definition of the operator norm, $\|\kappa(f)\|'_p \|g\|_p \geq |\int fg d\mu|$. Define the sign function $\text{sign}(x) = \mathbb{1}_{(0, \infty)}(x) - \mathbb{1}_{(-\infty, 0)}(x)$. Replacing g by $\tilde{g} := |g| \text{sign}(f)$ (note that $\|\tilde{g}\|_p = \|g\|_p$), we obtain

$$\|\kappa(f)\|'_p \|g\|_p \geq \left| \int f \tilde{g} d\mu \right| = \|fg\|_1 \quad (7.12)$$

First consider the case $q = 1$ and $f \in L^1(\mu)$. Applying (7.12) with $g \equiv 1 \in L^\infty(\mu)$ yields $\|\kappa(f)\|'_\infty \geq \|f\|_1$.

Now let $q \in (1, \infty)$. Let $g := |f|^{q-1}$. Since $\frac{q-1}{q} = \frac{1}{p}$, we have

$$\|\kappa(f)\|'_p \cdot \|g\|_p \geq \|fg\|_1 = \| |f|^q \|_1 = \|f\|_q^q = \|f\|_q \cdot \|f\|_q^{q-1} = \|f\|_q \cdot \|g\|_p.$$

Finally, let $q = \infty$. Without loss of generality, assume $\|f\|_\infty \in (0, \infty)$. Let $\varepsilon > 0$. Then there exists an $A_\varepsilon \in \mathcal{A}$ with $0 < \mu(A_\varepsilon) < \infty$ such that

$$A_\varepsilon \subset \{|f| > (1 - \varepsilon)\|f\|_\infty\}$$

If we let $g = \frac{1}{\mu(A_\varepsilon)} \mathbb{1}_{A_\varepsilon}$, then $\|g\|_1 = 1$ and $\|\kappa(f)\|'_1 \geq \|fg\|_1 \geq (1 - \varepsilon)\|f\|_\infty$.

Theorem 7.50 Let $p \in [1, \infty)$ and assume $\frac{1}{p} + \frac{1}{q} = 1$. Then $L^q(\mu)$ is isomorphic to its dual space $(L^p(\mu))'$ by virtue of the isometry κ .

Proof The proof makes use of the Radon-Nikodym theorem (Corollary 7.34). However, here we only sketch the proof since we do not want to go into the details of signed measures and signed contents. A signed content ν is an additive set function that is the difference $\nu = \nu^+ - \nu^-$ of two finite contents. This definition is parallel to that of a signed measure that is the difference of two finite measures.

As κ is an isometry, κ in particular is injective. Hence we only have to show that κ is surjective. Let $F \in (L^p(\mu))'$. Then $v(A) = F(\mathbb{1}_A)$ is a signed content on $\mathcal{A}$ and we have

$$|v(A)| \leq \|F\|'_p(\mu(A))^{1/p}.$$

Since μ is $\emptyset$ -continuous, v is also $\emptyset$ -continuous and is thus a signed measure on $\mathcal{A}$. We even have $v \ll \mu$. By the Radon-Nikodym theorem (Corollary 7.34) (applied to the measures v^- and v^+ ; see Exercise 7.5.1), v admits a density with respect to μ ; that is, a measurable function f with $v = f\mu$.

Let

$$\mathbb{E}_f := \{g : g \text{ is a simple function with } \mu(g \neq 0) < \infty\}$$

and let

$$\mathbb{E}_f^+ := \{g \in \mathbb{E}_f : g \geq 0\}$$

Then, for $g \in \mathbb{E}_f$,

$$F(g) = \int g f d\mu \quad (7.13)$$

In order to show that (7.13) holds for all $g \in \mathcal{L}^p(\mu)$, we first show $f \in \mathcal{L}^q(\mu)$. To this end, we distinguish two cases.

Case 1: $p = 1$ For every $\alpha > 0$,

$$\begin{aligned} \mu(\{|f| > \alpha\}) &\leq \frac{1}{\alpha} v(\{|f| > \alpha\}) \\ &= \frac{1}{\alpha} F(\mathbb{1}_{\{|f| > \alpha\}}) \leq \frac{1}{\alpha} \|F\|'_1 \cdot \|\mathbb{1}_{\{|f| > \alpha\}}\|_1 = \frac{1}{\alpha} \|F\|'_1 \cdot \mu(\{|f| > \alpha\}) \end{aligned}$$

This implies $\mu(\{|f| > \alpha\}) = 0$ if $\alpha > \|F\|'_1$; hence $\|f\|_\infty \leq \|F\|'_1 < \infty$.

Case 2: $p \in (1, \infty)$ By Theorem 1.96, there are $g_1, g_2, \dots \in \mathbb{E}_f^+$ such that $g_n \uparrow |f| \mu$ -a.e. Define $h_n = \text{sign}(f) (g_n)^{q-1} \in \mathbb{E}_f$; hence

$$\begin{aligned} \|g_n\|_q^q &\leq \int h_n f d\mu = F(h_n) \\ &\leq \|F\|'_p \cdot \|h_n\|_p = \|F\|'_p \cdot (\|g_n\|_q)^{q-1} \end{aligned}$$

Thus we have $\|g_n\|_q \leq \|F\|'_p$. Monotone convergence (Theorem 4.20) now yields $\|f\|_q \leq \|F\|'_p < \infty$; hence $f \in \mathcal{L}^q(\mu)$.

Concluding, the map $\tilde{F} : g \mapsto \int g f d\mu$ is in $(L^p(\mu))'$, and $\tilde{F}(g) = F(g)$ for every $g \in \mathbb{E}_f$. Since $\tilde{F}$ is continuous and $\mathbb{E}_f \subset L^p(\mu)$ is dense, we get $\tilde{F} = F$.

Remark 7.51 For $p = \infty$, the statement of Theorem 7.50 is false in general. (For finite $\mathcal{A}$, the claim is trivially true even for $p = \infty$.) For example, let $\Omega = \mathbb{N}$, $\mathcal{A} = 2^\Omega$ and let μ be the counting measure. Thus we consider sequence spaces $\ell^p = L^p(\mathbb{N}, 2^\mathbb{N}, \mu)$. For the subspace $\ell^K \subset \ell^\infty$ of convergent sequences, $F : \ell^K \rightarrow \mathbb{R}$, $(a_n)_{n \in \mathbb{N}} \mapsto \lim_{n \rightarrow \infty} a_n$ is a continuous linear functional. By the Hahn-Banach theorem of functional analysis (see, e.g., [87] or [173]), F can be extended to a continuous linear functional on ℓ^∞ . However, clearly there is no sequence $(b_n)_{n \in \mathbb{N}} \in \ell^1$ with $F((a_n)_{n \in \mathbb{N}}) = \sum_{m=1}^\infty a_m b_m$.

Exercise 7.6.1 Show that $\mathbb{E}_f \subset L^p(\mu)$ is dense if $p \in [1, \infty)$.

5.8 8 Conditional Expectations

If there is partial information on the outcome of a random experiment, the probabilities for the possible events may change. The concept of conditional probabilities and conditional expectations formalizes the corresponding calculus.

5.8.1 8.1 Elementary Conditional Probabilities

Example 8.1 We throw a die and consider the events

$$\begin{aligned} A &:= \{ \text{the face shows an odd number} \}, \\ B &:= \{ \text{the face shows three or smaller} \}. \end{aligned}$$

Clearly, $\mathbf{P}[A] = \frac{1}{2}$ and $\mathbf{P}[B] = \frac{1}{2}$. However, what is the probability that A occurs if we already know that B occurs?

We model the experiment on the probability space $(\Omega, \mathcal{A}, \mathbf{P})$, where $\Omega = \{1, \dots, 6\}$, $\mathcal{A} = 2^\Omega$ and $\mathbf{P}$ is the uniform distribution on Ω . Then

$$A = \{1, 3, 5\} \quad \text{and} \quad B = \{1, 2, 3\}.$$

If we know that B has occurred, it is plausible to assume the uniform distribution on the remaining possible outcomes; that is, on $\{1, 2, 3\}$. Thus we define a new probability measure $\mathbf{P}_B$ on $(B, 2^B)$ by

$$\mathbf{P}_B[C] = \frac{\#C}{\#B} \quad \text{for } C \subset B.$$

By assigning the points in $\Omega \setminus B$ probability zero (since they are impossible if B has occurred), we can extend $\mathbf{P}_B$ to a measure on Ω :

$$\mathbf{P}[C | B] := \mathbf{P}_B[C \cap B] = \frac{\#(C \cap B)}{\#B} \quad \text{for } C \subset \Omega.$$

In this way, we get

$$\mathbf{P}[A | B] = \frac{\#\{1, 3\}}{\#\{1, 2, 3\}} = \frac{2}{3}.$$

Motivated by this example, we make the following definition.

Definition 8.2 (Conditional probability) Let $(\Omega, \mathcal{A}, \mathbf{P})$ be a probability space and $B \in \mathcal{A}$. We define the conditional probability given B for any $A \in \mathcal{A}$ by

$$\mathbf{P}[A | B] = \begin{cases} \frac{\mathbf{P}[A \cap B]}{\mathbf{P}[B]}, & \text{if } \mathbf{P}[B] > 0 \\ 0, & \text{otherwise} \end{cases} \quad (8.1)$$

Remark 8.3 The specification in (8.1) for the case $\mathbf{P}[B] = 0$ is arbitrary and is of no importance.

Theorem 8.4 If $\mathbf{P}[B] > 0$, then $\mathbf{P}[\cdot | B]$ is a probability measure on $(\Omega, \mathcal{A})$.

Proof This is obvious.

Theorem 8.5 Let $A, B \in \mathcal{A}$ with $\mathbf{P}[A], \mathbf{P}[B] > 0$. Then

A, B are independent $\iff \mathbf{P}[A | B] = \mathbf{P}[A] \iff \mathbf{P}[B | A] = \mathbf{P}[B]$.

Proof This is trivial!

Theorem 8.6 (Summation formula) Let I be a countable set and let $(B_i)_{i \in I}$ be pairwise disjoint sets with $\mathbf{P}[\biguplus_{i \in I} B_i] = 1$. Then, for any $A \in \mathcal{A}$,

$$\mathbf{P}[A] = \sum_{i \in I} \mathbf{P}[A \mid B_i] \mathbf{P}[B_i] \quad (8.2)$$

Proof Due to the σ -additivity of $\mathbf{P}$, we have

$$\mathbf{P}[A] = \mathbf{P}\left[\biguplus_{i \in I} (A \cap B_i)\right] = \sum_{i \in I} \mathbf{P}[A \cap B_i] = \sum_{i \in I} \mathbf{P}[A \mid B_i] \mathbf{P}[B_i]$$

Theorem 8.7 (Bayes' formula) Let I be a countable set and let $(B_i)_{i \in I}$ be pairwise disjoint sets with $\mathbf{P}\left[\biguplus_{i \in I} B_i\right] = 1$. Then, for any $A \in \mathcal{A}$ with $\mathbf{P}[A] > 0$ and any $k \in I$,

$$\mathbf{P}[B_k \mid A] = \frac{\mathbf{P}[A \mid B_k] \mathbf{P}[B_k]}{\sum_{i \in I} \mathbf{P}[A \mid B_i] \mathbf{P}[B_i]} \quad (8.3)$$

Proof We have

$$\mathbf{P}[B_k \mid A] = \frac{\mathbf{P}[B_k \cap A]}{\mathbf{P}[A]} = \frac{\mathbf{P}[A \mid B_k] \mathbf{P}[B_k]}{\mathbf{P}[A]}.$$

Now use the expression in (8.2) for $\mathbf{P}[A]$.

Example 8.8 In the production of certain electronic devices, a fraction of 2% of the production is defective. A quick test detects a defective device with probability 95%; however, with probability 10% it gives a false alarm for an intact device.

If the test gives an alarm, what is the probability that the device just tested is indeed defective?

We formalize the description given above. Let

$$\begin{aligned} A &:= \{ \text{device is declared as defective} \} \\ B &:= \{ \text{device is defective} \} \end{aligned}$$

and

$$\begin{aligned} \mathbf{P}[B] &= 0.02, & \mathbf{P}[B^c] &= 0.98, \\ \mathbf{P}[A \mid B] &= 0.95, & \mathbf{P}[A \mid B^c] &= 0.1 \end{aligned}$$

Bayes' formula yields

$$\begin{aligned} \mathbf{P}[B \mid A] &= \frac{\mathbf{P}[A \mid B] \mathbf{P}[B]}{\mathbf{P}[A \mid B] \mathbf{P}[B] + \mathbf{P}[A \mid B^c] \mathbf{P}[B^c]} \\ &= \frac{0.95 \cdot 0.02}{0.95 \cdot 0.02 + 0.1 \cdot 0.98} = \frac{19}{117} \approx 0.162 \end{aligned}$$

On the other hand, the probability that a device that was not classified as defective is in fact defective is

$$\mathbf{P}[B \mid A^c] = \frac{0.05 \cdot 0.02}{0.05 \cdot 0.02 + 0.9 \cdot 0.98} = \frac{1}{883} \approx 0.00113$$

Now let $X \in \mathcal{L}^1(\mathbf{P})$. If $A \in \mathcal{A}$, then clearly also $\mathbb{1}_A X \in \mathcal{L}^1(\mathbf{P})$. We define

$$\mathbf{E}[X; A] := \mathbf{E}[\mathbb{1}_A X]. \quad (8.4)$$

If $\mathbf{P}[A] > 0$, then $\mathbf{P}[\cdot \mid A]$ is a probability measure. Since $\mathbb{1}_A X \in \mathcal{L}^1(\mathbf{P})$, we have $X \in \mathcal{L}^1(\mathbf{P}[\cdot \mid A])$. Hence we can define the expectation of X with respect to $\mathbf{P}[\cdot \mid A]$.

Definition 8.9 Let $X \in \mathcal{L}^1(\mathbf{P})$ and $A \in \mathcal{A}$. Then we define

$$\mathbf{E}[X | A] := \int X(\omega) \mathbf{P}[d\omega | A] = \begin{cases} \frac{\mathbf{E}[\mathbb{1}_A X]}{\mathbf{P}[A]}, & \text{if } \mathbf{P}[A] > 0 \\ 0, & \text{else} \end{cases} \quad (8.5)$$

Clearly, $\mathbf{P}[B | A] = \mathbf{E}[\mathbb{1}_B | A]$ for all $B \in \mathcal{A}$.

Consider now the situation that we studied with the summation formula for conditional probabilities. Hence, let I be a countable set and let $(B_i)_{i \in I}$ be pairwise disjoint events with $\biguplus_{i \in I} B_i = \Omega$. We define $\mathcal{F} := \sigma(B_i, i \in I)$. For $X \in \mathcal{L}^1(\mathbf{P})$, we define a map $\mathbf{E}[X | \mathcal{F}] : \Omega \rightarrow \mathbb{R}$ by

$$\mathbf{E}[X | \mathcal{F}](\omega) = \mathbf{E}[X | B_i] \iff B_i \ni \omega. \quad (8.6)$$

Lemma 8.10 The map $\mathbf{E}[X | \mathcal{F}]$ has the following properties.

- (i) $\mathbf{E}[X | \mathcal{F}]$ is $\mathcal{F}$ -measurable.
- (ii) $\mathbf{E}[X | \mathcal{F}] \in \mathcal{L}^1(\mathbf{P})$, and for any $A \in \mathcal{F}$, we have

$$\int_A \mathbf{E}[X | \mathcal{F}] d\mathbf{P} = \int_A X d\mathbf{P}$$

Proof (i) Let f be the map $f : \Omega \rightarrow I$ with

$$f(\omega) = i \iff B_i \ni \omega.$$

Further, let $g : I \rightarrow \mathbb{R}, i \mapsto \mathbf{E}[X | B_i]$. Since I is discrete, g is measurable. Since f is $\mathcal{F}$ -measurable, $\mathbf{E}[X | \mathcal{F}] = g \circ f$ is also $\mathcal{F}$ -measurable.

- (ii) Let $A \in \mathcal{F}$ and $J \subset I$ with $A = \biguplus_{j \in J} B_j$. Let $J' := \{i \in J : \mathbf{P}[B_i] > 0\}$. Hence

$$\int_A \mathbf{E}[X | \mathcal{F}] d\mathbf{P} = \sum_{i \in J'} \mathbf{P}[B_i] \mathbf{E}[X | B_i] = \sum_{i \in J'} \mathbf{E}[\mathbb{1}_{B_i} X] = \int_A X d\mathbf{P}$$

Exercise 8.1.1 (Lack of memory of the exponential distribution) Let $X > 0$ be a strictly positive random variable and let $\theta > 0$. Show that X is exponentially distributed if and only if

$$\mathbf{P}[X > t + s | X > s] = \mathbf{P}[X > t] \quad \text{for all } s, t \geq 0$$

In particular, $X \sim \exp_\theta$ if and only if $\mathbf{P}[X > t + s | X > s] = e^{-\theta t}$ for all $s, t \geq 0$.

Exercise 8.1.2 Consider a theater with n seats that is fully booked for this evening. Each of the n people entering the theater (one by one) has a seat reservation. However, the first person is absent-minded and takes a seat at random. Any subsequent person takes his or her reserved seat if it is free and otherwise picks a free seat at random.

- (i) What is the probability that the last person gets his or her reserved seat?
- (ii) What is the probability that the k th person gets his or her reserved seat?

5.8.2 8.2 Conditional Expectations

Let X be a random variable that is uniformly distributed on $[0, 1]$. Assume that if we know the value $X = x$, the random variables $Y_1, \dots, Y_n$ are independent and Ber_x -distributed. So far, with our machinery we can only deal with conditional probabilities of the type $\mathbf{P}[\cdot | X \in [a, b]]$, $a < b$ (since $X \in [a, b]$ has positive probability). How about $\mathbf{P}[Y_1 = \dots = Y_n = 1 | X = x]$? Intuitively, this should be x^n . We thus need a notion of conditional probabilities that allows us to deal with conditioning on events with probability zero and that is consistent with our intuition. In the next section, we will see that in the current example this can be done using transition kernels. First, however, we have to consider a more general situation.

In the following, $\mathcal{F} \subset \mathcal{A}$ will be a sub- σ -algebra and $X \in \mathcal{L}^1(\Omega, \mathcal{A}, \mathbf{P})$. In analogy with Lemma 8.10, we make the following definition.

Definition 8.11 (Conditional expectation) A random variable Y is called a conditional expectation of X given $\mathcal{F}$, symbolically $\mathbf{E}[X | \mathcal{F}] := Y$, if:

- (i) Y is $\mathcal{F}$ -measurable.
- (ii) For any $A \in \mathcal{F}$, we have $\mathbf{E}[X \mathbb{1}_A] = \mathbf{E}[Y \mathbb{1}_A]$.

For $B \in \mathcal{A}$, $\mathbf{P}[B | \mathcal{F}] := \mathbf{E}[\mathbb{1}_B | \mathcal{F}]$ is called a conditional probability of B given the σ -algebra $\mathcal{F}$.

Theorem 8.12 $\mathbf{E}[X | \mathcal{F}]$ exists and is unique (up to equality almost surely).

Since conditional expectations are defined only up to equality a.s., all equalities with conditional expectations are understood as equalities a.s., even if we do not say so explicitly.

Proof Uniqueness. Let Y and Y' be random variables that fulfill (i) and (ii). Let $A = \{Y > Y'\} \in \mathcal{F}$. Then, by (ii),

$$0 = \mathbf{E}[Y \mathbb{1}_A] - \mathbf{E}[Y' \mathbb{1}_A] = \mathbf{E}[(Y - Y') \mathbb{1}_A].$$

Since $(Y - Y') \mathbb{1}_A \geq 0$, we have $\mathbf{P}[A] = 0$; hence $Y \leq Y'$ almost surely. Similarly, we get $Y \geq Y'$ almost surely.

Existence. Let $X^+ = X \vee 0$ and $X^- = X^+ - X$. By

$$Q^\pm(A) := \mathbf{E}[X^\pm \mathbb{1}_A] \quad \text{for all } A \in \mathcal{F}$$

we define two finite measures on $(\Omega, \mathcal{F})$. Clearly, $Q^\pm \ll \mathbf{P}$; hence the RadonNikodym theorem (Corollary 7.34) yields the existence of $\mathcal{F}$ -measurable densities $Y^\pm$ such that

$$Q^\pm(A) = \int_A Y^\pm d\mathbf{P} = \mathbf{E}[Y^\pm \mathbb{1}_A]$$

Now define $Y = Y^+ - Y^-$.

Definition 8.13 If Y is a random variable and $X \in \mathcal{L}^1(\mathbf{P})$, then we define $\mathbf{E}[X | Y] := \mathbf{E}[X | \sigma(Y)]$.

Theorem 8.14 (Properties of the conditional expectation) Let $(\Omega, \mathcal{A}, \mathbf{P})$ and let X be as above. Let $\mathcal{G} \subset \mathcal{F} \subset \mathcal{A}$ be σ -algebras and let $Y \in \mathcal{L}^1(\Omega, \mathcal{A}, \mathbf{P})$. Then:

- (i) (Linearity) $\mathbf{E}[\lambda X + Y | \mathcal{F}] = \lambda \mathbf{E}[X | \mathcal{F}] + \mathbf{E}[Y | \mathcal{F}]$.
- (ii) (Monotonicity) If $X \geq Y$ a.s., then $\mathbf{E}[X | \mathcal{F}] \geq \mathbf{E}[Y | \mathcal{F}]$.
- (iii) If $\mathbf{E}[|XY|] < \infty$ and Y is measurable with respect to $\mathcal{F}$, then

$$\mathbf{E}[XY | \mathcal{F}] = Y \mathbf{E}[X | \mathcal{F}] \quad \text{and} \quad \mathbf{E}[Y | \mathcal{F}] = \mathbf{E}[Y | Y] = Y.$$

- (iv) (Tower property) $\mathbf{E}[\mathbf{E}[X | \mathcal{F}] | \mathcal{G}] = \mathbf{E}[\mathbf{E}[X | \mathcal{G}] | \mathcal{F}] = \mathbf{E}[X | \mathcal{G}]$.
- (v) (Triangle inequality) $\mathbf{E}[|X| | \mathcal{F}] \geq |\mathbf{E}[X | \mathcal{F}]|$.
- (vi) (Independence) If $\sigma(X)$ and $\mathcal{F}$ are independent, then $\mathbf{E}[X | \mathcal{F}] = \mathbf{E}[X]$.
- (vii) If $\mathbf{P}[A] \in \{0, 1\}$ for any $A \in \mathcal{F}$, then $\mathbf{E}[X | \mathcal{F}] = \mathbf{E}[X]$.
- (viii) (Dominated convergence) Assume $Y \in \mathcal{L}^1(\mathbf{P})$, $Y \geq 0$ and $(X_n)_{n \in \mathbb{N}}$ is a sequence of random variables with $|X_n| \leq Y$ for $n \in \mathbb{N}$ and such that $X_n \xrightarrow{n \rightarrow \infty} X$ a.s. Then

$$\lim_{n \rightarrow \infty} \mathbf{E}[X_n | \mathcal{F}] = \mathbf{E}[X | \mathcal{F}] \quad \text{a.s. and in } L^1(\mathbf{P}) \quad (8.7)$$

Proof (i) The right-hand side is $\mathcal{F}$ -measurable; hence, for $A \in \mathcal{F}$,

$$\begin{aligned} \mathbf{E}[\mathbb{1}_A(\lambda \mathbf{E}[X | \mathcal{F}] + \mathbf{E}[Y | \mathcal{F}])] &= \lambda \mathbf{E}[\mathbb{1}_A \mathbf{E}[X | \mathcal{F}]] + \mathbf{E}[\mathbb{1}_A \mathbf{E}[Y | \mathcal{F}]] \\ &= \lambda \mathbf{E}[\mathbb{1}_A X] + \mathbf{E}[\mathbb{1}_A Y] \\ &= \mathbf{E}[\mathbb{1}_A(\lambda X + Y)] \end{aligned}$$

(ii) Let $A = \{\mathbf{E}[X | \mathcal{F}] < \mathbf{E}[Y | \mathcal{F}]\} \in \mathcal{F}$. Since we have $X \geq Y$, we get $\mathbf{E}[\mathbb{1}_A(X - Y)] \geq 0$ and thus $\mathbf{P}[A] = 0$.

(iii) First assume $X \geq 0$ and $Y \geq 0$. For $n \in \mathbb{N}$, define $Y_n = 2^{-n} \lfloor 2^n Y \rfloor$. Then $Y_n \uparrow Y$ and $Y_n \mathbf{E}[X | \mathcal{F}] \uparrow Y \mathbf{E}[X | \mathcal{F}]$ (since $\mathbf{E}[X | \mathcal{F}] \geq 0$ by (ii)). By the monotone convergence theorem (Lemma 4.6(ii)),

$$\mathbf{E}[\mathbb{1}_A Y_n \mathbf{E}[X | \mathcal{F}]] \xrightarrow{n \rightarrow \infty} \mathbf{E}[\mathbb{1}_A Y \mathbf{E}[X | \mathcal{F}]].$$

On the other hand,

$$\begin{aligned} \mathbf{E}[\mathbb{1}_A Y_n \mathbf{E}[X | \mathcal{F}]] &= \sum_{k=1}^{\infty} \mathbf{E}[\mathbb{1}_A \mathbb{1}_{\{Y_n = k2^{-n}\}} k 2^{-n} \mathbf{E}[X | \mathcal{F}]] \\ &= \sum_{k=1}^{\infty} \mathbf{E}[\mathbb{1}_A \mathbb{1}_{\{Y_n = k2^{-n}\}} k 2^{-n} X] \\ &= \mathbf{E}[\mathbb{1}_A Y_n X] \xrightarrow{n \rightarrow \infty} \mathbf{E}[\mathbb{1}_A Y X] \end{aligned}$$

Hence $\mathbf{E}[\mathbb{1}_A Y \mathbf{E}[X | \mathcal{F}]] = \mathbf{E}[\mathbb{1}_A Y X]$. In the general case, write $X = X^+ - X^-$ and $Y = Y^+ - Y^-$ and exploit the linearity of the conditional expectation.

(iv) The second equality follows from (iii) with $Y = \mathbf{E}[X | \mathcal{G}]$ and $X = 1$. Now let $A \in \mathcal{G}$. Then, in particular, $A \in \mathcal{F}$; hence

$$\mathbf{E}[\mathbb{1}_A \mathbf{E}[\mathbf{E}[X | \mathcal{F}] | \mathcal{G}]] = \mathbf{E}[\mathbb{1}_A \mathbf{E}[X | \mathcal{F}]] = \mathbf{E}[\mathbb{1}_A X] = \mathbf{E}[\mathbb{1}_A \mathbf{E}[X | \mathcal{G}]].$$

(v) This follows from (i) and (ii) with $X = X^+ - X^-$.

(vi) Trivially, $\mathbf{E}[X]$ is measurable with respect to $\mathcal{F}$. Let $A \in \mathcal{F}$. Then X and $\mathbb{1}_A$ are independent; hence $\mathbf{E}[\mathbf{E}[X | \mathcal{F}] \mathbb{1}_A] = \mathbf{E}[X \mathbb{1}_A] = \mathbf{E}[X] \mathbf{E}[\mathbb{1}_A]$.

(vii) For any $A \in \mathcal{F}$ and $B \in \mathcal{A}$, we have $\mathbf{P}[A \cap B] = 0$ if $\mathbf{P}[A] = 0$, and $\mathbf{P}[A \cap B] = \mathbf{P}[B]$ if $\mathbf{P}[A] = 1$. Hence $\mathcal{F}$ and $\mathcal{A}$ are independent and thus $\mathcal{F}$ is independent of any sub- σ -algebra of $\mathcal{A}$. In particular, $\mathcal{F}$ and $\sigma(X)$ are independent. Hence the claim follows from (vi).

(viii) Let $|X_n| \leq Y$ for any $n \in \mathbb{N}$ and $X_n \xrightarrow{n \rightarrow \infty} X$ almost surely. Define $Z_n := \sup_{k \geq n} |X_k - X|$. Then $0 \leq Z_n \leq 2Y$ and $Z_n \xrightarrow{\text{a.s.}} 0$. By Corollary 6.26 (dominated convergence), we have $\mathbf{E}[Z_n] \xrightarrow{n \rightarrow \infty} 0$; hence, by the triangle inequality,

$$\mathbf{E}[|\mathbf{E}[X_n | \mathcal{F}] - \mathbf{E}[X | \mathcal{F}]|] \leq \mathbf{E}[\mathbf{E}[|X_n - X| | \mathcal{F}]] = \mathbf{E}[|X_n - X|] \leq \mathbf{E}[Z_n] \xrightarrow{n \rightarrow \infty} 0.$$

However, this is the $L^1(\mathbf{P})$ -convergence in (8.7). As $(Z_n)_{n \in \mathbb{N}}$ is decreasing, by (ii) also $(\mathbf{E}[Z_n | \mathcal{F}])_{n \in \mathbb{N}}$ decreases to some limit, say, Z . By Fatou's lemma,

$$\mathbf{E}[Z] \leq \lim_{n \rightarrow \infty} \mathbf{E}[\mathbf{E}[Z_n | \mathcal{F}]] = \lim_{n \rightarrow \infty} \mathbf{E}[Z_n] = 0.$$

Hence $Z = 0$ and thus $\mathbf{E}[Z_n | \mathcal{F}] \xrightarrow{n \rightarrow \infty} 0$ almost surely. However, by (v),

$$|\mathbf{E}[X_n | \mathcal{F}] - \mathbf{E}[X | \mathcal{F}]| \leq \mathbf{E}[Z_n | \mathcal{F}].$$

Remark 8.15 Intuitively, $\mathbf{E}[X | \mathcal{F}]$ is the best prediction we can make for the value of X if we only have the information of the σ -algebra $\mathcal{F}$. For example, if $\sigma(X) \subset \mathcal{F}$ (that is, if we know X already), then $\mathbf{E}[X | \mathcal{F}] = X$, as shown in (iii). At the other end of the spectrum is the case where X and $\mathcal{F}$ are independent; that is, where knowledge of $\mathcal{F}$ does not give any information on X . Here the best prediction for X is its mean; hence $\mathbf{E}[X] = \mathbf{E}[X | \mathcal{F}]$, as shown in (vi).

What exactly do we mean by "best prediction"? For square integrable random variables X , by the best prediction for X we will understand the $\mathcal{F}$ -measurable random variable that minimizes the L^2 -distance from X . The next corollary shows that the conditional expectation is in fact this minimizer.

Remark 8.16 Let $X : \Omega \rightarrow \mathbb{R}$ be a random variable such that $X^- \in \mathcal{L}^1(\mathbf{P})$. We can define the conditional expectation as the monotone limit

$$\mathbf{E}[X | \mathcal{F}] := \lim_{n \rightarrow \infty} \mathbf{E}[X_n | \mathcal{F}]$$

where $-X^- \leq X_1$ and $X_n \uparrow X$. Due to the monotonicity of the conditional expectation (Theorem 8.14(ii)) it is easy to show that the limit does not depend on the choice of the sequence (X_n) and that it fulfills the conditions of Definition 8.11. Analogously, we can define the conditional expectation $X^+ \in \mathcal{L}^1(\mathbf{P})$. For this generalization of the conditional expectation, we still have $\mathbf{E}[X | \mathcal{F}] \leq \mathbf{E}[Y | \mathcal{F}]$ a.s. if $Y \geq X$ a.s. (see Exercise 8.2.1).

Corollary 8.17 (Conditional expectation as projection) Let $\mathcal{F} \subset \mathcal{A}$ be a σ -algebra and let X be a random variable with $\mathbf{E}[X^2] < \infty$. Then $\mathbf{E}[X | \mathcal{F}]$ is the orthogonal projection of X on $\mathcal{L}^2(\Omega, \mathcal{F}, \mathbf{P})$. That is, for any $\mathcal{F}$ -measurable Y with $\mathbf{E}[Y^2] < \infty$,

$$\mathbf{E}[(X - Y)^2] \geq \mathbf{E}[(X - \mathbf{E}[X | \mathcal{F}])^2]$$

with equality if and only if $Y = \mathbf{E}[X | \mathcal{F}]$.

Proof First assume that $\mathbf{E}[\mathbf{E}[X | \mathcal{F}]^2] < \infty$. (In Theorem 8.20, we will see that we have $\mathbf{E}[\mathbf{E}[X | \mathcal{F}]^2] \leq \mathbf{E}[X^2]$, but here we want to keep the proof self-contained.) Let Y be $\mathcal{F}$ -measurable and assume $\mathbf{E}[Y^2] < \infty$. Then, by the Cauchy-Schwarz inequality, we have $\mathbf{E}[|XY|] < \infty$. Thus, using the tower property, we infer $\mathbf{E}[XY] = \mathbf{E}[\mathbf{E}[X | \mathcal{F}]Y]$ and $\mathbf{E}[X\mathbf{E}[X | \mathcal{F}]] = \mathbf{E}[\mathbf{E}[X\mathbf{E}[X | \mathcal{F}] | \mathcal{F}]] = \mathbf{E}[\mathbf{E}[X | \mathcal{F}]^2]$. Summing up, we have

$$\begin{aligned} & \mathbf{E}[(X - Y)^2] - \mathbf{E}[(X - \mathbf{E}[X | \mathcal{F}])^2] \\ &= \mathbf{E}[X^2 - 2XY + Y^2 - X^2 + 2X\mathbf{E}[X | \mathcal{F}] - \mathbf{E}[X | \mathcal{F}]^2] \\ &= \mathbf{E}[Y^2 - 2Y\mathbf{E}[X | \mathcal{F}] + \mathbf{E}[X | \mathcal{F}]^2] \\ &= \mathbf{E}[(Y - \mathbf{E}[X | \mathcal{F}])^2] \geq 0 \end{aligned}$$

For the case $\mathbf{E}[\mathbf{E}[X | \mathcal{F}]^2] < \infty$, we are done. Hence, it suffices to show that this condition follows from the assumption $\mathbf{E}[X^2] < \infty$. For $N \in \mathbb{N}$, define the truncated random variables $|X| \wedge N$. Clearly, we have $\mathbf{E}[\mathbf{E}[|X| \wedge N | \mathcal{F}]^2] \leq N^2$. By what we have shown already (with X replaced by $|X| \wedge N$ and with $Y = 0 \in \mathcal{L}^2(\Omega, \mathcal{F}, \mathbf{P})$), and using the elementary inequality $a^2 \leq 2(a - b)^2 + 2b^2$, $a, b \in \mathbb{R}$, we infer

$$\begin{aligned} \mathbf{E}[\mathbf{E}[|X| \wedge N | \mathcal{F}]^2] &\leq 2\mathbf{E}[(|X| \wedge N - \mathbf{E}[|X| \wedge N | \mathcal{F}])^2] + 2\mathbf{E}[(|X| \wedge N)^2] \\ &\leq 4\mathbf{E}[(|X| \wedge N)^2] \leq 4\mathbf{E}[X^2]. \end{aligned}$$

By Theorem 8.14(ii) and (viii), we get $\mathbf{E}[|X| \wedge N | \mathcal{F}] \uparrow \mathbf{E}[|X| | \mathcal{F}]$ for $N \rightarrow \infty$. By the triangle inequality (Theorem 8.14(v)) and the monotone convergence theorem (Theorem 4.20), we conclude

$$\mathbf{E}[\mathbf{E}[X | \mathcal{F}]^2] \leq \mathbf{E}[\mathbf{E}[|X| | \mathcal{F}]^2] = \lim_{N \rightarrow \infty} \mathbf{E}[\mathbf{E}[|X| \wedge N | \mathcal{F}]^2] \leq 4\mathbf{E}[X^2] < \infty.$$

This completes the proof.

Example 8.18 Let $X, Y \in \mathcal{L}^1(\mathbf{P})$ be independent. Then

$$\mathbf{E}[X + Y | Y] = \mathbf{E}[X | Y] + \mathbf{E}[Y | Y] = \mathbf{E}[X] + Y.$$

Example 8.19 Let $X_1, \dots, X_N$ be independent with $\mathbf{E}[X_i] = 0, i = 1, \dots, N$. For $n = 1, \dots, N$, define $\mathcal{F}_n := \sigma(X_1, \dots, X_n)$ and $S_n := X_1 + \dots + X_n$. Then, for $n \geq m$,

$$\begin{aligned}
\mathbf{E}[S_n | \mathcal{F}_m] &= \mathbf{E}[X_1 | \mathcal{F}_m] + \dots + \mathbf{E}[X_n | \mathcal{F}_m] \\
&= X_1 + \dots + X_m + \mathbf{E}[X_{m+1}] + \dots + \mathbf{E}[X_n] \\
&= S_m
\end{aligned}$$

By Theorem 8.14(iv), since $\sigma(S_m) \subset \mathcal{F}_m$, we have

$$\mathbf{E}[S_n | S_m] = \mathbf{E}[\mathbf{E}[S_n | \mathcal{F}_m] | S_m] = \mathbf{E}[S_m | S_m] = S_m$$

Next we show Jensen's inequality for conditional expectations.

Theorem 8.20 (Jensen's inequality) Let $I \subset \mathbb{R}$ be an interval, let $\varphi : I \rightarrow \mathbb{R}$ be convex and let X be an I -valued random variable on $(\Omega, \mathcal{A}, \mathbf{P})$. Further, let $\mathbf{E}[|X|] < \infty$ and let $\mathcal{F} \subset \mathcal{A}$ be a σ -algebra. Then

$$\infty \geq \mathbf{E}[\varphi(X) | \mathcal{F}] \geq \varphi(\mathbf{E}[X | \mathcal{F}])$$

Proof For the existence of $\mathbf{E}[\varphi(X) | \mathcal{F}]$ with values in $(-\infty, \infty]$ note that $\varphi(X)^- \in \mathcal{L}^1(\mathbf{P})$ and see Remark 8.16. By Exercise 8.2.2, we have $\mathbf{E}[X | \mathcal{F}] \in I$ a.s., hence $\varphi(\mathbf{E}[X | \mathcal{F}])$ is well-defined.

(Recall from Definition 1.68 the jargon words "almost surely on A ".) Note that $X = \mathbf{E}[X | \mathcal{F}]$ on the event $\{\mathbf{E}[X | \mathcal{F}] \text{ is a boundary point of } I\}$; hence here the claim is trivial. Indeed, without loss of generality, assume 0 is the left boundary of I and $A := \{\mathbf{E}[X | \mathcal{F}] = 0\}$. As X assumes values in $I \subset [0, \infty)$, we have $0 \leq \mathbf{E}[X \mathbb{1}_A] = \mathbf{E}[\mathbf{E}[X | \mathcal{F}] \mathbb{1}_A] = 0$; hence $X \mathbb{1}_A = 0$. The case of a right boundary point is similar.

Hence now consider the event $B := \{\mathbf{E}[X | \mathcal{F}] \text{ is an interior point of } I\}$. For every interior point $x \in I$, let $D^+\varphi(x)$ be the maximal slope of a tangent of φ at x ; i.e., the maximal number t with $\varphi(y) \geq (y - x)t + \varphi(x)$ for all $y \in I$ (see Theorem 7.7).

For each $x \in I^\circ$, there exists a $\mathbf{P}$ -null set N_x such that, for every $\omega \in B \setminus N_x$, we have

$$\begin{aligned}
\mathbf{E}[\varphi(X) | \mathcal{F}](\omega) &\geq \varphi(x) + \mathbf{E}[D^+\varphi(x)(X - x) | \mathcal{F}](\omega) \\
&= \varphi(x) + D^+\varphi(x)(\mathbf{E}[X | \mathcal{F}](\omega) - x) =: \psi_\omega(x)
\end{aligned} \tag{8.8}$$

Let $V := \mathbb{Q} \cap I^\circ$. Then $N := \bigcup_{x \in V} N_x$ is a $\mathbf{P}$ -null set and (8.8) holds for every $\omega \in B \setminus N$ and every $x \in V$.

The map $x \mapsto D^+\varphi(x)$ is right continuous (by Theorem 7.7(iv)). Therefore $x \mapsto \psi_\omega(x)$ is also right continuous. Hence, for every $\omega \in B \setminus N$, we have

$$\begin{aligned}
\varphi(\mathbf{E}[X | \mathcal{F}](\omega)) &= \psi_\omega(\mathbf{E}[X | \mathcal{F}](\omega)) \\
&\leq \sup_{x \in I^\circ} \psi_\omega(x) = \sup_{x \in V} \psi_\omega(x) \leq \mathbf{E}[\varphi(X) | \mathcal{F}](\omega)
\end{aligned}$$

Corollary 8.21 Let $p \in [1, \infty]$ and let $\mathcal{F} \subset \mathcal{A}$ be a sub- σ -algebra. Then the map

$$\mathcal{L}^p(\Omega, \mathcal{A}, \mathbf{P}) \rightarrow \mathcal{L}^p(\Omega, \mathcal{F}, \mathbf{P}), \quad X \mapsto \mathbf{E}[X | \mathcal{F}],$$

is a contraction (that is, $\|\mathbf{E}[X | \mathcal{F}]\|_p \leq \|X\|_p$) and thus continuous. Hence, for $X, X_1, X_2, \dots \in \mathcal{L}^p(\Omega, \mathcal{A}, \mathbf{P})$ with $\|X_n - X\|_p \xrightarrow{n \rightarrow \infty} 0$,

$$\|\mathbf{E}[X_n | \mathcal{F}] - \mathbf{E}[X | \mathcal{F}]\|_p \xrightarrow{n \rightarrow \infty} 0.$$

Proof For $p \in [1, \infty)$, use Jensen's inequality with $\varphi(x) = |x|^p$. For $p = \infty$, note that $|\mathbf{E}[X | \mathcal{F}]| \leq \mathbf{E}[|X| | \mathcal{F}] \leq \mathbf{E}[\|X\|_\infty | \mathcal{F}] = \|X\|_\infty$.

Corollary 8.22 Let $(X_i, i \in I)$ be uniformly integrable and let $(\mathcal{F}_j, j \in J)$ be a family of sub- σ -algebras of $\mathcal{A}$. Define $X_{i,j} := \mathbf{E}[X_i | \mathcal{F}_j]$. Then $(X_{i,j}, (i, j) \in I \times J)$ is uniformly integrable. In particular, for $X \in \mathcal{L}^1(\mathbf{P})$, the family $(\mathbf{E}[X | \mathcal{F}_j], j \in J)$ is uniformly integrable.

Proof By Theorem 6.19, there exists a monotone increasing convex function f with the property that $f(x)/x \rightarrow \infty, x \rightarrow \infty$ and $L := \sup_{i \in I} \mathbf{E}[f(|X_i|)] < \infty$. Then $x \mapsto f(|x|)$ is convex; hence, by Jensen's inequality,

$$\mathbf{E}[f(|X_{i,j}|)] = \mathbf{E}[f(|\mathbf{E}[X_i | \mathcal{F}_j]|)] \leq L < \infty$$

Thus $(X_{i,j}, (i, j) \in I \times J)$ is uniformly integrable by Theorem 6.19.

Example 8.23 Let μ and ν be finite measures with $\nu \ll \mu$. Let $f = d\nu/d\mu$ be the Radon-Nikodym derivative and let $I = \{\mathcal{F} \subset \mathcal{A} : \mathcal{F} \text{ is a } \sigma\text{-algebra}\}$. Consider the measures $\mu|_{\mathcal{F}}$ and $\nu|_{\mathcal{F}}$ that are restricted to $\mathcal{F}$. Then $\nu|_{\mathcal{F}} \ll \mu|_{\mathcal{F}}$ (since in $\mathcal{F}$ there are fewer μ -null sets); hence the Radon-Nikodym derivative $f_{\mathcal{F}} := d\nu|_{\mathcal{F}}/d\mu|_{\mathcal{F}}$ exists. Then $(f_{\mathcal{F}} : \mathcal{F} \in I)$ is uniformly integrable (with respect to μ). (For finite σ -algebras $\mathcal{F}$, this was shown in Example 7.39.) Indeed, let $\mathbf{P} = \mu/\mu(\Omega)$ and $\mathbf{Q} = \nu/\mu(\Omega)$. Then $f_{\mathcal{F}} = d\mathbf{Q}|_{\mathcal{F}}/d\mathbf{P}|_{\mathcal{F}}$. For any $F \in \mathcal{F}$, we thus have $\mathbf{E}[f_{\mathcal{F}} \mathbb{1}_F] = \int_F f_{\mathcal{F}} d\mathbf{P} = \mathbf{Q}(F) = \int_F f d\mathbf{P} = \mathbf{E}[f \mathbb{1}_F]$; hence $f_{\mathcal{F}} = \mathbf{E}[f | \mathcal{F}]$. By the preceding corollary, $(f_{\mathcal{F}} : \mathcal{F} \in I)$ is uniformly integrable with respect to $\mathbf{P}$ and thus also with respect to μ .

Exercise 8.2.1 Show the assertions of Remark 8.16.

Exercise 8.2.2 Let $I \subset \mathbb{R}$ be an arbitrary interval and let $X \in \mathcal{L}^1(\Omega, \mathcal{A}, \mathbf{P})$ be a random variable such that $X \in I$ a.s. For $\mathcal{F} \subset \mathcal{A}$, show that $\mathbf{E}[X | \mathcal{F}] \in I$ a.s. Is this statement still true if we require only $X^- \in \mathcal{L}^1(\Omega, \mathcal{A}, \mathbf{P})$ instead of $X \in \mathcal{L}^1(\Omega, \mathcal{A}, \mathbf{P})$?

Exercise 8.2.3 (Bayes' formula) Let $A \in \mathcal{A}$ and $B \in \mathcal{F} \subset \mathcal{A}$. Show that

$$\mathbf{P}[B | A] = \frac{\int_B \mathbf{P}[A | \mathcal{F}] d\mathbf{P}}{\int \mathbf{P}[A | \mathcal{F}] d\mathbf{P}}$$

If $\mathcal{F}$ is generated by pairwise disjoint sets $B_1, B_2, \dots$, then this is exactly Bayes' formula of Theorem 8.7.

Exercise 8.2.4 Give an example for $\mathbf{E}[\mathbf{E}[X | \mathcal{F}] | \mathcal{G}] \neq \mathbf{E}[\mathbf{E}[X | \mathcal{G}] | \mathcal{F}]$.

Exercise 8.2.5 Show the conditional Markov inequality: For monotone increasing $f : [0, \infty) \rightarrow [0, \infty)$ and $\varepsilon > 0$ with $f(\varepsilon) > 0$,

$$\mathbf{P}[|X| \geq \varepsilon | \mathcal{F}] \leq \frac{\mathbf{E}[f(|X|) | \mathcal{F}]}{f(\varepsilon)}$$

Exercise 8.2.6 Show the conditional Cauchy-Schwarz inequality: For square integrable random variables X, Y ,

$$\mathbf{E}[XY | \mathcal{F}]^2 \leq \mathbf{E}[X^2 | \mathcal{F}] \mathbf{E}[Y^2 | \mathcal{F}].$$

Exercise 8.2.7 Let $X_1, \dots, X_n$ be integrable i.i.d. random variables. Let $S_n = X_1 + \dots + X_n$. Show that

$$\mathbf{E}[X_i | S_n] = \frac{1}{n} S_n \quad \text{for every } i = 1, \dots, n$$

Exercise 8.2.8 Let X_1 and X_2 be independent and exponentially distributed with parameter $\theta > 0$. Compute $\mathbf{E}[X_1 \wedge X_2 | X_1]$.

Exercise 8.2.9 Let X and Y be real random variables with joint density f and let $h : \mathbb{R} \rightarrow \mathbb{R}$ be measurable with $\mathbf{E}[|h(X)|] < \infty$. Denote by λ the Lebesgue measure on $\mathbb{R}$.

(i) Show that almost surely

$$\mathbf{E}[h(X) | Y] = \frac{\int h(x)f(x, Y)\lambda(dx)}{\int f(x, Y)\lambda(dx)}$$

(ii) Let X and Y be independent and $\exp_\theta$ -distributed for some $\theta > 0$. Compute $\mathbf{E}[X | X + Y]$ and $\mathbf{P}[X \leq x | X + Y]$ for $x \geq 0$.

5.8.3 8.3 Regular Conditional Distribution

Let X be a random variable with values in a measurable space $(E, \mathcal{E})$. With our machinery, so far we can define the conditional probability $\mathbf{P}[A | X]$ for fixed $A \in \mathcal{A}$ only. However, we would like to define for every $x \in E$ a probability measure $\mathbf{P}[\cdot | X = x]$ such that for any $A \in \mathcal{A}$, we have $\mathbf{P}[A | X] = \mathbf{P}[A | X = x]$ on $\{X = x\}$. In this section, we show how to do this.

For example, we are interested in a two-stage random experiment. At the first stage, we manipulate a coin at random such that the probability of a success (i.e., "head") is X . At the second stage, we toss the coin n times independently with outcomes $Y_1, \dots, Y_n$. Hence the "conditional distribution of $(Y_1, \dots, Y_n)$ given $\{X = x\}$ " should be $(\text{Ber}_x)^{\otimes n}$.

Let X be as above and let Z be a $\sigma(X)$ -measurable real random variable. By the factorization lemma (Corollary 1.97 with $f = X$ and $g = Z$), there is a map $\varphi : E \rightarrow \mathbb{R}$ such that

$$\varphi \text{ is } \mathcal{E} - \mathcal{B}(\mathbb{R})\text{-measurable} \quad \text{and} \quad \varphi(X) = Z. \quad (8.9)$$

If X is surjective, then φ is determined uniquely. In this case, we denote $Z \circ X^{-1} := \varphi$ (even if the inverse map X^{-1} itself does not exist).

Definition 8.24 Let $Y \in \mathcal{L}^1(\mathbf{P})$ and $X : (\Omega, \mathcal{A}) \rightarrow (E, \mathcal{E})$. We define the conditional expectation of Y given $X = x$ by $\mathbf{E}[Y | X = x] := \varphi(x)$, where φ is the function from (8.9) with $Z = \mathbf{E}[Y | X]$.

Analogously, define $\mathbf{P}[A | X = x] = \mathbf{E}[\mathbb{1}_A | X = x]$ for $A \in \mathcal{A}$.

For a fixed set $B \in \mathcal{A}$ with $\mathbf{P}[B] > 0$, the conditional probability $\mathbf{P}[\cdot | B]$ is a probability measure. Is this true also for $\mathbf{P}[\cdot | X = x]$? The question is a bit tricky since for every given $A \in \mathcal{A}$, the expression $\mathbf{P}[A | X = x]$ is defined for almost all x only; that is, up to x in a null set that may, however, depend on A . Since there are uncountably many $A \in \mathcal{A}$ in general, we could not simply unite all the exceptional sets for any A . However, if the σ -algebra $\mathcal{A}$ can be approximated by countably many A sufficiently well, then there is hope.

Our first task is to give precise definitions. Then we present the theorem that justifies our hope.

Definition 8.25 (Transition kernel, Markov kernel) Let $(\Omega_1, \mathcal{A}_1), (\Omega_2, \mathcal{A}_2)$ be measurable spaces. A map $\kappa : \Omega_1 \times \mathcal{A}_2 \rightarrow [0, \infty]$ is called a (σ) -finite transition kernel (from Ω_1 to Ω_2) if:

- (i) $\omega_1 \mapsto \kappa(\omega_1, A_2)$ is $\mathcal{A}_1$ -measurable for any $A_2 \in \mathcal{A}_2$.
- (ii) $A_2 \mapsto \kappa(\omega_1, A_2)$ is a (σ) -finite measure on $(\Omega_2, \mathcal{A}_2)$ for any $\omega_1 \in \Omega_1$.

If in (ii) the measure is a probability measure for all $\omega_1 \in \Omega_1$, then κ is called a stochastic kernel or a Markov kernel. If in (ii) we also have $\kappa(\omega_1, \Omega_2) \leq 1$ for any $\omega_1 \in \Omega_1$, then κ is called sub-Markov or substochastic.

Remark 8.26 It is sufficient to check property (i) in Definition 8.25 for sets A_2 from a π -system $\mathcal{E}$ that generates $\mathcal{A}_2$ and that either contains Ω_2 or a sequence $E_n \uparrow \Omega_2$. Indeed, in this case,

$$\mathcal{D} := \{A_2 \in \mathcal{A}_2 : \omega_1 \mapsto \kappa(\omega_1, A_2) \text{ is } \mathcal{A}_1\text{-measurable}\}$$

is a λ -system (exercise!). Since $\mathcal{E} \subset \mathcal{D}$, by the $\pi - \lambda$ theorem (Theorem 1.19), $\mathcal{D} = \sigma(\mathcal{E}) = \mathcal{A}_2$.

Example 8.27

(i) Let $(\Omega_1, \mathcal{A}_1)$ and $(\Omega_2, \mathcal{A}_2)$ be discrete measurable spaces and let $(K_{ij})_{\substack{i \in \Omega_1 \\ j \in \Omega_2}}$ be a matrix with nonnegative entries and finite row sums

$$K_i := \sum_{j \in \Omega_2} K_{ij} < \infty \quad \text{for } i \in \Omega_1.$$

Then we can define a finite transition kernel from Ω_1 to Ω_2 by $\kappa(i, A) = \sum_{j \in A} K_{ij}$. κ is stochastic if $K_i = 1$ for all $i \in \Omega_1$. It is substochastic if $K_i \leq 1$ for all $i \in \Omega_1$.

(ii) If μ_2 is a finite measure on Ω_2 , then $\kappa(\omega_1, \cdot) \equiv \mu_2$ is a finite transition kernel.

(iii) $\kappa(x, \cdot) = \text{Poi}_x$ is a stochastic kernel from $[0, \infty)$ to $\mathbb{N}_0$ (note that $x \mapsto \text{Poi}_x(A)$ is continuous and hence measurable for all $A \subset \mathbb{N}_0$).

(iv) Let μ be a distribution on $\mathbb{R}^n$ and let X be a random variable with $\mathbf{P}_X = \mu$. Then $\kappa(x, \cdot) = \mathbf{P}[X + x \in \cdot] = \delta_x * \mu$ defines a stochastic kernel from $\mathbb{R}^n$ to $\mathbb{R}^n$. Indeed, the sets $(-\infty, y], y \in \mathbb{R}^n$ form an $\cap$ -stable generator of $\mathcal{B}(\mathbb{R}^n)$ and $x \mapsto \kappa(x, (-\infty, y]) = \mu((-\infty, y - x])$ is left continuous and hence measurable. Hence, by Remark 8.26, $x \mapsto \kappa(x, A)$ is measurable for all $A \in \mathcal{B}(\mathbb{R}^n)$.

Definition 8.28 Let Y be a random variable with values in a measurable space $(E, \mathcal{E})$ and let $\mathcal{F} \subset \mathcal{A}$ be a sub- σ -algebra. A stochastic kernel $\kappa_{Y, \mathcal{F}}$ from $(\Omega, \mathcal{F})$ to $(E, \mathcal{E})$ is called a regular conditional distribution of Y given $\mathcal{F}$ if

$$\kappa_{Y, \mathcal{F}}(\omega, B) = \mathbf{P}[\{Y \in B\} \mid \mathcal{F}](\omega)$$

for $\mathbf{P}$ -almost all $\omega \in \Omega$ and for all $B \in \mathcal{E}$; that is, if

$$\int \mathbb{1}_B(Y) \mathbb{1}_A d\mathbf{P} = \int \kappa_{Y, \mathcal{F}}(\cdot, B) \mathbb{1}_A d\mathbf{P} \quad \text{for all } A \in \mathcal{F}, B \in \mathcal{E} \quad (8.10)$$

Consider the special case where $\mathcal{F} = \sigma(X)$ for a random variable X (with values in an arbitrary measurable space $(E', \mathcal{E}')$). Then the stochastic kernel

$$(x, A) \mapsto \kappa_{Y, X}(x, A) = \mathbf{P}[\{Y \in A\} \mid X = x] = \kappa_{Y, \sigma(X)}(X^{-1}(x), A)$$

(the function from the factorization lemma with an arbitrary value for $x \notin X(\Omega)$) is called a regular conditional distribution of Y given X .

Theorem 8.29 (Regular conditional distributions in $\mathbb{R}$) Let $Y : (\Omega, \mathcal{A}) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$ be real-valued. Then there exists a regular conditional distribution $\kappa_{Y, \mathcal{F}}$ of Y given $\mathcal{F}$.

Proof The strategy of the proof consists in constructing a measurable version of the distribution function of the conditional distribution of Y by first defining it for rational values (up to a null set) and then extending it to the real numbers.

For $r \in \mathbb{Q}$, let $F(r, \cdot)$ be a version of the conditional probability $\mathbf{P}[Y \in (-\infty, r] \mid \mathcal{F}]$. For $r \leq s$, clearly $\mathbb{1}_{\{Y \in (-\infty, r]\}} \leq \mathbb{1}_{\{Y \in (-\infty, s]\}}$. Hence, by Theorem 8.14(ii) (monotonicity of the conditional expectation), there is a null set $A_{r,s} \in \mathcal{F}$ with

$$F(r, \omega) \leq F(s, \omega) \quad \text{for all } \omega \in \Omega \setminus A_{r,s}. \quad (8.11)$$

By Theorem 8.14(viii) (dominated convergence), there are null sets $(B_r)_{r \in \mathbb{Q}} \in \mathcal{F}$ and $C \in \mathcal{F}$ such that

$$\lim_{n \rightarrow \infty} F\left(r + \frac{1}{n}, \omega\right) = F(r, \omega) \quad \text{for all } \omega \in \Omega \setminus B_r \quad (8.12)$$

as well as

$$\inf_{n \in \mathbb{N}} F(-n, \omega) = 0 \quad \text{and} \quad \sup_{n \in \mathbb{N}} F(n, \omega) = 1 \quad \text{for all } \omega \in \Omega \setminus C \quad (8.13)$$

Let $N := \left(\bigcup_{r,s \in \mathbb{Q}} A_{r,s} \right) \cup \left(\bigcup_{r \in \mathbb{Q}} B_r \right) \cup C$. For $\omega \in \Omega \setminus N$, define

$$\tilde{F}(z, \omega) := \inf \{ F(r, \omega) : r \in \mathbb{Q}, r > z \} \quad \text{for all } z \in \mathbb{R}$$

By construction, $\tilde{F}(\cdot, \omega)$ is monotone increasing and right continuous. By (8.11) and (8.12), we have

$$\tilde{F}(z, \omega) = F(z, \omega) \quad \text{for all } z \in \mathbb{Q} \text{ and } \omega \in \Omega \setminus N. \quad (8.14)$$

Therefore, by (8.13), $\tilde{F}(\cdot, \omega)$ is a distribution function for any $\omega \in \Omega \setminus N$. For $\omega \in N$, define $\tilde{F}(\cdot, \omega) = F_0$, where F_0 is an arbitrary but fixed distribution function.

For any $\omega \in \Omega$, let $\kappa(\omega, \cdot)$ be the probability measure on $(\Omega, \mathcal{A})$ with distribution function $\tilde{F}(\cdot, \omega)$. Then, for $r \in \mathbb{Q}$ and $B = (-\infty, r]$,

$$\omega \mapsto \kappa(\omega, B) = F(r, \omega) \mathbb{1}_{N^c}(\omega) + F_0(r) \mathbb{1}_N(\omega) \quad (8.15)$$

is $\mathcal{F}$ -measurable. Now $\{(-\infty, r], r \in \mathbb{Q}\}$ is a π -system that generates $\mathcal{B}(\mathbb{R})$. By Remark 8.26, measurability holds for all $B \in \mathcal{B}(\mathbb{R})$ and hence κ is identified as a stochastic kernel.

We still have to show that κ is a version of the conditional distribution. For $A \in \mathcal{F}$, $r \in \mathbb{Q}$ and $B = (-\infty, r]$, by (8.15),

$$\int_A \kappa(\omega, B) \mathbf{P}[d\omega] = \int_A \mathbf{P}[Y \in B \mid \mathcal{F}] d\mathbf{P} = \mathbf{P}[A \cap \{Y \in B\}]$$

As functions of B , both sides are finite measures on $\mathcal{B}(\mathbb{R})$ that coincide on the $\cap$ stable generator $\{(-\infty, r], r \in \mathbb{Q}\}$. By the uniqueness theorem (Lemma 1.42), we thus have equality for all $B \in \mathcal{B}(\mathbb{R})$. Hence $\mathbf{P}$ -a.s. $\kappa(\cdot, B) = \mathbf{P}[Y \in B \mid \mathcal{F}]$ and thus $\kappa = \kappa_{Y, \mathcal{F}}$.

Example 8.30 Let Z_1, Z_2 be independent Poisson random variables with parameters $\lambda_1, \lambda_2 \geq 0$. One can show (exercise!) that (with $Y = Z_1$ and $X = Z_1 + Z_2$)

$$\mathbf{P}[Z_1 = k \mid Z_1 + Z_2 = n] = b_{n,p}(k) \quad \text{for } k = 0, \dots, n,$$

where $p = \frac{\lambda_1}{\lambda_1 + \lambda_2}$.

This example could still be treated by elementary means. The full strength of the result is displayed in the following examples.

Example 8.31 Let X and Y be real random variables with joint density f (with respect to Lebesgue measure λ^2 on $\mathbb{R}^2$). For $x \in \mathbb{R}$, define

$$f_X(x) = \int_{\mathbb{R}} f(x, y) \lambda(dy)$$

Clearly, $f_X(x) > 0$ for $\mathbf{P}_X$ -a.a. $x \in \mathbb{R}$ and f_X^{-1} is the density of the absolutely continuous part of the Lebesgue measure λ with respect to $\mathbf{P}_X$. The regular conditional distribution of Y given X has density

$$\frac{\mathbf{P}[Y \in dy \mid X = x]}{dy} = f_{Y|X}(x, y) := \frac{f(x, y)}{f_X(x)} \quad \text{for } \mathbf{P}_X[dx]\text{-a.a. } x \in \mathbb{R}. \quad (8.16)$$

Indeed, by Fubini's theorem (Theorem 14.16), the map $x \mapsto \int_B f_{Y|X}(x, y) \lambda(dy)$ is measurable for all $B \in \mathcal{B}(\mathbb{R})$ and for $A, B \in \mathcal{B}(\mathbb{R})$, we have

$$\begin{aligned}
& \int_A \mathbf{P}[X \in dx] \int_B f_{Y|X}(x, y) \lambda(dy) \\
&= \int_A \mathbf{P}[X \in dx] f_X(x)^{-1} \int_B f(x, y) \lambda(dy) \\
&= \int_A \lambda(dx) \int_B f(x, y) \lambda(dy) \\
&= \int_{A \times B} f d\lambda^2 = \mathbf{P}[X \in A, Y \in B]
\end{aligned}$$

Example 8.32 Let $\mu_1, \mu_2 \in \mathbb{R}$, $\sigma_1, \sigma_2 > 0$ and let Z_1, Z_2 be independent and $\mathcal{N}_{\mu_i, \sigma_i^2}$ -distributed ($i = 1, 2$). Then there exists a regular conditional distribution

$$\mathbf{P}[Z_1 \in \cdot \mid Z_1 + Z_2 = x] \quad \text{for } x \in \mathbb{R}.$$

If we define $X = Z_1 + Z_2$ and $Y = Z_1$, then $(X, Y) \sim \mathcal{N}_{\mu, \Sigma}$ is bivariate normally distributed with covariance matrix

$$\Sigma := \begin{pmatrix} \sigma_1^2 + \sigma_2^2 & \sigma_1^2 \\ \sigma_1^2 & \sigma_1^2 \end{pmatrix}$$

and with

$$\mu := \begin{pmatrix} \mu_1 + \mu_2 \\ \mu_1 \end{pmatrix}$$

Note that

$$\Sigma^{-1} = \begin{pmatrix} \sigma_1^2 & \sigma_2^2 \end{pmatrix}^{-1} \begin{pmatrix} \sigma_1^2 & -\sigma_1^2 \\ -\sigma_1^2 & \sigma_1^2 + \sigma_2^2 \end{pmatrix} = \begin{pmatrix} \sigma_1^2 & \sigma_2^2 \end{pmatrix}^{-1} B^T B,$$

where $B = \begin{pmatrix} \sigma_1 & -\sigma_1 \\ 0 & \sigma_2 \end{pmatrix}$. Hence (X, Y) has the density (see Example 1.105(ix))

$$\begin{aligned}
f(x, y) &= \det(2\pi\Sigma)^{-1/2} \exp \left(-\frac{1}{2\sigma_1^2\sigma_2^2} \left\| B \begin{pmatrix} x - (\mu_1 + \mu_2) \\ y - \mu_1 \end{pmatrix} \right\|^2 \right) \\
&= (4\pi^2\sigma_1^2\sigma_2^2)^{-1/2} \exp \left(-\frac{\sigma_1^2(y - (x - \mu_2))^2 + \sigma_2^2(y - \mu_1)^2}{2\sigma_1^2\sigma_2^2} \right) \\
&= C_x \exp \left(-(y - \mu_x)^2 / 2\sigma_x^2 \right)
\end{aligned}$$

Here C_x is a normalising constant and

$$\mu_x = \mu_1 + \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2} (x - \mu_1 - \mu_2) \quad \text{and} \quad \sigma_x^2 = \frac{\sigma_1^2\sigma_2^2}{\sigma_1^2 + \sigma_2^2}.$$

By (8.16), $\mathbf{P}[Z_1 \in \cdot \mid Z_1 + Z_2 = x]$ has the density

$$y \mapsto f_{Y|X}(x, y) = \frac{C_x}{f_X(x)} \exp \left(-\frac{(y - \mu_x)^2}{2\sigma_x^2} \right),$$

hence

$$\mathbf{P}[Z_1 \in \cdot \mid Z_1 + Z_2 = x] = \mathcal{N}_{\mu_x, \sigma_x^2} \quad \text{for almost all } x \in \mathbb{R}$$

Example 8.33 If X and Y are independent real random variables, then for $\mathbf{P}_X$ almost all $x \in \mathbb{R}$

$$\mathbf{P}[X + Y \in \cdot \mid X = x] = \delta_x * \mathbf{P}_Y$$

The situation is not completely satisfying as we have made the very restrictive assumption that Y is real-valued. Originally we were also interested in the situation where Y takes values in $\mathbb{R}^n$ or in even more general spaces. We now extend the result to a larger class of ranges for Y .

Definition 8.34 Two measurable spaces $(E, \mathcal{E})$ and $(E', \mathcal{E}')$ are called isomorphic if there exists a bijective map $\varphi : E \rightarrow E'$ such that φ is $\mathcal{E} - \mathcal{E}'$ -measurable and the inverse map φ^{-1} is $\mathcal{E}' - \mathcal{E}$ -measurable. Then we say that φ is an isomorphism of measurable spaces. If in addition μ and μ' are measures on $(E, \mathcal{E})$ and $(E', \mathcal{E}')$ and if $\mu' = \mu \circ \varphi^{-1}$, then φ is an isomorphism of measure spaces, and the measure spaces $(E, \mathcal{E}, \mu)$ and $(E', \mathcal{E}', \mu')$ are called isomorphic.

Definition 8.35 A measurable space $(E, \mathcal{E})$ is called a Borel space if there exists a Borel set $B \in \mathcal{B}(\mathbb{R})$ such that $(E, \mathcal{E})$ and $(B, \mathcal{B}(B))$ are isomorphic.

A separable topological space whose topology is induced by a complete metric is called a Polish space. In particular, $\mathbb{R}^d, \mathbb{Z}^d, \mathbb{R}^{\mathbb{N}}, (C([0, 1]), \|\cdot\|_{\infty})$ and so forth are Polish. Closed subsets of Polish spaces are again Polish. We come back to Polish spaces in the context of convergence of measures in Chapter 13. Without proof, we present the following topological result (see, e.g., [37, Theorem 13.1.1]).

Theorem 8.36 Let E be a Polish space with Borel σ -algebra $\mathcal{E}$. Then $(E, \mathcal{E})$ is a Borel space.

Theorem 8.37 (Regular conditional distribution) Let $\mathcal{F} \subset \mathcal{A}$ be a sub- σ -algebra. Let Y be a random variable with values in a Borel space $(E, \mathcal{E})$ (hence, for example, E Polish, $E = \mathbb{R}^d, E = \mathbb{R}^{\infty}, E = C([0, 1])$, etc.). Then there exists a regular conditional distribution $\kappa_{Y, \mathcal{F}}$ of Y given $\mathcal{F}$.

Proof Let $B \in \mathcal{B}(\mathbb{R})$ and let $\varphi : E \rightarrow B$ be an isomorphism of measurable spaces. By Theorem 8.29, we obtain the regular conditional distribution $\kappa_{Y', \mathcal{F}}$ of the real random variable $Y' = \varphi \circ Y$. Now define $\kappa_{Y, \mathcal{F}}(\omega, A) = \kappa_{Y', \mathcal{F}}(\omega, \varphi(A))$ for $A \in \mathcal{E}$.

To conclude, we pick up again the example with which we started. Now we can drop the quotation marks from the statement and write it down formally. Hence, let X be uniformly distributed on $[0, 1]$. Given $X = x$, let $(Y_1, \dots, Y_n)$ be independent and Ber_x -distributed. Define $Y = (Y_1, \dots, Y_n)$. By Theorem 8.37 (with $E = \{0, 1\}^n \subset \mathbb{R}^n$), a regular conditional distribution exists:

$$\kappa_{Y, X}(x, \cdot) = \mathbf{P}[Y \in \cdot \mid X = x] \quad \text{for } x \in [0, 1].$$

Indeed, for almost all $x \in [0, 1]$,

$$\mathbf{P}[Y \in \cdot \mid X = x] = (\text{Ber}_x)^{\otimes n}$$

Theorem 8.38 Let X be a random variable on $(\Omega, \mathcal{A}, \mathbf{P})$ with values in a Borel space $(E, \mathcal{E})$. Let $\mathcal{F} \subset \mathcal{A}$ be a σ -algebra and let $\kappa_{X, \mathcal{F}}$ be a regular conditional distribution of X given $\mathcal{F}$. Further, let $f : E \rightarrow \mathbb{R}$ be measurable and $\mathbf{E}[|f(X)|] < \infty$. Then

$$\mathbf{E}[f(X) \mid \mathcal{F}](\omega) = \int f(x) \kappa_{X, \mathcal{F}}(\omega, dx) \quad \text{for } \mathbf{P}\text{-almost all } \omega \quad (8.17)$$

Proof We check that the right-hand side in (8.17) has the properties of the conditional expectation.

It is enough to consider the case $f \geq 0$. By approximating f by simple functions, we see that the right-hand side in (8.17) is $\mathcal{F}$ -measurable (see Lemma 14.20 for a formal argument). Hence, by Theorem 1.96, there exist sets $A_1, A_2, \dots \in \mathcal{E}$ and numbers $\alpha_1, \alpha_2, \dots \geq 0$ such that

$$g_n := \sum_{i=1}^n \alpha_i \mathbb{1}_{A_i} \xrightarrow{n \rightarrow \infty} f.$$

Now, for any $n \in \mathbb{N}$ and $B \in \mathcal{F}$,

$$\begin{aligned} \mathbf{E}[g_n(X) \mathbb{1}_B] &= \sum_{i=1}^n \alpha_i \mathbf{P}[\{X \in A_i\} \cap B] \\ &= \sum_{i=1}^n \alpha_i \int_B \mathbf{P}[\{X \in A_i\} \mid \mathcal{F}] \mathbf{P}[d\omega] \\ &= \sum_{i=1}^n \alpha_i \int_B \kappa_{X, \mathcal{F}}(\omega, A_i) \mathbf{P}[d\omega] \\ &= \int_B \sum_{i=1}^n \alpha_i \kappa_{X, \mathcal{F}}(\omega, A_i) \mathbf{P}[d\omega] \\ &= \int_B \left(\int g_n(x) \kappa_{X, \mathcal{F}}(\omega, dx) \right) \mathbf{P}[d\omega] \end{aligned}$$

By the monotone convergence theorem, for almost all ω , the inner integral converges to $\int f(x) \kappa_{X, \mathcal{F}}(\omega, dx)$. Applying the monotone convergence theorem once more, we get

$$\mathbf{E}[f(X) \mathbb{1}_B] = \lim_{n \rightarrow \infty} \mathbf{E}[g_n(X) \mathbb{1}_B] = \int_B \int f(x) \kappa_{X, \mathcal{F}}(\omega, dx) \mathbf{P}[d\omega]$$

Exercise 8.3.1 Let $(E, \mathcal{E})$ be a Borel space and let μ be an atom-free measure (that is, $\mu(\{x\}) = 0$ for any $x \in E$). Show that for any $A \in \mathcal{E}$ and any $n \in \mathbb{N}$, there exist pairwise disjoint sets $A_1, \dots, A_n \in \mathcal{E}$ with $\biguplus_{k=1}^n A_k = A$ and $\mu(A_k) = \mu(A)/n$ for any $k = 1, \dots, n$.

Exercise 8.3.2 Let $p, q \in (1, \infty)$ with $\frac{1}{p} + \frac{1}{q} = 1$ and let $X \in \mathcal{L}^p(\mathbf{P})$ and $Y \in \mathcal{L}^q(\mu)$. Let $\mathcal{F} \subset \mathcal{A}$ be a σ -algebra. Use the preceding theorem to show the conditional version of Hölder's inequality:

$$\mathbf{E}[|XY| \mid \mathcal{F}] \leq \mathbf{E}[|X|^p \mid \mathcal{F}]^{1/p} \mathbf{E}[|Y|^q \mid \mathcal{F}]^{1/q} \quad \text{almost surely.}$$

Exercise 8.3.3 Assume the random variable (X, Y) is uniformly distributed on the disc $B := \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$ and on $[-1, 1]^2$, respectively.

(i) In both cases, determine the conditional distribution of Y given $X = x$.

(ii) Let $R := \sqrt{X^2 + Y^2}$ and $\Theta = \arctan(Y/X)$. In both cases, determine the conditional distribution of Θ given $R = r$.

Exercise 8.3.4 Let $A \subset \mathbb{R}^n$ be a Borel measurable set of finite Lebesgue measure $\lambda(A) \in (0, \infty)$ and let X be uniformly distributed on A (see Example 1.75). Let $B \subset A$ be measurable with $\lambda(B) > 0$. Show that the conditional distribution of X given $\{X \in B\}$ is the uniform distribution on B .

Exercise 8.3.5 (Borel's paradox) Consider the Earth as a ball (as widely accepted nowadays). Let X be a random point that is uniformly distributed on the surface. Let Θ be the longitude and let Φ be the latitude of X . A little differently from the usual convention, assume that Θ takes values in $[0, \pi)$ and Φ in $[-\pi, \pi)$. Hence, for fixed Θ , a complete great circle is described when Φ runs through its domain. Now, given Θ , is Φ uniformly distributed on $[-\pi, \pi)$? One could conjecture that any point on the great circle is equally likely. However, this is not the case! If we thicken the great circle slightly such that its longitudes range from Θ to $\Theta + \varepsilon$ (for a small ε), on the equator it is thicker (measured in meters) than at the poles. If we let $\varepsilon \rightarrow 0$, intuitively we should get the conditional probabilities as proportional to the thickness

(in meters).

- (i) Show that $\mathbf{P}[\{\Phi \in \cdot\} \mid \Theta = \theta]$ for almost all θ has the density $\frac{1}{4}|\cos(\phi)|$ for $\phi \in [-\pi, \pi)$.
(ii) Show that $\mathbf{P}[\{\Theta \in \cdot\} \mid \Phi = \phi] = \mathcal{U}_{[0, \pi)}$ for almost all ϕ .

Hint: Show that Θ and Φ are independent, and compute the distributions of Θ and Φ .

Exercise 8.3.6 (Rejection sampling for generating random variables) Let E be a countable set and let P and Q be probability measures on E . Assume there is a $c > 0$ with

$$f(e) := \frac{Q(\{e\})}{P(\{e\})} \leq c \quad \text{for all } e \in E \text{ with } P(\{e\}) > 0$$

Let $X_1, X_2, \dots$ be independent random variables with distribution P . Let $U_1, U_2, \dots$ be i.i.d. random variables that are independent of $X_1, X_2, \dots$ and that are uniformly distributed on $[0, 1]$. Let N be the smallest (random) nonnegative integer n such that $U_n \leq f(X_n)/c$ and define $Y := X_N$.

Show that Y has distribution Q .

Remark. This method for generating random variables with a given distribution Q is called rejection sampling, as it can also be described as follows. The random variable X_1 is a proposal for the value of Y . This proposal is accepted with probability $f(X_1)/c$ and is rejected otherwise. If the first proposal is rejected, the game starts afresh with proposal X_2 and so on.

Exercise 8.3.7 Let E be a Polish space and let $P, Q \in \mathcal{M}_1(\mathbb{R})$. Let $c > 0$ with $f := \frac{dQ}{dP} \leq cP$ -almost surely. Show the statement analogous to Exercise 8.3.6.

Exercise 8.3.8 Show that $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ and $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ are isomorphic. Conclude that every Borel set $B \in \mathcal{B}(\mathbb{R}^n)$ is a Borel space.

5.9 9 Martingales

One of the most important concepts of modern probability theory is the martingale, which formalizes the notion of a fair game. In this chapter, we first lay the foundations for the treatment of general stochastic processes. We then introduce martingales and the discrete stochastic integral. We close with an application to a model from mathematical finance.

5.9.1 9.1 Processes, Filtrations, Stopping Times

We introduce the fundamental technical terms for the investigation of stochastic processes (including martingales). In order to be able to recycle the terms later in a more general context, we go for greater generality than is necessary for the treatment of martingales only.

In the following, let (E, τ) be a Polish space with Borel σ -algebra $\mathcal{E}$. Further, let $(\Omega, \mathcal{F}, \mathbf{P})$ be a probability space and let $I \subset \mathbb{R}$ be arbitrary. We are mostly interested in the cases $I = \mathbb{N}_0, I = \mathbb{Z}, I = [0, \infty)$ and I an interval.

Definition 9.1 (Stochastic process) Let $I \subset \mathbb{R}$. A family of random variables $X = (X_t, t \in I)$ (on $(\Omega, \mathcal{F}, \mathbf{P})$) with values in $(E, \mathcal{E})$ is called a stochastic process with index set (or time set) I and range E .

Remark 9.2 Sometimes families of random variables with more general index sets are called stochastic processes. We come back to this with the Poisson point process in Chapter 24.

Remark 9.3 Following a certain tradition, we will often denote a stochastic process by $X = (X_t)_{t \in I}$ if we want to emphasize the "time evolution" aspect rather than the formal notion of a family of random variables. Formally, both objects are of course the same.

Example 9.4 Let $I = \mathbb{N}_0$ and let $(Y_n, n \in \mathbb{N})$ be a family of i.i.d. $\text{Rad}_{1/2}$ -random variables on a probability space $(\Omega, \mathcal{F}, \mathbf{P})$; that is, random variables with

$$\mathbf{P}[Y_n = 1] = \mathbf{P}[Y_n = -1] = \frac{1}{2}$$

Let $E = \mathbb{Z}$ (with the discrete topology) and let

$$X_t = \sum_{n=1}^t Y_n \quad \text{for all } t \in \mathbb{N}_0.$$

($X_t, t \in \mathbb{N}_0$) is called a symmetric simple random walk on $\mathbb{Z}$.

Example 9.5 The Poisson process $X = (X_t)_{t \geq 0}$ with intensity $\alpha > 0$ (see Section 5.5) is a stochastic process with range $\mathbb{N}_0$.

We introduce some further terms.

Definition 9.6 If X is a random variable (or a stochastic process), we write $\mathcal{L}[X] = \mathbf{P}_X$ for the distribution of X . If $\mathcal{G} \subset \mathcal{F}$ is a σ -algebra, then we write $\mathcal{L}[X \mid \mathcal{G}]$ for the regular conditional distribution of X given $\mathcal{G}$.

Definition 9.7 An E -valued stochastic process $X = (X_t)_{t \in I}$ is called

- (i) real-valued if $E = \mathbb{R}$,
- (ii) a process with independent increments if X is real-valued and for all $n \in \mathbb{N}$ and all $t_0, \dots, t_n \in I$ with $t_0 < t_1 < \dots < t_n$, we have that

$$(X_{t_i} - X_{t_{i-1}})_{i=1, \dots, n} \quad \text{is independent,}$$

- (iii) a Gaussian process if X is real-valued and for all $n \in \mathbb{N}$ and $t_1, \dots, t_n \in I$,

$$(X_{t_1}, \dots, X_{t_n}) \quad \text{is } n\text{-dimensional normally distributed, and}$$

- (iv) integrable (respectively square integrable) if X is real-valued and $\mathbf{E}[|X_t|] < \infty$ (respectively $\mathbf{E}[(X_t)^2] < \infty$) for all $t \in I$.

Now assume that $I \subset \mathbb{R}$ is closed under addition. Then X is called

- (v) stationary if $\mathcal{L}[(X_{s+t})_{t \in I}] = \mathcal{L}[(X_t)_{t \in I}]$ for all $s \in I$, and
- (vi) a process with stationary increments if X is real-valued and

$$\mathcal{L}[X_{s+t+r} - X_{t+r}] = \mathcal{L}[X_{s+r} - X_r] \quad \text{for all } r, s, t \in I$$

(If $0 \in I$, then it is enough to consider $r = 0$.)

Example 9.8

(i) The Poisson process with intensity θ and the random walk on $\mathbb{Z}$ are processes with stationary independent increments.

(ii) If $X_t, t \in I$, are i.i.d. random variables, then $(X_t)_{t \in I}$ is stationary.

(iii) Let $(X_n)_{n \in \mathbb{Z}}$ be real-valued and stationary and let $k \in \mathbb{N}$ and $c_0, \dots, c_k \in \mathbb{R}$. Define

$$Y_n := \sum_{i=0}^k c_i X_{n-i}.$$

Then $Y = (Y_n)_{n \in \mathbb{Z}}$ is a stationary process. If $c_0, \dots, c_k \geq 0$ and $c_0 + \dots + c_k = 1$, then Y is called the moving average of X (with weights $c_0, \dots, c_k$).

The following two definitions make sense also for more general index sets I that are partially ordered. However, we restrict ourselves to the case $I \subset \mathbb{R}$.

Definition 9.9 (Filtration) Let $\mathbb{F} = (\mathcal{F}_t, t \in I)$ be a family of σ -algebras with $\mathcal{F}_t \subset \mathcal{F}$ for all $t \in I$. $\mathbb{F}$ is called a filtration if $\mathcal{F}_s \subset \mathcal{F}_t$ for all $s, t \in I$ with $s \leq t$.

Definition 9.10 A stochastic process $X = (X_t, t \in I)$ is called adapted to the filtration $\mathbb{F}$ if X_t is $\mathcal{F}_t$ -measurable for all $t \in I$. If $\mathcal{F}_t = \sigma(X_s, s \leq t)$ for all $t \in I$, then we denote by $\mathbb{F} = \sigma(X)$ the filtration that is generated by X .

Remark 9.11 Clearly, a stochastic process is always adapted to the filtration it generates. Hence the generated filtration is the smallest filtration to which the process is adapted.

Definition 9.12 (Predictable) A stochastic process $X = (X_n, n \in \mathbb{N}_0)$ is called predictable (or previsible) with respect to the filtration $\mathbb{F} = (\mathcal{F}_n, n \in \mathbb{N}_0)$ if X_0 is constant and if, for every $n \in \mathbb{N}$

$$X_n \text{ is } \mathcal{F}_{n-1}\text{-measurable.}$$

Example 9.13 Let $I = \mathbb{N}_0$ and let $Y_1, Y_2, \dots$ be real random variables. For $n \in \mathbb{N}_0$, define $X_n := \sum_{m=1}^n Y_m$. Let

$$\mathcal{F}_0 = \{\emptyset, \Omega\} \quad \text{and} \quad \mathcal{F}_n = \sigma(Y_1, \dots, Y_n) \quad \text{for } n \in \mathbb{N}.$$

Then $\mathbb{F} = (\mathcal{F}_n, n \in \mathbb{N}_0) = \sigma(Y)$ is the filtration generated by $Y = (Y_n)_{n \in \mathbb{N}}$ and X is adapted to $\mathbb{F}$; hence $\sigma(X) \subset \mathbb{F}$. Clearly, $(Y_1, \dots, Y_n)$ is measurable with respect to $\sigma(X_1, \dots, X_n)$; hence $\sigma(Y) \subset \sigma(X)$, and thus also $\mathbb{F} = \sigma(X)$.

Now let $\tilde{X}_n := \sum_{m=1}^n \mathbb{1}_{[0, \infty)}(Y_m)$. Then $\tilde{X}$ is also adapted to $\mathbb{F}$; however, in general, $\mathbb{F} \not\supseteq \sigma(\tilde{X})$.

Example 9.14 Let $I = \mathbb{N}_0$ and let $D_1, D_2, \dots$ be i.i.d. $\text{Rad}_{1/2}$ -distributed random variables (that is, $\mathbf{P}[D_i = -1] = \mathbf{P}[D_i = 1] = \frac{1}{2}$ for all $i \in \mathbb{N}$). Let $D = (D_i)_{i \in \mathbb{N}}$ and $\mathbb{F} = \sigma(D)$. We interpret D_i as the result of a bet that gives a gain or loss of one euro for every euro we put at stake. Just before each gamble we decide how much money we bet. Let H_n be the number of euros to bet in the n th gamble. Clearly, H_n may only depend on the results of the gambles that happened earlier, but not on D_m for any $m \geq n$. To put it differently, there must be a function $F_n : \{-1, 1\}^{n-1} \rightarrow \mathbb{N}$ such that $H_n = F_n(D_1, \dots, D_{n-1})$. (For example, for the Petersburg game (Example 4.22) we had $F_n(x_1, \dots, x_{n-1}) = 2^{n-1} \mathbb{1}_{\{x_1=x_2=\dots=x_{n-1}=0\}}$.) Hence H is predictable. On the other hand, any predictable H has the form $H_n = F_n(D_1, \dots, D_{n-1})$, $n \in \mathbb{N}$, for certain functions $F_n : \{-1, 1\}^{n-1} \rightarrow \mathbb{N}$. Hence any predictable H is an admissible gambling strategy.

Definition 9.15 (Stopping time) A random variable τ with values in $I \cup \{\infty\}$ is called a stopping time (with respect to $\mathbb{F}$) if for any $t \in I$

$$\{\tau \leq t\} \in \mathcal{F}_t$$

The idea is that $\mathcal{F}_t$ reflects the knowledge of an observer at time t . Whether or not $\{\tau \leq t\}$ is true can thus be determined on the basis of the information available at time t .

Theorem 9.16 Let I be countable. τ is a stopping time if and only if $\{\tau = t\} \in \mathcal{F}_t$ for all $t \in I$.

Proof This is left as an exercise!

Example 9.17 Let $I = \mathbb{N}_0$ (or, more generally, let $I \subset [0, \infty)$ be right-discrete; that is, $t < \inf I \cap (t, \infty)$ for all $t \geq 0$, and hence I in particular is countable) and let $K \subset \mathbb{R}$ be measurable. Let X be an adapted real-valued stochastic process. Consider the first time that X is in K :

$$\tau_K := \inf \{t \in I : X_t \in K\}.$$

It is intuitively clear that τ_K should be a stopping time since we can determine by observation up to time t whether $\{\tau_K \leq t\}$ occurs. Formally, we argue that $\{X_s \in K\} \in \mathcal{F}_s \subset \mathcal{F}_t$ for all $s \leq t$. Hence also the countable union of these sets is in $\mathcal{F}_t$:

$$\{\tau_K \leq t\} = \bigcup_{s \in I \cap [0, t]} \{X_s \in K\} \in \mathcal{F}_t$$

Consider now the random time $\tilde{\tau} := \sup \{t \in I : X_t \in K\}$ of the last visit of X to K . For a fixed time t , on the basis of previous observations, we cannot determine whether X is already in K for the last time. For this we would have to rely on prophecy. Hence, in general, $\tilde{\tau}$ is not a stopping time.

Lemma 9.18 Let $I \subset [0, \infty)$ be closed under addition and let σ and τ be stopping times. Then:

- (i) $\sigma \vee \tau$ and $\sigma \wedge \tau$ are stopping times.
- (ii) If $\sigma, \tau \geq 0$, then $\sigma + \tau$ is also a stopping time.
- (iii) If $s \geq 0$, then $\tau + s$ is a stopping time. However, in general, $\tau - s$ is not.

Before we present the (simple) formal proof, we state that in particular (i) and (iii) are properties we would expect of stopping times. With (i), the interpretation is clear. For (iii), note that $\tau - s$ peeks into the future by s time units (in fact, $\{\tau - s \leq t\} \in \mathcal{F}_{t+s}$), while $\tau + s$ looks back s time units. For stopping times, however, only retrospection is allowed.

Proof (i) For $t \in I$, we have $\{\sigma \vee \tau \leq t\} = \{\sigma \leq t\} \cap \{\tau \leq t\} \in \mathcal{F}_t$ and $\{\sigma \wedge \tau \leq t\} = \{\sigma \leq t\} \cup \{\tau \leq t\} \in \mathcal{F}_t$.

(ii) Let $t \in I$. By (i), $\tau \wedge t$ and $\sigma \wedge t$ are stopping times for any $t \in I$. In particular, $\{\tau \wedge t \leq s\} \in \mathcal{F}_s \subset \mathcal{F}_t$ for any $s \leq t$. On the other hand, we have $\tau \wedge t \leq s$ for $s > t$. Hence $\tau' := (\tau \wedge t) + \mathbb{1}_{\{\tau > t\}}$ and $\sigma' := (\sigma \wedge t) + \mathbb{1}_{\{\sigma > t\}}$ (and thus $\tau' + \sigma'$) are $\mathcal{F}_t$ -measurable. We conclude $\{\tau + \sigma \leq t\} = \{\tau' + \sigma' \leq t\} \in \mathcal{F}_t$.

(iii) For $\tau + s$, this is a consequence of (ii) (with the stopping time $\sigma \equiv s$). For $\tau - s$, since τ is a stopping time, we have $\{\tau - s \leq t\} = \{\tau \leq t + s\} \in \mathcal{F}_{t+s}$. However, in general, $\mathcal{F}_{t+s}$ is a strict superset of $\mathcal{F}_t$; hence $\tau - s$ is not a stopping time. $\square$

Definition 9.19 Let τ be a stopping time. Then

$$\mathcal{F}_\tau := \{A \in \mathcal{F} : A \cap \{\tau \leq t\} \in \mathcal{F}_t \text{ for any } t \in I\}$$

is called the σ -algebra of τ -past.

Example 9.20 Let $I = \mathbb{N}_0$ (or let $I \subset [0, \infty)$ be right-discrete; compare Example 9.17) and let X be an adapted real-valued stochastic process. Let $K \in \mathbb{R}$ and let $\tau = \inf \{t : X_t \geq K\}$ be the stopping time of first entrance in $[K, \infty)$. Consider the events $A = \{\sup \{X_t : t \in I\} > K - 5\}$ and $B = \{\sup \{X_t : t \in I\} > K + 5\}$.

Clearly, $\{\tau \leq t\} \subset A$ for all $t \in I$; hence $A \cap \{\tau \leq t\} = \{\tau \leq t\} \in \mathcal{F}_t$. Thus $A \in \mathcal{F}_\tau$. However, in general, $B \notin \mathcal{F}_\tau$ since up to time τ , we cannot decide whether X will ever exceed $K + 5$. $\square$

Lemma 9.21 If σ and τ are stopping times with $\sigma \leq \tau$, then $\mathcal{F}_\sigma \subset \mathcal{F}_\tau$.

Proof Let $A \in \mathcal{F}_\sigma$ and $t \in I$. Then $A \cap \{\sigma \leq t\} \in \mathcal{F}_t$. Now $\{\tau \leq t\} \in \mathcal{F}_t$ since τ is a stopping time. Since $\sigma \leq \tau$, we thus get

$$A \cap \{\tau \leq t\} = (A \cap \{\sigma \leq t\}) \cap \{\tau \leq t\} \in \mathcal{F}_t.$$

$\square$

Definition 9.22 If $\tau < \infty$ is a stopping time, then we define $X_\tau(\omega) := X_{\tau(\omega)}(\omega)$.

Lemma 9.23 Let I be countable, let X be adapted and let $\tau < \infty$ be a stopping time. Then X_τ is measurable with respect to $\mathcal{F}_\tau$.

Proof Let A be measurable and $t \in I$. Hence $\{\tau = s\} \cap X_s^{-1}(A) \in \mathcal{F}_s \subset \mathcal{F}_t$ for all $s \leq t$. Thus

$$X_\tau^{-1}(A) \cap \{\tau \leq t\} = \bigcup_{\substack{s \in I \\ s \leq t}} (\{\tau = s\} \cap X_s^{-1}(A)) \in \mathcal{F}_t$$

$\square$

For uncountable I and for fixed ω , in general, the map $I \rightarrow E, t \mapsto X_t(\omega)$ is not measurable; hence neither is the composition X_τ always measurable. Here one needs assumptions on the regularity of the paths $t \mapsto X_t(\omega)$; for example, right continuity. We come back to this point in Chapter 21 and leave this as a warning for the time being.

5.9.2 9.2 Martingales

Everyone who does not own a casino would agree without hesitation that the successive payment of gains $Y_1, Y_2, \dots$, such that $Y_1, Y_2, \dots$ are i.i.d. with $\mathbf{E}[Y_1] = 0$, could be considered a fair game consisting of consecutive rounds. In this case, the process X of partial sums $X_n = Y_1 + \dots + Y_n$ is integrable and $\mathbf{E}[X_n | \mathcal{F}_m] = X_m$ if $m < n$ (where $\mathbb{F} = \sigma(X)$). We want to use this equation for the conditional expectations as the defining equation for a fair game that in the following will be called a martingale. Note that, in particular, this definition does not require that the individual payments be independent or identically distributed. This makes the notion quite a bit more flexible. The momentousness of the following concept will become manifest only gradually.

Definition 9.24 Let $(\Omega, \mathcal{F}, \mathbf{P})$ be a probability space, $I \subseteq \mathbb{R}$ a filtration. Let $X = (X_t)_{t \in I}$ be a real-valued, adapted stochastic process. X is called (with respect to $\mathbb{F}$)

- martingale if $\mathbf{E}[X_t | \mathcal{F}_s] = X_s$ for all $s, t \in I$ with $s < t$,
- submartingale if $\mathbf{E}[X_t | \mathcal{F}_s] \geq X_s$ for all $s, t \in I$ with $s < t$,
- supermartingale if $\mathbf{E}[X_t | \mathcal{F}_s] \leq X_s$ for all $s, t \in I$ with $s < t$.

Remark 9.25 Clearly, for a martingale, the map $t \mapsto \mathbf{E}[X_t]$ is constant, for submartingales it is monotone increasing and for supermartingales it is monotone decreasing. $\square$

Remark 9.26 The etymology of the term martingale has not been resolved completely. The French *la martingale* (originally Provençal *martegalo*, named after the town Martiques) in equitation means "a piece of rein used in jumping and cross country riding". Sometimes the ramified shape, in particular of the running martingale (French *la martingale à anneaux*), is considered as emblematic for the doubling strategy in the Petersburg game.

This doubling strategy itself is the second meaning of *la martingale*. Starting here, a shift in the meaning towards the mathematical notion seems plausible. A different derivation, in contrast to the appearance, is based on the function of the rein, which is to "check the upward movement of the horse's head". Thus the notion of a martingale might first have been used for general gambling strategies (checking the movements of chance) and later for the doubling strategy in particular.

Remark 9.27 If $I = \mathbb{N}$, $I = \mathbb{N}_0$ or $I = \mathbb{Z}$, then it is enough to consider at each instant s only $t = s + 1$. In fact, by the tower property of the conditional expectation (Theorem 8.14(iv)), we get

$$\mathbf{E}[X_{s+2} | \mathcal{F}_s] = \mathbf{E}[\mathbf{E}[X_{s+2} | \mathcal{F}_{s+1}] | \mathcal{F}_s].$$

Thus, if the defining equality (or inequality) holds for any time step of size one, by induction it holds for all times. $\square$

Remark 9.28 If we do not explicitly mention the filtration $\mathbb{F}$, we tacitly assume that $\mathbb{F}$ is generated by X ; that is, $\mathcal{F}_t = \sigma(X_s, s \leq t)$. $\square$

Remark 9.29 Let $\mathbb{F}$ and $\mathbb{F}'$ be filtrations with $\mathcal{F}_t \subset \mathcal{F}'_t$ for all t , and let X be an $\mathbb{F}'$ (sub-, super-) martingale that is adapted to $\mathbb{F}$. Then X is also a (sub-, super-) martingale with respect to the smaller filtration $\mathbb{F}$. Indeed, for $s < t$ and for the case of a submartingale,

$$\mathbf{E}[X_t | \mathcal{F}_s] = \mathbf{E}[\mathbf{E}[X_t | \mathcal{F}'_s] | \mathcal{F}_s] \geq \mathbf{E}[X_s | \mathcal{F}_s] = X_s.$$

In particular, an $\mathbb{F}$ -(sub-, super-) martingale X is always a (sub-, super-) martingale with respect to its own filtration $\sigma(X)$. $\square$

Example 9.30 Let $Y_1, \dots, Y_N$ be independent random variables with $\mathbf{E}[Y_t] = 0$ for all $t = 1, \dots, N$. Let $\mathcal{F}_t := \sigma(Y_1, \dots, Y_t)$ and $X_t := \sum_{s=1}^t Y_s$. Then X is adapted and integrable, and $\mathbf{E}[Y_r | \mathcal{F}_s] = 0$ for $r > s$. Hence, for $t > s$,

$$\mathbf{E}[X_t | \mathcal{F}_s] = \mathbf{E}[X_s | \mathcal{F}_s] + \mathbf{E}[X_t - X_s | \mathcal{F}_s] = X_s + \sum_{r=s+1}^t \mathbf{E}[Y_r | \mathcal{F}_s] = X_s$$

Thus, X is an $\mathbb{F}$ -martingale.

Similarly, X is a submartingale if $\mathbf{E}[Y_t] \geq 0$ for all t , and a supermartingale if $\mathbf{E}[Y_t] \leq 0$ for all t . $\square$

Example 9.31 Consider the situation of the preceding example; however, now with $\mathbf{E}[Y_t] = 1$ and $X_t = \prod_{s=1}^t Y_s$ for $t \in \mathbb{N}_0$. By Theorem 5.4, $Y_1 \cdot Y_2$ is integrable.

Inductively, we get $\mathbf{E}[|X_t|] < \infty$ for all $t \in \mathbb{N}_0$. Evidently, X is adapted to $\mathbb{F}$ and for all $s \in \mathbb{N}_0$, we have

$$\mathbf{E}[X_{s+1} | \mathcal{F}_s] = \mathbf{E}[X_s Y_{s+1} | \mathcal{F}_s] = X_s \mathbf{E}[Y_{s+1} | \mathcal{F}_s] = X_s$$

Hence X is an $\mathbb{F}$ -martingale.

Theorem 9.32

- (i) X is a supermartingale if and only if $(-X)$ is a submartingale.
- (ii) Let X and Y be martingales and let $a, b \in \mathbb{R}$. Then $(aX + bY)$ is a martingale.
- (iii) Let X and Y be supermartingales and $a, b \geq 0$. Then $(aX + bY)$ is a supermartingale.
- (iv) Let X and Y be supermartingales. Then $Z := X \wedge Y = (\min(X_t, Y_t))_{t \in I}$ is a supermartingale.
- (v) If $(X_t)_{t \in \mathbb{N}_0}$ is a supermartingale and $\mathbf{E}[X_T] \geq \mathbf{E}[X_0]$ for some $T \in \mathbb{N}_0$, then $(X_t)_{t \in \{0, \dots, T\}}$ is a martingale. If there exists a sequence $T_N \rightarrow \infty$ with $\mathbf{E}[X_{T_N}] \geq \mathbf{E}[X_0]$, then X is a martingale.

Proof (i), (ii) and (iii) These are evident.

(iv) Since $|Z_t| \leq |X_t| + |Y_t|$, we have $\mathbf{E}[|Z_t|] < \infty$ for all $t \in I$. Due to monotonicity of the conditional expectation (Theorem 8.14(ii)), for $t > s$, we have $\mathbf{E}[Z_t | \mathcal{F}_s] \leq \mathbf{E}[Y_t | \mathcal{F}_s] \leq Y_s$ and $\mathbf{E}[Z_t | \mathcal{F}_s] \leq \mathbf{E}[X_t | \mathcal{F}_s] \leq X_s$; hence

$$\mathbf{E}[Z_t | \mathcal{F}_s] \leq X_s \wedge Y_s = Z_s$$

(v) For $t \leq T$, let $Y_t := \mathbf{E}[X_T | \mathcal{F}_t]$. Then Y is a martingale and $Y_t \leq X_t$. Hence

$$\mathbf{E}[X_0] \leq \mathbf{E}[X_T] = \mathbf{E}[Y_T] = \mathbf{E}[Y_t] \leq \mathbf{E}[X_t] \leq \mathbf{E}[X_0]$$

(The first inequality holds by assumption.) We infer that $Y_t = X_t$ almost surely for all t and thus $(X_t)_{t \in \{0, \dots, T\}}$ is a martingale.

Let $T_N \rightarrow \infty$ with $\mathbf{E}[X_{T_N}] \geq \mathbf{E}[X_0]$ for all $N \in \mathbb{N}$. Then, for any $t > s \geq 0$, there is an $N \in \mathbb{N}$ with $T_N > t$. Hence, $\mathbf{E}[X_t | \mathcal{F}_s] = \mathbf{E}[X_s]$ and X is a martingale.

Remark 9.33 Many statements about supermartingales hold mutatis mutandis for submartingales. For example, in the preceding theorem, claim (i) holds with the words "submartingale" and "supermartingale" interchanged, claim (iv) holds for submartingales if the minimum is replaced by a maximum, and so on. We often do not give the statements both for submartingales and for supermartingales. Instead, we choose representatively one case. Note,

however, that those statements that we make explicitly about martingales usually cannot be adapted easily to sub- or supermartingales (such as (ii) in the preceding theorem).

Corollary 9.34 Let X be a submartingale and $a \in \mathbb{R}$. Then $(X - a)^+$ is a submartingale.

Proof Clearly, 0 and $Y = X - a$ are submartingales. By (iv), $(X - a)^+ = Y \vee 0$ is also a submartingale.

Theorem 9.35 Let X be a martingale and let $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ be a convex function.

(i) If

$$\mathbf{E} [\varphi (X_t)^+] < \infty \quad \text{for all } t \in I, \quad (9.1)$$

then $(\varphi (X_t))_{t \in I}$ is a submartingale.

(ii) If $t^* := \sup(I) \in I$, then $\mathbf{E} [\varphi (X_{t^*})^+] < \infty$ implies (9.1).

(iii) In particular, if $p \geq 1$ and $\mathbf{E} [|X_t|^p] < \infty$ for all $t \in I$, then $(|X_t|^p)_{t \in I}$ is a submartingale.

Proof (i) We always have $\mathbf{E} [\varphi (X_t)^-] < \infty$ (Theorem 7.9); hence, by assumption, $\mathbf{E} [|\varphi (X_t)|] < \infty$ for all $t \in I$. Jensen's inequality (Theorem 8.20) then yields, for $t > s$,

$$\mathbf{E} [\varphi (X_t) \mid \mathcal{F}_s] \geq \varphi (\mathbf{E} [X_t \mid \mathcal{F}_s]) = \varphi (X_s)$$

(ii) Since φ is convex, so is $x \mapsto \varphi(x)^+$. Furthermore, by assumption, we have $\mathbf{E} [\varphi (X_{t^*})^+] < \infty$; hence Jensen's inequality implies that, for all $t \in I$,

$$\begin{aligned} \mathbf{E} [\varphi (X_t)^+] &= \mathbf{E} [\varphi (\mathbf{E} [X_{t^*} \mid \mathcal{F}_t])^+] \leq \mathbf{E} [\mathbf{E} [\varphi (X_{t^*})^+ \mid \mathcal{F}_t]] \\ &= \mathbf{E} [\varphi (X_{t^*})^+] < \infty \end{aligned}$$

(iii) This is evident since $x \mapsto |x|^p$ is convex.

Example 9.36 (See Example 9.4.) Symmetric simple random walk X on $\mathbb{Z}$ is a square integrable martingale. Hence $(X_n^2)_{n \in \mathbb{N}_0}$ is a submartingale.

Exercise 9.2.1 Let Y be a random variable with $\mathbf{E}[|Y|] < \infty$ and let $\mathbb{F}$ be a filtration as well as

$$X_t := \mathbf{E}[Y \mid \mathcal{F}_t] \quad \text{for all } t \in I$$

Show that X is an $\mathbb{F}$ -martingale.

Exercise 9.2.2 Let $(X_n)_{n \in \mathbb{N}_0}$ be a predictable $\mathbb{F}$ -martingale. Show that $X_n = X_0$ almost surely for all $n \in \mathbb{N}_0$.

Exercise 9.2.3 Show that the claim of Theorem 9.35 continues to hold if X is only a submartingale but if φ is in addition assumed to be monotone increasing. Give an example that shows that the monotonicity of φ is essential. (Compare Corollary 9.34.)

Exercise 9.2.4 (Azuma's inequality) Show the following.

(i) If X is a random variable with $|X| \leq 1$ a.s., then there is a random variable Y with values in $\{-1, +1\}$ and with $\mathbf{E}[Y \mid X] = X$.

(ii) For X as in (i) with $\mathbf{E}[X] = 0$, infer that (using Jensen's inequality)

$$\mathbf{E} [e^{\lambda X}] \leq \cosh(\lambda) \leq e^{\lambda^2/2} \quad \text{for all } \lambda \in \mathbb{R}.$$

(iii) If $(M_n)_{n \in \mathbb{N}_0}$ is a martingale with $M_0 = 0$ and if there is a sequence $(c_k)_{k \in \mathbb{N}}$ of nonnegative numbers with $|M_n - M_{n-1}| \leq c_n$ a.s. for all $n \in \mathbb{N}$, then

$$\mathbf{E} [e^{\lambda M_n}] \leq \exp \left(\frac{1}{2} \lambda^2 \sum_{k=1}^n c_k^2 \right)$$

(iv) Under the assumptions of (iii), Azuma's inequality holds:

$$\mathbf{P}[|M_n| \geq \lambda] \leq 2 \exp\left(-\frac{\lambda^2}{2 \sum_{k=1}^n c_k^2}\right) \quad \text{for all } \lambda \geq 0$$

Hint: Use Markov's inequality for $f(x) = e^{\gamma x}$ and choose the optimal γ .

5.9.3 9.3 Discrete Stochastic Integral

So far we have encountered a martingale as the process of partial sums of gains of a fair game. This game can also be the price of a stock that is traded at discrete times on a stock exchange. With this interpretation, it is particularly evident that it is natural to construct new stochastic processes by considering investment strategies for the stock. The value of the portfolio, which is the new stochastic process, changes as the stock price changes. It is the price multiplied by the number of stocks in the portfolio. In order to describe such processes formally, we introduce the following notion.

Definition 9.37 (Discrete stochastic integral) Let $(X_n)_{n \in \mathbb{N}_0}$ be an $\mathbb{F}$ -adapted real process and let $(H_n)_{n \in \mathbb{N}}$ be a real-valued and $\mathbb{F}$ -predictable process. The discrete stochastic integral of H with respect to X is the stochastic process $H \cdot X$ defined by

$$(H \cdot X)_n := \sum_{m=1}^n H_m (X_m - X_{m-1}) \quad \text{for } n \in \mathbb{N}_0 \quad (9.2)$$

If X is a martingale, then $H \cdot X$ is also called the martingale transform of X .

Remark 9.38 Clearly, $H \cdot X$ is adapted to $\mathbb{F}$.

Let X be a (possibly unfair) game where $X_n - X_{n-1}$ is the gain per euro in the n th round. We interpret H_n as the number of euros we bet in the n th game. H is then a gambling strategy. Clearly, the value of H_n has to be decided at time $n-1$; that is, before the result of X_n is known. In other words, H must be predictable.

Now assume that X is a fair game (that is, a martingale) and H is locally bounded (that is, each H_n is bounded). Then (since $\mathbf{E}[X_{n+1} - X_n | \mathcal{F}_n] = 0$)

$$\begin{aligned} \mathbf{E}[(H \cdot X)_{n+1} | \mathcal{F}_n] &= \mathbf{E}[(H \cdot X)_n + H_{n+1}(X_{n+1} - X_n) | \mathcal{F}_n] \\ &= (H \cdot X)_n + H_{n+1} \mathbf{E}[X_{n+1} - X_n | \mathcal{F}_n] \\ &= (H \cdot X)_n. \end{aligned}$$

Thus $H \cdot X$ is a martingale. The following theorem says that the converse also holds; that is, X is a martingale if, for sufficiently many predictable processes, the stochastic integral is a martingale.

Theorem 9.39 (Stability theorem) Let $(X_n)_{n \in \mathbb{N}_0}$ be an adapted, real-valued stochastic process with $\mathbf{E}[|X_0|] < \infty$.

(i) X is a martingale if and only if, for any locally bounded predictable process H , the stochastic integral $H \cdot X$ is a martingale.

(ii) X is a submartingale (supermartingale) if and only if $H \cdot X$ is a submartingale (supermartingale) for any locally bounded predictable $H \geq 0$.

Proof (i) " $\implies$ " This has been shown in the discussion above.

" $\implies$ " Fix an $n_0 \in \mathbb{N}$, and let $H_n = \mathbb{1}_{\{n=n_0\}}$. Then $(H \cdot X)_{n_0-1} = 0$; hence

$$0 = \mathbf{E}[(H \cdot X)_{n_0} | \mathcal{F}_{n_0-1}] = \mathbf{E}[X_{n_0} | \mathcal{F}_{n_0-1}] - X_{n_0-1}$$

(ii) This is similar to (i).

The preceding theorem says, in particular, that we cannot find any locally bounded gambling strategy that transforms a martingale (or, if we are bound to nonnegative gambling strategies,

as we are in real life, a supermartingale) into a submartingale. Quite the contrary is suggested by the many invitations to play all kinds of "sure winning systems" in lotteries.

Example 9.40 (Petersburg game) We continue Example 9.14 (see also Example 4.22). Define $X_n := D_1 + \dots + D_n$ for $n \in \mathbb{N}_0$. Then X is a martingale. The gambling strategy $H_n := 2^{n-1} \mathbb{1}_{\{D_1=D_2=\dots=D_{n-1}=-1\}}$ for $n \in \mathbb{N}$ and $H_0 = 1$ is predictable and locally bounded. Let $S_n = \sum_{i=1}^n H_i D_i = (H \cdot X)_n$ be the gain after n rounds. Then S is a martingale by the preceding theorem. In particular, we get (as shown already in Example 4.22) that $\mathbf{E}[S_n] = 0$ for all $n \in \mathbb{N}$. We will come back later to the point that this superficially contrasts with $S_n \xrightarrow{n \rightarrow \infty} 1$ a.s. (see Example 11.6).

For the moment, note that the martingale $S' = (1 - S_n)_{n \in \mathbb{N}_0}$, just like the one in Example 9.31, has the structure of a product of independent random variables with expectation 1. In fact, $S'_n = \prod_{i=1}^n (1 - D_i)$.

5.9.4 9.4 Discrete Martingale Representation Theorem and the Crr Model

By virtue of the stochastic integral, we have transformed a martingale X via a gambling strategy H into a new martingale $H \cdot X$. Let us change the perspective and ask: For fixed X , which are the martingales Y (with $Y_0 = 0$) that can be obtained as discrete stochastic integrals of X with a suitable gambling strategy $H = H(Y)$? Possibly all martingales Y ? This is not the case, in general, as the example below indicates. However, we will see that all martingales can be represented as stochastic integrals if the increments $X_{n+1} - X_n$ can take only two values (given $X_1, \dots, X_n$). In this case, we give a representation theorem and use it to discuss the fair price for a European call option in the stock market model of Cox-Ross-Rubinstein. This model is rather simple and describes an idealized market (no transaction costs, fractional numbers of stocks tradeable and so on). For extensive literature on stochastic aspects of mathematical finance, we refer to the textbooks [9, 42, 48, 57, 86, 102, 121] or [159].

Example 9.41 Consider the very simple martingale $X = (X_n)_{n=0,1}$ with only two time points. Let $X_0 = 0$ almost surely and $\mathbf{P}[X_1 = -1] = \mathbf{P}[X_1 = 0] = \mathbf{P}[X_1 = 1] = \frac{1}{3}$. Let $Y_0 = 0$. Further, let $Y_1 = 2$ if $X_1 = 1$ and $Y_1 = -1$ otherwise. Then Y is manifestly a $\sigma(X)$ -martingale. However, there is no number H_1 such that $H_1 X_1 = Y_1$. $\square$

Let $T \in \mathbb{N}$ be a fixed time. If $(Y_n)_{n=0,1,\dots,T}$ is an $\mathbb{F}$ -martingale, then $Y_n = \mathbf{E}[Y_T | \mathcal{F}_n]$ for all $n \leq T$. An $\mathbb{F}$ -martingale Y is thus determined uniquely by the terminal values Y_T (and vice versa). Let X be a martingale. As $(H \cdot X)$ is a martingale, the representation problem for martingales is thus reduced to the problem of representing an integrable random variable $V := Y_T$ as $v_0 + (H \cdot X)_T$, where $v_0 = \mathbf{E}[Y_T]$.

We saw that, in general, this is not possible if the differences $X_{n+1} - X_n$ take three (or more) different values. Hence we now consider the case where only two values are possible. Here, at each time step, a system of two linear equations with two unknowns has to be solved. In the case where $X_{n+1} - X_n$ takes three values, the system has three equations and is thus overdetermined.

Definition 9.42 (Binary model) A stochastic process $X_0, \dots, X_T$ is called binary splitting or a binary model if there exist random variables $D_1, \dots, D_T$ with values in $\{-1, +1\}$ and functions $f_n : \mathbb{R}^{n-1} \times \{-1, +1\} \rightarrow \mathbb{R}$ for $n = 1, \dots, T$, as well as $x_0 \in \mathbb{R}$ such that $X_0 = x_0$ and

$$X_n = f_n(X_1, \dots, X_{n-1}, D_n) \quad \text{for any } n = 1, \dots, T.$$

By $\mathbb{F} = \sigma(X)$, we denote the filtration generated by X .

Note that X_n depends only on $X_1, \dots, X_{n-1}$ and D_n but not on the full information inherent in the values $D_1, \dots, D_n$. If the latter were the case, a ramification into more than two values in one time step would be possible.

Theorem 9.43 (Representation theorem) Let X be a binary model and let V_T be an $\mathcal{F}_T$ -measurable random variable. Then there exists a bounded predictable process H and $av_0 \in \mathbb{R}$ with $V_T = v_0 + (H \cdot X)_T$.

Note that $\mathbb{F}$ is the filtration generated by X , not the, possibly larger, filtration generated by $D_1, \dots, D_T$. For the latter case, the claim of the theorem would be incorrect since, loosely speaking, with H we can bet on X but not on the D_i .

Proof We show that there exist $\mathcal{F}_{T-1}$ -measurable random variables V_{T-1} and H_T such that $V_T = V_{T-1} + H_T(X_T - X_{T-1})$. By a backward induction, this yields the claim.

Since V_T is $\mathcal{F}_T$ -measurable, by the factorization lemma (Corollary 1.97) there exists a function $g_T : \mathbb{R}^T \rightarrow \mathbb{R}$ with $V_T = g_T(X_1, \dots, X_T)$. Define

$$X_T^\pm = f_T(X_1, \dots, X_{T-1}, \pm 1) \quad \text{and} \quad V_T^\pm = g_T(X_1, \dots, X_{T-1}, X_T^\pm).$$

Each of these four random variables is manifestly $\mathcal{F}_{T-1}$ -measurable. Hence we are looking for solutions V_{T-1} and H_T of the following system of linear equations:

$$\begin{aligned} V_{T-1} + H_T(X_T^- - X_{T-1}) &= V_T^- \\ V_{T-1} + H_T(X_T^+ - X_{T-1}) &= V_T^+ \end{aligned} \tag{9.3}$$

By construction, $X_T^+ - X_T^- \neq 0$ if $V_T^+ - V_T^- \neq 0$. Hence we can solve (9.3) and get

$$H_T := \begin{cases} \frac{V_T^+ - V_T^-}{X_T^+ - X_T^-}, & \text{if } X_T^+ \neq X_T^- \\ 0, & \text{else} \end{cases}$$

and $V_{T-1} = V_T^+ - H_T(X_T^+ - X_{T-1}) = V_T^- - H_T(X_T^- - X_{T-1})$.

We now want to interpret X as the market price of a stock and V_T as the payment of a financial derivative on X , a so-called contingent claim or, briefly, claim. For example, V_T could be a (European) call option with maturity T and strike price $K \geq 0$. In this case, we have $V_T = (X_T - K)^+$. Economically speaking, the European call gives the buyer the right (but not the obligation) to buy one stock at time T at price K (from the issuer of the option). As typically the option is exercised only if the market price at time T is larger than K (and then gives a profit of $X_T - K$ as the stock could be sold at price X_T on the market), the value of the option is indeed $V_T = (X_T - K)^+$.

At the stock exchanges, not only are stocks traded but also derivatives on stocks. Hence, what is the fair price $\pi(V_T)$ for which a trader would offer (and buy) the contingent claim V_T ? If there exists a strategy H and a v_0 such that $V_T = v_0 + (H \cdot X)_T$, then the trader can sell the claim for v_0 (at time 0) and replicate the claim by building a portfolio that follows the trading strategy H . In this case, the claim V_T is called replicable and the strategy H is called a hedging strategy, or briefly a hedge. A market in which every claim can be replicated is called complete. In this sense, the binary model is a complete market.

If there was a second strategy H' and a second v'_0 with $v'_0 + (H' \cdot X)_T = V_T$, then, in particular, $v_0 - v'_0 = ((H' - H) \cdot X)_T$. If we had $v_0 > v'_0$, then the trader could follow the strategy $H' - H$ (which gives a final payment of $V_T - V_T = 0$) and make a sure profit of $v_0 - v'_0$. In the opposite case, $v_0 < v'_0$, the strategy $H - H'$ ensures a risk-free profit. Such a risk-free profit (or free lunch in economic jargon) is called an arbitrage. It is reasonable to assume that a market gives no opportunity for an arbitrage. Hence the fair price $\pi(V_T)$ is determined uniquely once there is one trading strategy H and a v_0 such that $V_T = v_0 + (H \cdot X)_T$.

Until now, we have not assumed that X is a martingale. However, if X is a martingale, then $(H \cdot X)$ is a martingale with $(H \cdot X)_0 = 0$; hence clearly $\mathbf{E}[(H \cdot X)_T] = 0$. Thus

$$\pi(V_T) = v_0 = \mathbf{E}[V_T] \quad (9.4)$$

Since, in this case, v_0 does not depend on the trading strategy and is hence unique, the market is automatically arbitrage-free. A finite market is thus arbitrage-free if and only if there exists an equivalent martingale (to be defined below). In this case, uniqueness of this martingale is equivalent to completeness of the market ("the fundamental theorem of asset pricing" by Harrison-Pliska [68]). In larger markets, equivalence holds only with a somewhat more flexible notion of arbitrage (see [30]).

Now if X is not a martingale, then in some cases, we can replace X by a different process X' that is a martingale and such that the distributions $\mathbf{P}_X$ and $\mathbf{P}_{X'}$ are equivalent; that is, have the same null sets. Such a process is called an equivalent martingale, and $\mathbf{P}_{X'}$ is called an equivalent martingale measure. A trading strategy that replicates V_T with respect to X also replicates V_T with respect to X' . In particular, the fair price does not change if we pass to the equivalent martingale X' . Thus we can compute $\pi(V_T)$ by applying (9.4) to the equivalent martingale.

While here this is only of interest in that it simplifies the computation of fair prices, it has an economic interpretation as a measure for the market prices that we would see if all traders were risk-neutral; that is, for traders who price a future payment by its mean value. Typically, however, traders are risk-averse and thus real market prices include a discount due to the inherent risk.

Now we consider one model in greater detail.

Definition 9.44 Let $T \in \mathbb{N}$, $a \in (-1, 0)$ and $b > 0$ as well as $p \in (0, 1)$. Further, let $D_1, \dots, D_T$ be i.i.d. Rad_p random variables (that is, $\mathbf{P}[D_1 = 1] = 1 - \mathbf{P}[D_1 = -1] = p$). We let $X_0 = x_0 > 0$ and for $n = 1, \dots, T$, define

$$X_n = \begin{cases} (1+b)X_{n-1}, & \text{if } D_n = +1 \\ (1+a)X_{n-1}, & \text{if } D_n = -1 \end{cases}$$

X is called the multi-period binomial model or the Cox-Ross-Rubinstein model (without interest returns).

As we have shown already, the CRR model is complete. Further, with the choice $p = p^* := \frac{a}{a-b}$, we can change X into a martingale. Hence the model is also arbitrage-free (for all $p \in (0, 1)$). Now we want to compute the price of a European call option $V_T := (X_T - K)^+$ explicitly. To this end, we can assume $p = p^*$. Letting $A := \min \{i \in \mathbb{N}_0 : (1+b)^i(1+a)^{T-i}x_0 > K\}$, we get

$$\begin{aligned} \pi(V_T) &= \mathbf{E}_{p^*}[V_T] = \sum_{i=0}^T b_{T,p^*}(\{i\}) [(1+b)^i(1+a)^{T-i}x_0 - K]^+ \\ &= x_0 \sum_{i=A}^T \binom{T}{i} (p^*)^i (1-p^*)^{T-i} [(1+b)^i(1+a)^{T-i}] \\ &\quad - K \sum_{i=A}^T b_{T,p^*}(\{i\}). \end{aligned}$$

If we define $p' = (1+b)p^*$, then $p' \in (0, 1)$ and $1-p' = (1-p^*)(1+a)$. We thus obtain the Cox-Ross-Rubinstein formula

$$\pi(V_T) = x_0 b_{T,p'}(\{A, \dots, T\}) - K b_{T,p^*}(\{A, \dots, T\}). \quad (9.5)$$

This is the discrete analogue of the celebrated Black-Scholes formula for option pricing in certain time-continuous markets.

5.10 10 Optional Sampling Theorems

In Chapter 9 we saw that martingales are transformed into martingales if we apply certain admissible gambling strategies. In this chapter, we establish a similar stability property for martingales that are stopped at a random time. In order also to obtain these results for submartingales and supermartingales, in the first section, we start with a decomposition theorem for adapted processes. We show the optional sampling and optional stopping theorems in the second section. The chapter finishes with the investigation of random stopping times with an infinite time horizon.

5.10.1 10.1 Doob Decomposition and Square Variation

Let $X = (X_n)_{n \in \mathbb{N}_0}$ be an adapted process with $\mathbf{E}[|X_n|] < \infty$ for all $n \in \mathbb{N}_0$. We will decompose X into a sum consisting of a martingale and a predictable process. To this end, for $n \in \mathbb{N}_0$, define

$$M_n := X_0 + \sum_{k=1}^n (X_k - \mathbf{E}[X_k | \mathcal{F}_{k-1}]) \quad (10.1)$$

and

$$A_n := \sum_{k=1}^n (\mathbf{E}[X_k | \mathcal{F}_{k-1}] - X_{k-1}).$$

Evidently, $X_n = M_n + A_n$. By construction, A is predictable with $A_0 = 0$, and M is a martingale since

$$\mathbf{E}[M_n - M_{n-1} | \mathcal{F}_{n-1}] = \mathbf{E}[X_n - \mathbf{E}[X_n | \mathcal{F}_{n-1}] | \mathcal{F}_{n-1}] = 0.$$

Theorem 10.1 (Doob decomposition) Let $X = (X_n)_{n \in \mathbb{N}_0}$ be an adapted integrable process. Then there exists a unique decomposition $X = M + A$, where A is predictable with $A_0 = 0$ and M is a martingale. This representation of X is called the Doob decomposition. X is a submartingale if and only if A is monotone increasing.

Proof We only have to show uniqueness of the decomposition. Hence, let $X = M + A = M' + A'$ be two such decompositions. Then $M - M' = A' - A$ is a predictable martingale; hence (see Exercise 9.2.2) $M_n - M'_n = M_0 - M'_0 = 0$ for all $n \in \mathbb{N}_0$.

Example 10.2 Let $I = \mathbb{N}_0$ or $I = \{0, \dots, N\}$. Let $(X_n)_{n \in I}$ be a square integrable $\mathbb{F}$ -martingale (that is, $\mathbf{E}[X_n^2] < \infty$ for all $n \in I$). By Theorem 9.35, $Y := (X_n^2)_{n \in I}$ is a submartingale. Let $Y = M + A$ be the Doob decomposition of Y . Then $(X_n^2 - A_n)_{n \in I}$ is a martingale. Furthermore, $\mathbf{E}[X_{i-1}X_i | \mathcal{F}_{i-1}] = X_{i-1}\mathbf{E}[X_i | \mathcal{F}_{i-1}] = X_{i-1}^2$; hence (as in (10.1))

$$\begin{aligned} A_n &= \sum_{i=1}^n (\mathbf{E}[X_i^2 | \mathcal{F}_{i-1}] - X_{i-1}^2) \\ &= \sum_{i=1}^n (\mathbf{E}[(X_i - X_{i-1})^2 | \mathcal{F}_{i-1}] - 2X_{i-1}^2 + 2\mathbf{E}[X_{i-1}X_i | \mathcal{F}_{i-1}]) \\ &= \sum_{i=1}^n \mathbf{E}[(X_i - X_{i-1})^2 | \mathcal{F}_{i-1}] \end{aligned}$$

Definition 10.3 Let $(X_n)_{n \in I}$ be a square integrable $\mathbb{F}$ -martingale. The unique predictable process A for which $(X_n^2 - A_n)_{n \in I}$ becomes a martingale is called the square variation process of X and is denoted by $(\langle X \rangle_n)_{n \in I} := A$.

By the preceding example, we conclude the following theorem.

Theorem 10.4 Let X be as in Definition 10.3. Then, for $n \in \mathbb{N}_0$,

$$\langle X \rangle_n = \sum_{i=1}^n \mathbf{E} [(X_i - X_{i-1})^2 \mid \mathcal{F}_{i-1}] \quad (10.2)$$

and

$$\mathbf{E} [\langle X \rangle_n] = \text{Var} [X_n - X_0]. \quad (10.3)$$

Remark 10.5 If Y and A are as in Example 10.2, then A is monotone increasing since $(X_n^2)_{n \in I}$ is a submartingale (see Theorem 10.1). Therefore, A is sometimes called the increasing process of Y .

Example 10.6 Let $Y_1, Y_2, \dots$ be independent, square integrable, centered random variables. Then $X_n := Y_1 + \dots + Y_n$ defines a square integrable martingale with $\langle X \rangle_n = \sum_{i=1}^n \mathbf{E} [Y_i^2]$. In fact, $A_n = \sum_{i=1}^n \mathbf{E} [Y_i^2 \mid Y_1, \dots, Y_{i-1}] = \sum_{i=1}^n \mathbf{E} [Y_i^2]$ (as in Example 10.2).

Note that in order for $\langle X \rangle$ to have the simple form as in Example 10.6, it is not enough for the random variables $Y_1, Y_2, \dots$ to be uncorrelated.

Example 10.7 Let $Y_1, Y_2, \dots$ be independent, square integrable random variables with $\mathbf{E} [Y_n] = 1$ for all $n \in \mathbb{N}$. Let $X_n := \prod_{i=1}^n Y_i$ for $n \in \mathbb{N}_0$. Then $X = (X_n)_{n \in \mathbb{N}_0}$ is a square integrable martingale with respect to $\mathbb{F} = \sigma(X)$ (why?) and

$$\mathbf{E} [(X_n - X_{n-1})^2 \mid \mathcal{F}_{n-1}] = \mathbf{E} [(Y_n - 1)^2 X_{n-1}^2 \mid \mathcal{F}_{n-1}] = \text{Var} [Y_n] X_{n-1}^2.$$

Hence $\langle X \rangle_n = \sum_{i=1}^n \text{Var} [Y_i] X_{i-1}^2$. We see that the square variation process can indeed be a truly random process.

Example 10.8 Let $(X_n)_{n \in \mathbb{N}_0}$ be the one-dimensional symmetric simple random walk

$$X_n = \sum_{i=1}^n R_i \quad \text{for all } n \in \mathbb{N}_0,$$

where $R_1, R_2, R_3, \dots$ are i.i.d. and $\sim \text{Rad}_{1/2}$; that is,

$$\mathbf{P} [R_i = 1] = 1 - \mathbf{P} [R_i = -1] = \frac{1}{2}.$$

Clearly, X is a martingale and hence $|X|$ is a submartingale. Let $|X| = M + A$ be Doob's decomposition of $|X|$. Then

$$A_n = \sum_{i=1}^n (\mathbf{E} [|X_i| \mid \mathcal{F}_{i-1}] - |X_{i-1}|).$$

Now

$$|X_i| = \begin{cases} |X_{i-1}| + R_i, & \text{if } X_{i-1} > 0, \\ |X_{i-1}| - R_i, & \text{if } X_{i-1} < 0, \\ 1, & \text{if } X_{i-1} = 0. \end{cases}$$

Therefore,

$$\mathbf{E} [|X_i| \mid \mathcal{F}_{i-1}] = \begin{cases} |X_{i-1}|, & \text{if } |X_{i-1}| \neq 0 \\ 1, & \text{if } |X_{i-1}| = 0 \end{cases}$$

The process

$$A_n = \# \{i \leq n-1 : |X_i| = 0\}$$

is the so-called local time of X at 0 . We conclude that (since $\mathbf{P}[X_{2j} = 0] = \binom{2j}{j}4^{-j}$ and $\mathbf{P}[X_{2j+1} = 0] = 0$)

$$\begin{aligned} \mathbf{E}[|X_n|] &= \mathbf{E}[\# \{i \leq n-1 : X_i = 0\}] \\ &= \sum_{i=0}^{n-1} \mathbf{P}[X_i = 0] = \sum_{j=0}^{\lfloor (n-1)/2 \rfloor} \binom{2j}{j} 4^{-j}. \end{aligned}$$

Example 10.9 We want to generalize the preceding example further. Evidently, we did not use (except in the last formula) the fact that X is a random walk. Rather, we just used the fact that the differences $(\Delta X)_n := X_n - X_{n-1}$ take only the values -1 and +1 . Hence, now let X be a martingale with $|X_n - X_{n-1}| = 1$ almost surely for all $n \in \mathbb{N}$ and with $X_0 = x_0 \in \mathbb{Z}$ almost surely. Let $f : \mathbb{Z} \rightarrow \mathbb{R}$ be an arbitrary map. Then $Y := (f(X_n))_{n \in \mathbb{N}_0}$ is an integrable adapted process (since $|f(X_n)| \leq \max_{x \in \{x_0-n, \dots, x_0+n\}} |f(x)|$). In order to compute Doob's decomposition of Y , define the first and second discrete derivatives of f :

$$f'(x) := \frac{f(x+1) - f(x-1)}{2}$$

and

$$f''(x) := f(x-1) + f(x+1) - 2f(x).$$

Further, let $F'_n := f'(X_{n-1})$ and $F''_n := f''(X_{n-1})$. By computing the cases $X_n = X_{n-1} - 1$ and $X_n = X_{n-1} + 1$ separately, we see that for all $n \in \mathbb{N}$

$$\begin{aligned} f(X_n) - f(X_{n-1}) &= \frac{f(X_{n-1}+1) - f(X_{n-1}-1)}{2} (X_n - X_{n-1}) \\ &\quad + \frac{1}{2} f(X_{n-1}-1) + \frac{1}{2} f(X_{n-1}+1) - f(X_{n-1}) \\ &= f'(X_{n-1})(X_n - X_{n-1}) + \frac{1}{2} f''(X_{n-1}) \\ &= F'_n \cdot (X_n - X_{n-1}) + \frac{1}{2} F''_n \end{aligned}$$

Summing up, we get the discrete Itô formula:

$$\begin{aligned} f(X_n) &= f(x_0) + \sum_{i=1}^n f'(X_{i-1})(X_i - X_{i-1}) + \sum_{i=1}^n \frac{1}{2} f''(X_{i-1}) \\ &= f(x_0) + (F' \cdot X)_n + \sum_{i=1}^n \frac{1}{2} F''_i \end{aligned} \tag{10.4}$$

Here $F' \cdot X$ is the discrete stochastic integral (see Definition 9.37). Now $M := f(x_0) + F' \cdot X$ is a martingale by Theorem 9.39 since F' is predictable (and since $|F'_n| \leq \max_{x \in \{x_0-n, \dots, x_0+n\}} |f'(x)|$), and $A := (\sum_{i=1}^n \frac{1}{2} F''_i)_{n \in \mathbb{N}_0}$ is predictable. Hence $f(X) := (f(X_n))_{n \in \mathbb{N}_0} = M + A$ is the Doob decomposition of $f(X)$. In particular, $f(X)$ is a submartingale if $f''(x) \geq 0$ for all $x \in \mathbb{Z}$; that is, if f is convex. We knew this already from Theorem 9.35; however, here we could also quantify how much $f(X)$ differs from a martingale.

In the special cases $f(x) = x^2$ and $f(x) = |x|$, the second derivative is $f''(x) = 2$ and $f''(x) = 2 \cdot \mathbb{1}_{\{0\}}(x)$, respectively. Thus, from (10.4), we recover the statements of Theorem 10.4 and Example 10.8.

Later we will derive a formula similar to (10.4) for stochastic processes in continuous time (see Section 25.3).

5.10.2 10.2 Optional Sampling and Optional Stopping

Lemma 10.10 Let $I \subset \mathbb{R}$ be countable, let $(X_t)_{t \in I}$ be a martingale, let $T \in I$ and let τ be a stopping time with $\tau \leq T$. Then $X_\tau = \mathbf{E}[X_T | \mathcal{F}_\tau]$ and, in particular, $\mathbf{E}[X_\tau] = \mathbf{E}[X_0]$.

Proof It is enough to show that $\mathbf{E}[X_T \mathbb{1}_A] = \mathbf{E}[X_\tau \mathbb{1}_A]$ for all $A \in \mathcal{F}_\tau$. By the definition of $\mathcal{F}_\tau$, we have $\{\tau = t\} \cap A \in \mathcal{F}_t$ for all $t \in I$. Hence

$$\begin{aligned} \mathbf{E}[X_\tau \mathbb{1}_A] &= \sum_{t \leq T} \mathbf{E}[X_t \mathbb{1}_{\{\tau=t\} \cap A}] = \sum_{t \leq T} \mathbf{E}[\mathbf{E}[X_T | \mathcal{F}_t] \mathbb{1}_{\{\tau=t\} \cap A}] \\ &= \sum_{t \leq T} \mathbf{E}[X_T \mathbb{1}_A \mathbb{1}_{\{\tau=t\}}] = \mathbf{E}[X_T \mathbb{1}_A] \end{aligned}$$

Theorem 10.11 (Optional sampling theorem) Let $X = (X_n)_{n \in \mathbb{N}_0}$ be a supermartingale and let $\sigma \leq \tau$ be stopping times.

(i) Assume there exists a $T \in \mathbb{N}$ with $\tau \leq T$. Then

$$X_\sigma \geq \mathbf{E}[X_\tau | \mathcal{F}_\sigma]$$

and, in particular, $\mathbf{E}[X_\sigma] \geq \mathbf{E}[X_\tau]$. If X is a martingale, then equality holds in each case.

(ii) If X is nonnegative and if $\tau < \infty$ a.s., then we have $\mathbf{E}[X_\tau] \leq \mathbf{E}[X_0] < \infty$, $\mathbf{E}[X_\sigma] \leq \mathbf{E}[X_0] < \infty$ and $X_\sigma \geq \mathbf{E}[X_\tau | \mathcal{F}_\sigma]$.

(iii) Assume that, more generally, X is only adapted and integrable. Then X is a martingale if and only if $\mathbf{E}[X_\tau] = \mathbf{E}[X_0]$ for any bounded stopping time τ .

Proof (i) Let $X = M + A$ be Doob's decomposition of X . Hence A is predictable and monotone decreasing, $A_0 = 0$, and M is a martingale. Applying Lemma 10.10 to M yields

$$\begin{aligned} X_\sigma &= A_\sigma + M_\sigma = \mathbf{E}[A_\sigma + M_T | \mathcal{F}_\sigma] \\ &\geq \mathbf{E}[A_\tau + M_T | \mathcal{F}_\sigma] = \mathbf{E}[A_\tau + \mathbf{E}[M_T | \mathcal{F}_\tau] | \mathcal{F}_\sigma] \\ &= \mathbf{E}[A_\tau + M_\tau | \mathcal{F}_\sigma] = \mathbf{E}[X_\tau | \mathcal{F}_\sigma] \end{aligned}$$

Here we used $\mathcal{F}_\tau \supset \mathcal{F}_\sigma$, the tower property and the monotonicity of the conditional expectation (see Theorem 8.14).

(ii) We have $X_{\tau \wedge n} \xrightarrow{n \rightarrow \infty} X_\tau$ almost surely. By (i), we get $\mathbf{E}[X_{\tau \wedge n}] \leq \mathbf{E}[X_0]$ for any $n \in \mathbb{N}$. Using Fatou's lemma, we infer

$$\mathbf{E}[X_\tau] \leq \liminf_{n \rightarrow \infty} \mathbf{E}[X_{\tau \wedge n}] \leq \mathbf{E}[X_0] < \infty$$

Similarly, we can show that $\mathbf{E}[X_\sigma] \leq \mathbf{E}[X_0]$.

Now, let $m, n \in \mathbb{N}$ with $m \geq n$. Part (i) applied to the bounded stopping times $\tau \wedge m \geq \sigma \wedge n$ yields

$$X_{\sigma \wedge n} \geq \mathbf{E}[X_{\tau \wedge m} | \mathcal{F}_{\sigma \wedge n}]$$

Now $\{\sigma < n\} \cap A \in \mathcal{F}_{\sigma \wedge n}$ for $A \in \mathcal{F}_\sigma$. Hence

$$\mathbf{E}[X_\sigma \mathbb{1}_{\{\sigma < n\} \cap A}] = \mathbf{E}[X_{\sigma \wedge n} \mathbb{1}_{\{\sigma < n\} \cap A}] \geq \mathbf{E}[X_{\tau \wedge m} \mathbb{1}_{\{\sigma < n\} \cap A}]$$

Using Fatou's lemma, we get

$$\mathbf{E} [X_\tau \mathbb{1}_{\{\sigma < n\} \cap A}] \leq \liminf_{m \rightarrow \infty} \mathbf{E} [X_{\tau \wedge m} \mathbb{1}_{\{\sigma < n\} \cap A}] \leq \mathbf{E} [X_\sigma \mathbb{1}_{\{\sigma < n\} \cap A}].$$

Monotone convergence (for $n \rightarrow \infty$) thus yields $\mathbf{E} [X_\tau \mathbb{1}_A] \leq \mathbf{E} [X_\sigma \mathbb{1}_A]$.

(iii) If X is a martingale, then the claim follows from Lemma 10.10. Now assume that $\mathbf{E} [X_\tau] = \mathbf{E} [X_0]$ for any bounded stopping time τ . Let $t > s$ and $A \in \mathcal{F}_s$. It is enough to show that $\mathbf{E} [X_t \mathbb{1}_A] = \mathbf{E} [X_s \mathbb{1}_A]$. Define $\tau = s \mathbb{1}_A + t \mathbb{1}_{A^c}$. Then τ is a bounded stopping time. However, by assumption,

$$\mathbf{E} [X_t \mathbb{1}_A] = \mathbf{E} [X_t] - \mathbf{E} [X_t \mathbb{1}_{A^c}] = \mathbf{E} [X_0] - \mathbf{E} [X_\tau] + \mathbf{E} [X_s \mathbb{1}_A] = \mathbf{E} [X_s \mathbb{1}_A]$$

Corollary 10.12 Let X be a martingale (respectively a submartingale), and assume $(\tau_N)_{N \in \mathbb{N}}$ is a monotone increasing sequence of bounded stopping times (hence $\tau_N \leq T_N$, $N \in \mathbb{N}$ for some $T_N \in \mathbb{N}$). Then $(X_{\tau_N})_{N \in \mathbb{N}}$ is a martingale (respectively a submartingale) with respect to the filtration $(\mathcal{F}_{\tau_N})_{N \in \mathbb{N}}$.

Definition 10.13 (Stopped process) Let $I \subset \mathbb{R}$ be countable, let $(X_t)_{t \in I}$ be adapted and let τ be a stopping time. We define the stopped process X^τ by

$$X_t^\tau = X_{\tau \wedge t} \quad \text{for any } t \in I.$$

Further, let $\mathbb{F}^\tau$ be the filtration $\mathbb{F}^\tau = (\mathcal{F}_t^\tau)_{t \in I} = (\mathcal{F}_{\tau \wedge t})_{t \in I}$.

Remark 10.14 X^τ is adapted both to $\mathbb{F}$ and to $\mathbb{F}^\tau$.

Theorem 10.15 (Optional stopping) Let $(X_n)_{n \in \mathbb{N}_0}$ be a (sub-, super-) martingale with respect to $\mathbb{F}$ and let τ be a stopping time. Then X^τ is a (sub-, super-) martingale both with respect to $\mathbb{F}$ and with respect to $\mathbb{F}^\tau$.

Proof We give the proof only for the case where X is a submartingale. The other cases are similar since there $(-X)$ is a submartingale.

For each $n \in \mathbb{N}_0$, we have

$$\begin{aligned} \mathbf{E} [|X_n^\tau|] &\leq \mathbf{E} [\max \{|X_m| : m \leq n\}] \\ &\leq \mathbf{E} [|X_0|] + \dots + \mathbf{E} [|X_n|] < \infty. \end{aligned}$$

Hence X^τ is integrable.

Let X be a submartingale. Since $\{\tau > n-1\} \in \mathcal{F}_{n-1}$, we have

$$\begin{aligned} \mathbf{E} [X_n^\tau - X_{n-1}^\tau \mid \mathcal{F}_{n-1}] &= \mathbf{E} [X_{\tau \wedge n} - X_{\tau \wedge (n-1)} \mid \mathcal{F}_{n-1}] \\ &= \mathbf{E} [(X_n - X_{n-1}) \mathbb{1}_{\{\tau > n-1\}} \mid \mathcal{F}_{n-1}] \\ &= \mathbb{1}_{\{\tau > n-1\}} \mathbf{E} [X_n - X_{n-1} \mid \mathcal{F}_{n-1}] \\ &\geq 0, \quad \text{since } X \text{ is an } \mathbb{F}\text{-submartingale.} \end{aligned}$$

Therefore, X^τ is an $\mathbb{F}$ -submartingale. As X^τ is adapted to $\mathbb{F}^\tau$ and since $\mathbb{F}^\tau$ is the smaller filtration, X^τ is also an $\mathbb{F}^\tau$ -submartingale (see Remark 9.29).

Example 10.16 Let X be a symmetric simple random walk on $\mathbb{Z}$ (see Example 10.8). Let $a, b \in \mathbb{Z}$ with $a < 0, b > 0$ and let

$$\begin{aligned} \tau_a &= \inf \{t \geq 0 : X_t = a\}, \quad \tau_b = \inf \{t \geq 0 : X_t = b\} \quad \text{and} \\ \tau_{a,b} &= \tau_a \wedge \tau_b. \end{aligned}$$

$\tau_{a,b}$ is a stopping time by Lemma 9.18. Let $A = \{\tau_{a,b} = \tau_a\}$ be the event where X hits a before hitting b . We want to compute $\mathbf{P}[A]$. By Exercise 2.3.1, almost surely $\limsup_{n \rightarrow \infty} X_n = \infty$ and $\liminf_{n \rightarrow \infty} X_n = -\infty$. Therefore, almost surely $\tau_a < \infty$ and $\tau_b < \infty$. By the optional stopping theorem, $X^{\tau_{a,b}}$ is a martingale. Since $\tau_{a,b} \wedge n \xrightarrow{n \rightarrow \infty} \tau_{a,b}$ almost surely, we

get $X_n \xrightarrow{\tau_{a,b} \rightarrow \infty} X_{\tau_{a,b}}$ almost surely. As $|X_n^{\tau_{a,b}}|$ is bounded by $b - a$, we can infer that $X_n^{\tau_{a,b}} \xrightarrow{n \rightarrow \infty} X_{\tau_{a,b}}$ also in L^1 . Thus

$$\begin{aligned} 0 &= \lim_{n \rightarrow \infty} \mathbf{E} [X_n^{\tau_{a,b}}] = \mathbf{E} [X_{\tau_{a,b}}] = a \cdot \mathbf{P} [\tau_{a,b} = \tau_a] + b \cdot \mathbf{P} [\tau_{a,b} = \tau_b] \\ &= b + (a - b) \mathbf{P} [\tau_{a,b} = \tau_a] \end{aligned}$$

We conclude that $\mathbf{P} [\tau_{a,b} = \tau_a] = \frac{b}{b-a}$.

Example 10.17 Finally, we use our machinery in order to compute $\mathbf{E} [\tau_{a,b}]$ and $\mathbf{E} [\tau_a]$. The square variation process $\langle X \rangle$ (compare Definition 10.3) is given by

$$\langle X \rangle_n = \sum_{i=1}^n \mathbf{E} [(X_i - X_{i-1})^2 \mid \mathcal{F}_{i-1}] = n$$

hence $(X_n^2 - n)_{n \in \mathbb{N}_0}$ is a martingale. By the optional stopping theorem,

$$0 = \mathbf{E} [X_{\tau_{a,b} \wedge n}^2 - (\tau_{a,b} \wedge n)] \quad \text{for all } n \in \mathbb{N}_0$$

Monotone convergence yields

$$\mathbf{E} [\tau_{a,b}] = \mathbf{E} [X_{\tau_{a,b}}^2] = a^2 \mathbf{P} [\tau_{a,b} = \tau_a] + b^2 \mathbf{P} [\tau_{a,b} = \tau_b] = |a| \cdot b.$$

In order to compute $\mathbf{E} [\tau_a]$, note that $\tau_{a,b} \uparrow \tau_a$ almost surely if $b \rightarrow \infty$. The monotone convergence theorem thus yields $\mathbf{E} [\tau_a] = \lim_{b \rightarrow \infty} \mathbf{E} [\tau_{a,b}] = \infty$.

Remark 10.18 Evidently, $X_{\tau_b} = b > 0$. Hence $X_0 < \mathbf{E} [X_{\tau_b} \mid \mathcal{F}_0] = b$. The claim of the optional sampling theorem may thus fail, in general, if the stopping time is unbounded.

Example 10.19 (Gambler's ruin problem) Consider a game of two persons, A and B . In each round, a coin is tossed. Depending on the outcome, either A gets a euro from B or vice versa. The game endures until one of the players is ruined. For simplicity, we assume that in the beginning A has $k_A \in \mathbb{N}$ euros while B has $k_B = N - k_A$ euros, where $N \in \mathbb{N}$, $N \geq k_A$. We want to know the probability of B 's ruin. In Example 10.16 we saw that for a fair coin this probability is k_A/N . Now we allow the coin to be unfair.

Hence, let $Y_1, Y_2, \dots$ be i.i.d. and $\sim \text{Rad}_p$ (that is, $\mathbf{P} [Y_i = 1] = 1 - \mathbf{P} [Y_i = -1] = p$) for some $p \in (0, 1) \setminus \{\frac{1}{2}\}$. Denote by $X_n := k_B + \sum_{i=1}^n Y_i$ the running total for B after n rounds, where formally we assume that the game continues even after one player is ruined. We define as above τ_0, τ_N and $\tau_{0,N}$ as the times of first entrance of X into $\{0\}, \{N\}$ and $\{0, N\}$, respectively. The ruin probability of B thus is $p_B^N := \mathbf{P} [\tau_{0,N} = \tau_0]$. Since X is not a martingale (except for the case $p = \frac{1}{2}$ that was excluded), we use a trick to construct a martingale. Define a new process Z by $Z_n := r^{X_n} = r^{k_B} \prod_{i=1}^n r^{Y_i}$, where $r > 0$ has to be chosen so that Z becomes a martingale. By Example 9.31, this is the case if and only if

$$\mathbf{E} [r^{Y_1}] = pr + (1 - p)r^{-1} = 1;$$

hence, if $r = 1$ or $r = \frac{1-p}{p}$. Evidently, the choice $r = 1$ is useless (as Z does not yield any information on X); hence we assume $r = \frac{1-p}{p}$. We thus get

$$\tau_0 = \inf \{n \in \mathbb{N}_0 : Z_n = 1\} \quad \text{and} \quad \tau_N = \inf \{n \in \mathbb{N}_0 : Z_n = r^N\}.$$

(Note that here we cannot argue as above in order to show that $\tau_0 < \infty$ and $\tau_N < \infty$ almost surely. In fact, for $p \neq \frac{1}{2}$, only one of the statements holds. However, using, for example, the strong law of large numbers, we obtain that $\liminf_{n \rightarrow \infty} X_n = \infty$ (and thus $\tau_N < \infty$) almost surely if $p > \frac{1}{2}$. Similarly, $\tau_0 < \infty$ almost surely if $p < \frac{1}{2}$.) As in Example 10.16, the optional

stopping theorem yields $r^{k_B} = Z_0 = \mathbf{E}[Z_{\tau_0, N}] = p_B^N + (1 - p_B^N) r^N$. Therefore, the probability of B 's ruin is

$$p_B^N = \frac{r^{k_B} - r^N}{1 - r^N} \quad (10.5)$$

If the game is advantageous for B (that is, $p > \frac{1}{2}$), then $r < 1$. In this case, in the limit $N \rightarrow \infty$ (with constant k_B),

$$p_B^\infty := \lim_{N \rightarrow \infty} p_B^N = r^{k_B} \quad (10.6)$$

Exercise 10.2.1 Let X be a square integrable martingale with square variation process $\langle X \rangle$. Let τ be a finite stopping time. Show the following:

(i) If $\mathbf{E}[\langle X \rangle_\tau] < \infty$, then

$$\mathbf{E}[(X_\tau - X_0)^2] = \mathbf{E}[\langle X \rangle_\tau] \quad \text{and} \quad \mathbf{E}[X_\tau] = \mathbf{E}[X_0]. \quad (10.7)$$

(ii) If $\mathbf{E}[\langle X \rangle_\tau] = \infty$, then both equalities in (10.7) may fail.

Exercise 10.2.2 We consider a situation that is more general than the one in the preceding example by assuming only that $Y_1, Y_2, \dots$ are i.i.d. integrable random variables that are not almost surely constant (and $X_n = Y_1 + \dots + Y_n$). We further assume that there is a $\delta > 0$ such that $\mathbf{E}[\exp(\theta Y_1)] < \infty$ for all $\theta \in (-\delta, \delta)$. Define a map $\psi : (-\delta, \delta) \rightarrow \mathbb{R}$ by $\theta \mapsto \log(\mathbf{E}[\exp(\theta Y_1)])$ and the process Z^θ by $Z_n^\theta := \exp(\theta X_n - n\psi(\theta))$ for $n \in \mathbb{N}_0$. Show the following:

(i) Z^θ is a martingale for all $\theta \in (-\delta, \delta)$.

(ii) ψ is strictly convex.

(iii) $\mathbf{E}[\sqrt{Z_n^\theta}] \xrightarrow{n \rightarrow \infty} 0$ for $\theta \neq 0$.

(iv) $Z_n^\theta \xrightarrow{n \rightarrow \infty} 0$ almost surely.

We may interpret Y_n as the difference between the premiums and the payments of an insurance company at time n . If the initial capital of the company is $k_0 > 0$, then $k_0 + X_n$ is the account balance at time n . We are interested in the ruin probability

$$p(k_0) = \mathbf{P}[\inf\{X_n + k_0 : n \in \mathbb{N}_0\} < 0]$$

depending on the initial capital.

It can be assumed that the premiums are calculated such that $\mathbf{E}[Y_1] > 0$. Show that if the equation $\psi(\theta) = 0$ has a solution $\theta^* \neq 0$, then $\theta^* < 0$. Show further that in this case, the Cramér-Lundberg inequality holds:

$$p(k_0) \leq \exp(\theta^* k_0). \quad (10.8)$$

Equality holds if $k_0 \in \mathbb{N}$ and if Y_i assumes only the values -1 and 1. In this case, we get (10.6) with $r = \exp(\theta^*)$.

5.10.3 10.3 Uniform Integrability and Optional Sampling

We extend the optional sampling theorem to unbounded stopping times. We will see that this is possible if the underlying martingale is uniformly integrable (compare Definition 6.16).

Lemma 10.20 Let $(X_n)_{n \in \mathbb{N}_0}$ be a uniformly integrable martingale. Then the family $(X_\tau : \tau \text{ is a finite stopping time})$ is uniformly integrable.

Proof By Theorem 6.19, there exists a monotone increasing, convex function $f : [0, \infty) \rightarrow [0, \infty)$ with

$$\liminf_{x \rightarrow \infty} f(x)/x = \infty \quad \text{and} \quad L := \sup_{n \in \mathbb{N}_0} \mathbf{E}[f(|X_n|)] < \infty$$

If $\tau < \infty$ is a finite stopping time, then by the optional sampling theorem for bounded stopping times (Theorem 10.11 with $\tau = n$ and $\sigma = \tau \wedge n$), $\mathbf{E}[X_n | \mathcal{F}_{\tau \wedge n}] = X_{\tau \wedge n}$. Since $\{\tau \leq n\} \in \mathcal{F}_{\tau \wedge n}$, Jensen's inequality yields

$$\begin{aligned} \mathbf{E}[f(|X_\tau|) \mathbb{1}_{\{\tau \leq n\}}] &= \mathbf{E}[f(|X_{\tau \wedge n}|) \mathbb{1}_{\{\tau \leq n\}}] \\ &\leq \mathbf{E}[\mathbf{E}[f(|X_n|) | \mathcal{F}_{\tau \wedge n}] \mathbb{1}_{\{\tau \leq n\}}] \\ &= \mathbf{E}[f(|X_n|) \mathbb{1}_{\{\tau \leq n\}}] \leq L. \end{aligned}$$

Hence $\mathbf{E}[f(|X_\tau|)] \leq L$. By Theorem 6.19, the family

$$\{X_\tau, \tau \text{ is a finite stopping time}\}$$

is uniformly integrable.

Theorem 10.21 (Optional sampling and uniform integrability) Let $(X_n, n \in \mathbb{N}_0)$ be a uniformly integrable martingale (respectively supermartingale) and let $\sigma \leq \tau$ be finite stopping times. Then $\mathbf{E}[|X_\tau|] < \infty$ and $X_\sigma = \mathbf{E}[X_\tau | \mathcal{F}_\sigma]$ (respectively $X_\sigma \geq \mathbf{E}[X_\tau | \mathcal{F}_\sigma]$).

Proof First let X be a martingale. We have $\{\sigma \leq n\} \cap F \in \mathcal{F}_{\sigma \wedge n}$ for all $F \in \mathcal{F}_\sigma$. Hence, by the optional sampling theorem (Theorem 10.11),

$$\mathbf{E}[X_{\tau \wedge n} \mathbb{1}_{\{\sigma \leq n\} \cap F}] = \mathbf{E}[X_{\sigma \wedge n} \mathbb{1}_{\{\sigma \leq n\} \cap F}].$$

By Lemma 10.20, $(X_{\sigma \wedge n}, n \in \mathbb{N}_0)$ and thus $(X_{\sigma \wedge n} \mathbb{1}_{\{\sigma \leq n\} \cap F}, n \in \mathbb{N}_0)$ are uniformly integrable. Similarly, this holds for X_τ . Therefore, by Theorem 6.25,

$$\begin{aligned} \mathbf{E}[X_\tau \mathbb{1}_F] &= \lim_{n \rightarrow \infty} \mathbf{E}[X_{\tau \wedge n} \mathbb{1}_{\{\sigma \leq n\} \cap F}] = \lim_{n \rightarrow \infty} \mathbf{E}[X_{\sigma \wedge n} \mathbb{1}_{\{\sigma \leq n\} \cap F}] \\ &= \mathbf{E}[X_\sigma \mathbb{1}_F]. \end{aligned}$$

We conclude that $\mathbf{E}[X_\tau | \mathcal{F}_\sigma] = X_\sigma$.

Now let X be a supermartingale and let $X = M + A$ be its Doob decomposition; that is, M is a martingale and $A \leq 0$ is predictable and decreasing. Since

$$\mathbf{E}[|A_n|] = \mathbf{E}[-A_n] \leq \mathbf{E}[|X_n - X_0|] \leq \mathbf{E}[|X_0|] + \sup_{m \in \mathbb{N}_0} \mathbf{E}[|X_m|] < \infty,$$

we have $A_n \downarrow A_\infty$ for some $A_\infty \leq 0$ with $\mathbf{E}[-A_\infty] < \infty$ (by the monotone convergence theorem). Hence A and thus $M = X - A$ are uniformly integrable (Theorem 6.18(ii)). Therefore,

$$\mathbf{E}[|X_\tau|] \leq \mathbf{E}[-A_\tau] + \mathbf{E}[|M_\tau|] \leq \mathbf{E}[-A_\infty] + \mathbf{E}[|M_\tau|] < \infty.$$

Furthermore,

$$\begin{aligned} \mathbf{E}[X_\tau | \mathcal{F}_\sigma] &= \mathbf{E}[M_\tau | \mathcal{F}_\sigma] + \mathbf{E}[A_\tau | \mathcal{F}_\sigma] \\ &= M_\sigma + A_\sigma + \mathbf{E}[(A_\tau - A_\sigma) | \mathcal{F}_\sigma] \\ &\leq M_\sigma + A_\sigma = X_\sigma \end{aligned}$$

Corollary 10.22 Let X be a uniformly integrable martingale (respectively supermartingale) and let $\tau_1 \leq \tau_2 \leq \dots$ be finite stopping times. Then $(X_{\tau_n})_{n \in \mathbb{N}}$ is a martingale (respectively supermartingale).

5.11 11 Martingale Convergence Theorems and Their Applications

We became familiar with martingales $X = (X_n)_{n \in \mathbb{N}_0}$ as fair games and found that under certain transformations (optional stopping, discrete stochastic integral) martingales turn into martingales. In this chapter, we will see that under weak conditions (non-negativity or uniform integrability) martingales converge almost surely. Furthermore, the martingale structure implies L^p -convergence under assumptions that are (formally) weaker than those of Chapter 7. The basic ideas of this chapter are Doob's inequality (Theorem 11.2) and the upcrossing inequality (Lemma 11.3).

5.11.1 11.1 Doob's Inequality

With Kolmogorov's inequality (Theorem 5.28), we became acquainted with an inequality that bounds the probability of large values of the maximum of a square integrable process with independent centered increments. Here we want to improve this inequality in two directions. On the one hand, we replace the independent increments by the assumption that the process of partial sums is a martingale. On the other hand, we can manage with less than second moments; alternatively, we can get better bounds if we have higher moments.

Let $I \subset \mathbb{N}_0$ and let $X = (X_n)_{n \in I}$ be a stochastic process. For $n \in \mathbb{N}$, we denote

$$X_n^* = \sup \{X_k : k \leq n\} \quad \text{and} \quad |X|_n^* = \sup \{|X_k| : k \leq n\}.$$

Lemma 11.1 If X is a submartingale, then, for all $\lambda > 0$,

$$\lambda \mathbf{P}[X_n^* \geq \lambda] \leq \mathbf{E}[X_n \mathbb{1}_{\{X_n^* \geq \lambda\}}] \leq \mathbf{E}[|X_n| \mathbb{1}_{\{X_n^* \geq \lambda\}}].$$

Proof The second inequality is trivial. For the first one, let

$$\tau := \inf \{k \in I : X_k \geq \lambda\} \wedge n.$$

By Theorem 10.11 (optional sampling theorem),

$$\begin{aligned} \mathbf{E}[X_n] &\geq \mathbf{E}[X_\tau] = \mathbf{E}[X_\tau \mathbb{1}_{\{X_n^* \geq \lambda\}}] + \mathbf{E}[X_\tau \mathbb{1}_{\{X_n^* < \lambda\}}] \\ &\geq \lambda \mathbf{P}[X_n^* \geq \lambda] + \mathbf{E}[X_n \mathbb{1}_{\{X_n^* < \lambda\}}]. \end{aligned}$$

(Note that $\tau = n$ if $X_n^* < \lambda$.) Now subtract $\mathbf{E}[X_n \mathbb{1}_{\{X_n^* < \lambda\}}]$.

Theorem 11.2 (Doob's L^p -inequality) Let X be a martingale or a positive submartingale.

(i) For any $p \geq 1$ and $\lambda > 0$,

$$\lambda^p \mathbf{P}[|X|_n^* \geq \lambda] \leq \mathbf{E}[|X_n|^p]$$

(ii) For any $p > 1$,

$$\mathbf{E}[|X_n|^p] \leq \mathbf{E}[(|X|_n^*)^p] \leq \left(\frac{p}{p-1}\right)^p \mathbf{E}[|X_n|^p].$$

Proof We follow the proof in [144].

(i) By Theorem 9.35, $(|X_n|^p)_{n \in I}$ is a submartingale, and the claim follows by Lemma 11.1.

(ii) The first inequality is trivial. For the second inequality, we may assume that $\mathbf{E}[|X_n|^p] < \infty$. Note that, by Lemma 11.1,

$$\lambda \mathbf{P}[|X|_n^* \geq \lambda] \leq \mathbf{E}[|X_n| \mathbb{1}_{\{|X|_n^* \geq \lambda\}}].$$

Hence, for any $K > 0$,

$$\begin{aligned}
\mathbf{E}[(|X|_n^* \wedge K)^p] &= \mathbf{E} \left[\int_0^{|X|_n^* \wedge K} p\lambda^{p-1} d\lambda \right] \\
&= \mathbf{E} \left[\int_0^K p\lambda^{p-1} \mathbb{1}_{\{|X|_n^* \geq \lambda\}} d\lambda \right] \\
&= \int_0^K p\lambda^{p-1} \mathbf{P}[|X|_n^* \geq \lambda] d\lambda \\
&\leq \int_0^K p\lambda^{p-2} \mathbf{E}[|X|_n \mathbb{1}_{\{|X|_n^* \geq \lambda\}}] d\lambda \\
&= p\mathbf{E} \left[|X|_n \int_0^{|X|_n^* \wedge K} \lambda^{p-2} d\lambda \right] \\
&= \frac{p}{p-1} \mathbf{E}[|X|_n \cdot (|X|_n^* \wedge K)^{p-1}]
\end{aligned}$$

Hölder's inequality then yields

$$\mathbf{E}[(|X|_n^* \wedge K)^p] \leq \frac{p}{p-1} \mathbf{E}[(|X|_n^* \wedge K)^p]^{(p-1)/p} \cdot \mathbf{E}[|X|_n^p]^{1/p}.$$

We raise both sides to the p th power and divide by $\mathbf{E}[(|X|_n^* \wedge K)^p]^{p-1}$ (here we need the truncation at K to make sure we divide by a finite number) to obtain

$$\mathbf{E}[(|X|_n^* \wedge K)^p] \leq \left(\frac{p}{p-1} \right)^p \mathbf{E}[|X|_n^p].$$

Finally, let $K \rightarrow \infty$. $\square$

Exercise 11.1.1 Let $(X_n)_{n \in \mathbb{N}_0}$ be a submartingale or a supermartingale. Use Theorem 11.2 and Doob's decomposition to show that, for all $n \in \mathbb{N}$ and $\lambda > 0$,

$$\lambda \mathbf{P}[|X|_n^* \geq \lambda] \leq 12\mathbf{E}[|X_0|] + 9\mathbf{E}[|X_n|].$$

5.11.2 11.2 Martingale Convergence Theorems

In this section, we present the usual martingale convergence theorems and give a few small examples. We start with the core of the martingale convergence theorems, the so-called upcrossing inequality.

Let $\mathbb{F} = (\mathcal{F}_n)_{n \in \mathbb{N}_0}$ be a filtration and $\mathcal{F}_\infty = \sigma(\bigcup_{n \in \mathbb{N}_0} \mathcal{F}_n)$. Let $(X_n)_{n \in \mathbb{N}_0}$ be realvalued and adapted to $\mathbb{F}$. Let $a, b \in \mathbb{R}$ with $a < b$. If we think of X as a stock price, it would be a sensible trading strategy to buy the stock when its price has fallen below a and to sell it when it exceeds b at least if we knew for sure that the price would always rise above the level b again. Each time the price makes such an upcrossing from a to b , we make a profit of at least $b - a$. If we get a bound on the maximal profit we can make, dividing it by $b - a$ gives a bound on the maximal number of such upcrossings. If this number is finite for all $a < b$, then the price has to converge as $n \rightarrow \infty$.

Let us get into the details. Define stopping times $\sigma_0 \equiv 0$ and

$$\begin{aligned}
\tau_k &:= \inf \{n \geq \sigma_{k-1} : X_n \leq a\} \quad \text{for } k \in \mathbb{N}, \\
\sigma_k &:= \inf \{n \geq \tau_k : X_n \geq b\} \quad \text{for } k \in \mathbb{N}.
\end{aligned}$$

Note that $\tau_k = \infty$ if $\sigma_{k-1} = \infty$, and $\sigma_k = \infty$ if $\tau_k = \infty$. We say that X has its k th upcrossing over $[a, b]$ between τ_k and σ_k if $\sigma_k < \infty$. For $n \in \mathbb{N}$, define

$$U_n^{a,b} := \sup \{k \in \mathbb{N}_0 : \sigma_k \leq n\}$$

as the number of upcrossings over $[a, b]$ until time n .

Lemma 11.3 (Upcrossing inequality) Let $(X_n)_{n \in \mathbb{N}_0}$ be a submartingale. Then

$$\mathbf{E} [U_n^{a,b}] \leq \frac{\mathbf{E} [(X_n - a)^+] - \mathbf{E} [(X_0 - a)^+]}{b - a}$$

Proof Recall the discrete stochastic integral $H \cdot X$ from Definition 9.37. Formally, the intimated trading strategy H is described for $m \in \mathbb{N}_0$ by

$$H_m := \begin{cases} 1, & \text{if } m \in \{\tau_k + 1, \dots, \sigma_k\} \text{ for some } k \in \mathbb{N} \\ 0, & \text{else.} \end{cases}$$

H is nonnegative and predictable since, for all $m \in \mathbb{N}$,

$$\{H_m = 1\} = \bigcup_{k=1}^{\infty} (\{\tau_k \leq m-1\} \cap \{\sigma_k > m-1\})$$

and each of the events is in $\mathcal{F}_{m-1}$. Define $Y = \max(X, a)$. If $k \in \mathbb{N}$ and $\sigma_k < \infty$, then clearly $Y_{\sigma_i} - Y_{\tau_i} = Y_{\sigma_i} - a \geq b - a$ for all $i \leq k$; hence

$$(H \cdot Y)_{\sigma_k} = \sum_{i=1}^k \sum_{j=\tau_i+1}^{\sigma_i} (Y_j - Y_{j-1}) = \sum_{i=1}^k (Y_{\sigma_i} - Y_{\tau_i}) \geq k(b - a).$$

For $j \in \{\sigma_k, \dots, \tau_{k+1}\}$, we have $(H \cdot Y)_j = (H \cdot Y)_{\sigma_k}$. On the other hand, for $j \in \{\tau_k + 1, \dots, \sigma_k\}$, we have $(H \cdot Y)_j \geq (H \cdot Y)_{\tau_k} = (H \cdot Y)_{\sigma_{k-1}}$. Hence $(H \cdot Y)_n \geq (b - a)U_n^{a,b}$ for all $n \in \mathbb{N}$.

By Corollary 9.34, Y is a submartingale, and (by Theorem 9.39) so are $H \cdot Y$ and $(1 - H) \cdot Y$. Now $Y_n - Y_0 = (1 \cdot Y)_n = (H \cdot Y)_n + ((1 - H) \cdot Y)_n$; hence

$$\mathbf{E} [Y_n - Y_0] \geq \mathbf{E} [(H \cdot Y)_n] \geq (b - a)\mathbf{E} [U_n^{a,b}].$$

Theorem 11.4 (Martingale convergence theorem) Let $(X_n)_{n \in \mathbb{N}_0}$ be a submartingale with $\sup \{\mathbf{E} [X_n^+] : n \geq 0\} < \infty$. Then there exists an $\mathcal{F}_\infty$ -measurable random variable X_∞ with $\mathbf{E} [|X_\infty|] < \infty$ and $X_n \xrightarrow{n \rightarrow \infty} X_\infty$ almost surely.

Proof Let $a < b$. Since $\mathbf{E} [(X_n - a)^+] \leq |a| + \mathbf{E} [X_n^+]$, by Lemma 11.3,

$$\mathbf{E} [U_n^{a,b}] \leq \frac{|a| + \mathbf{E} [X_n^+]}{b - a}$$

Manifestly, the monotone limit $U^{a,b} := \lim_{n \rightarrow \infty} U_n^{a,b}$ exists. By assumption, we have $\mathbf{E} [U^{a,b}] = \lim_{n \rightarrow \infty} \mathbf{E} [U_n^{a,b}] < \infty$. In particular, $\mathbf{P} [U^{a,b} < \infty] = 1$. Define the $\mathcal{F}_\infty$ -measurable events

$$C^{a,b} = \left\{ \liminf_{n \rightarrow \infty} X_n < a \right\} \cap \left\{ \limsup_{n \rightarrow \infty} X_n > b \right\} \subset \{U^{a,b} = \infty\}$$

and

$$C = \bigcup_{\substack{a, b \in \mathbb{Q} \\ a < b}} C^{a,b}$$

Then $\mathbf{P}[C^{a,b}] = 0$ and thus also $\mathbf{P}[C] = 0$. However, by construction, $(X_n)_{n \in \mathbb{N}}$ is convergent on C^c . Hence there exists the almost sure limit $X_\infty = \lim_{n \rightarrow \infty} X_n$. Each X_n is $\mathcal{F}_\infty$ -measurable; hence X_∞ also is $\mathcal{F}_\infty$ -measurable.

By Fatou's lemma,

$$\mathbf{E}[X_\infty^+] \leq \sup \{ \mathbf{E}[X_n^+] : n \geq 0 \} < \infty$$

On the other hand (since X is a submartingale), again by Fatou's lemma,

$$\begin{aligned} \mathbf{E}[X_\infty^-] &\leq \liminf_{n \rightarrow \infty} \mathbf{E}[X_n^-] = \liminf_{n \rightarrow \infty} (\mathbf{E}[X_n^+] - \mathbf{E}[X_n]) \\ &\leq \sup \{ \mathbf{E}[X_n^+] : n \in \mathbb{N}_0 \} - \mathbf{E}[X_0] < \infty \end{aligned}$$

Corollary 11.5 If X is a nonnegative supermartingale, then there is an $\mathcal{F}_\infty$ -measurable random variable $X_\infty \geq 0$ with $\mathbf{E}[X_\infty] \leq \mathbf{E}[X_0]$ and $X_n \xrightarrow{n \rightarrow \infty} X_\infty$ a.s.

Proof The preceding theorem with $(-X)$ establishes X_∞ as the almost sure limit. Fatou's lemma yields

$$\mathbf{E}[X_\infty] \leq \liminf_{n \rightarrow \infty} \mathbf{E}[X_n] \leq \mathbf{E}[X_0]$$

Example 11.6 Let S_n be the account balance in the Petersburg game after the n th round (see Example 9.40). Then S is a martingale and $S_n \leq 1$ almost surely for any n . Hence the assumptions of Theorem 11.4 are fulfilled and $(S_n)_{n \in \mathbb{N}_0}$ converges to a finite random variable almost surely for $n \rightarrow \infty$. Since the account changes as long as stakes are put up (that is, as long as $S_n < 1$), we get $\lim_{n \rightarrow \infty} S_n = 1$ almost surely.

Since $\mathbf{E}[S_n] = 0$ for all $n \in \mathbb{N}_0$, this convergence cannot hold in L^1 . This observation tallies with the fact that S is not uniformly integrable.

For uniformly integrable martingales, a stronger convergence theorem holds.

Theorem 11.7 (Convergence theorem for uniformly integrable martingales) Let $(X_n)_{n \in \mathbb{N}_0}$ be a uniformly integrable $\mathbb{F}$ - (sub-, super-) martingale. Then there exists an $\mathcal{F}_\infty$ -measurable integrable random variable X_∞ with $X_n \xrightarrow{n \rightarrow \infty} X_\infty$ a.s. and in L^1 . Furthermore:

- $X_n = \mathbf{E}[X_\infty | \mathcal{F}_n]$ for all $n \in \mathbb{N}$ if X is a martingale.
- $X_n \leq \mathbf{E}[X_\infty | \mathcal{F}_n]$ for all $n \in \mathbb{N}$ if X is a submartingale.
- $X_n \geq \mathbf{E}[X_\infty | \mathcal{F}_n]$ for all $n \in \mathbb{N}$ if X is a supermartingale.

Remark 11.8 The statement of Theorem 11.7 can be reformulated as: The process $(X_n)_{n \in \mathbb{N}_0 \cup \{\infty\}}$ is a (sub-, super-) martingale with respect to $(\mathcal{F}_n)_{n \in \mathbb{N}_0 \cup \{\infty\}}$.

Proof We give the proof for the case where X is a submartingale. Uniform integrability implies $\sup \{ \mathbf{E}[X_n^+] : n \geq 0 \} < \infty$. By Theorem 11.4, the almost sure limit X_∞ exists. Hence $\mathbf{E}[|X_n - X_\infty|] \xrightarrow{n \rightarrow \infty} 0$ by Theorem 6.25. By Corollary 8.21, the L^1 convergence of (X_n) implies the L^1 -convergence of the conditional expectations: $\mathbf{E}[|\mathbf{E}[X_n | \mathcal{F}_m] - \mathbf{E}[X_\infty | \mathcal{F}_m]|] \xrightarrow{n \rightarrow \infty} 0$. Thus, by the triangle inequality,

$$\begin{aligned} &|\mathbf{E}[(\mathbf{E}[X_\infty | \mathcal{F}_m] - X_m)^-] - \mathbf{E}[(\mathbf{E}[X_n | \mathcal{F}_m] - X_m)^-]| \\ &\leq \mathbf{E}[|\mathbf{E}[X_\infty | \mathcal{F}_m] - \mathbf{E}[X_n | \mathcal{F}_m]|] \xrightarrow{n \rightarrow \infty} 0 \end{aligned}$$

As X is a submartingale, we have $(\mathbf{E}[X_n | \mathcal{F}_m] - X_m)^- = 0$ for $n \geq m$. Therefore, $\mathbf{E}[(\mathbf{E}[X_\infty | \mathcal{F}_m] - X_m)^-] = 0$ and thus $\mathbf{E}[X_\infty | \mathcal{F}_m] - X_m \geq 0$ almost surely.

Corollary 11.9 Let $X \geq 0$ be a martingale and let $X_\infty = \lim_{n \rightarrow \infty} X_n$. Then $\mathbf{E}[X_\infty] = \mathbf{E}[X_0]$ if and only if X is uniformly integrable.

Proof This is a direct consequence of Theorem 6.25.

Let $p \in [1, \infty)$. A real-valued stochastic process $(X_i)_{i \in I}$ is called L^p -bounded if $\sup_{i \in I} \mathbf{E}[|X_i|^p] < \infty$ (Definition 6.20). In general, for $(|X_i|^p)_{i \in I}$ to be uniformly integrable it is not enough that $(X_i)_{i \in I}$ be L^p -bounded. However, if X is a martingale and if $p > 1$, then Doob's inequality implies that the statements are equivalent. In particular, in this case, almost sure convergence implies convergence in L^p .

Theorem 11.10 (L^p -convergence theorem for martingales) Let $p > 1$ and let $(X_n)_{n \in \mathbb{N}_0}$ be an L^p -bounded martingale. Then there exists an $\mathcal{F}_\infty$ -measurable random variable X_∞ with $\mathbf{E}[|X_\infty|^p] < \infty$ and $X_n \xrightarrow{n \rightarrow \infty} X_\infty$ almost surely and in L^p . In particular, $(|X_n|^p)_{n \in \mathbb{N}_0}$ is uniformly integrable.

Proof By Corollary 6.21, X is uniformly integrable. Hence the almost sure limit X_∞ exists. By Doob's inequality (Theorem 11.2), for all $n \in \mathbb{N}$,

$$\mathbf{E}[\sup\{|X_k|^p : k \leq n\}] \leq \left(\frac{p}{p-1}\right)^p \mathbf{E}[|X_n|^p].$$

Therefore,

$$\mathbf{E}[\sup\{|X_k|^p : k \in \mathbb{N}_0\}] \leq \left(\frac{p}{p-1}\right)^p \sup\{\mathbf{E}[|X_n|^p] : n \in \mathbb{N}_0\} < \infty.$$

Hence, in particular, $(|X_n|^p)_{n \in \mathbb{N}_0}$ is uniformly integrable.

Since $|X_n - X_\infty|^p \leq 2^p \sup\{|X_n|^p : n \in \mathbb{N}_0\}$, dominated convergence yields

$$\mathbf{E}[|X_\infty|^p] < \infty \quad \text{and} \quad \mathbf{E}[|X_n - X_\infty|^p] \xrightarrow{n \rightarrow \infty} 0.$$

For the case of square integrable martingales, there is a convenient criterion for L^2 -boundedness that we record as a corollary (see Definition 10.3).

Corollary 11.11 Let X be a square integrable martingale with square variation process $\langle X \rangle$. Then the following four statements are equivalent:

- (i) $\sup_{n \in \mathbb{N}} \mathbf{E}[X_n^2] < \infty$.
- (ii) $\lim_{n \rightarrow \infty} \mathbf{E}[\langle X \rangle_n] < \infty$.
- (iii) X converges in L^2 .
- (iv) X converges almost surely and in L^2 .

Proof "(i) $\iff$ (ii)" Since $\text{Var}[X_n - X_0] = \mathbf{E}[\langle X \rangle_n]$ (see Theorem 10.4), X is bounded in L^2 if and only if (ii) holds.

"(iv) $\implies$ (iii) $\implies$ (i)" This is trivial.

"(i) $\implies$ (iv)" This is the statement of Theorem 11.10.

Remark 11.12 In general, the statement of Theorem 11.10 fails for $p = 1$. See Exercise 11.2.1.

Lemma 11.13 Let X be a square integrable martingale with square variation process $\langle X \rangle$, and let τ be a stopping time. Then the stopped process X^τ has square variation process $\langle X^\tau \rangle = \langle X \rangle^\tau := (\langle X \rangle_{\tau \wedge n})_{n \in \mathbb{N}_0}$.

Proof This is left as an exercise.

If in Corollary 11.11 we do not assume that the expectations of the square variation are bounded but only that the square variation is almost surely bounded, then we still get that X converges almost surely (albeit not in L^2).

Theorem 11.14 If X is a square integrable martingale with $\sup_{n \in \mathbb{N}} \langle X \rangle_n < \infty$ almost surely, then X converges almost surely.

Proof Without loss of generality, we can assume that $X_0 = 0$, otherwise consider the martingale $(X_n - X_0)_{n \in \mathbb{N}_0}$, which has the same square variation process. For $K > 0$, let

$$\tau_K := \inf\{n \in \mathbb{N} : \langle X \rangle_{n+1} \geq K\}$$

This is a stopping time since $\langle X \rangle$ is predictable. Evidently, $\sup_{n \in \mathbb{N}} \langle X \rangle_{\tau_K \wedge n} \leq K$ almost surely. By Corollary 11.11, the stopped process X^{τ_K} converges almost surely (and in L^2) to a random variable that we denote by $X_\infty^{\tau_K}$. By assumption, $\mathbf{P}[\tau_K = \infty] \rightarrow 1$ for $K \rightarrow \infty$; hence X converges almost surely.

Example 11.15 Let X be a symmetric simple random walk on $\mathbb{Z}$. That is, $X_n = \sum_{k=1}^n R_k$, where $R_1, R_2, \dots$ are i.i.d. and $\sim \text{Rad}_{1/2}$:

$$\mathbf{P}[R_1 = 1] = \mathbf{P}[R_1 = -1] = \frac{1}{2}$$

Then X is a martingale; however, $\limsup_{n \rightarrow \infty} X_n = \infty$ and $\liminf_{n \rightarrow \infty} X_n = -\infty$. Therefore, X does not even converge improperly. By the martingale convergence theorem, this is consonant with the fact that X is not uniformly integrable.

Example 11.16 (Voter model, due to [28, 75]) Consider a simple model that describes the behavior of opportunistic voters who are capable of only one out of two opinions, say 0 and 1. Let $\Lambda \subset \mathbb{Z}^d$ be a set that we interpret as the sites at each of which there is one voter. For simplicity, assume that $\Lambda = \{0, \dots, L-1\}^d$ for some $L \in \mathbb{N}$. Let $x(i) \in \{0, 1\}$ be the opinion of the voter at site $i \in \Lambda$ and denote by $x \in \{0, 1\}^\Lambda$ a generic state of the whole population. We now assume that the individual opinions may change at discrete time steps. At any time $n \in \mathbb{N}_0$, one site I_n out of Λ is chosen at random and the individual at that site reconsiders his or her opinion. To this end, the voter chooses a neighbor $I_n + N_n \in \Lambda$ (with periodic boundary conditions; that is, with addition modulo L in each coordinate) at random and adopts his or her opinion. We thus get a random sequence $(X_n)_{n \in \mathbb{N}_0}$ of states in $\{0, 1\}^\Lambda$ that represents the random evolution of the opinions of the whole colony. See Fig. 11.1 for a computer simulation of the voter model.

For a formal description of this model, let $(I_n)_{n \in \mathbb{N}}$ and $(N_n)_{n \in \mathbb{N}}$ be independent random variables. For any $n \in \mathbb{N}$, I_n is uniformly distributed on Λ and N_n is uniformly distributed on the set $\mathcal{N} := \{i \in \mathbb{Z}^d : \|i\|_2 = 1\}$ of the $2d$ nearest neighbors of the origin. Furthermore, $x = X_0 \in \{0, 1\}^\Lambda$ is the initial state. The states at later times are defined inductively by

$$X_n(i) = \begin{cases} X_{n-1}(i), & \text{if } I_n \neq i \\ X_{n-1}(I_n + N_n), & \text{if } I_n = i \end{cases}$$

Of course, the behavior over small periods of time is determined by the perils of randomness. However, in the long run, we might see certain patterns. To be more specific, the question is: In the long run, will there be a consensus of all individuals or will competing opinions persist?

Let $M_n := \sum_{i \in \Lambda} X_n(i)$ be the total number of individuals of opinion 1 at time n . Let $\mathbb{F}$ be the filtration $\mathbb{F} = (\mathcal{F}_n)_{n \in \mathbb{N}_0}$, where $\mathcal{F}_n = \sigma(I_k, N_k : k \leq n)$ for all $n \in \mathbb{N}_0$. Then M is adapted to $\mathbb{F}$ and

$$\begin{aligned} \mathbf{E}[M_n | \mathcal{F}_{n-1}] &= M_{n-1} - \mathbf{E}[X_{n-1}(I_n) | \mathcal{F}_{n-1}] + \mathbf{E}[X_{n-1}(I_n + N_n) | \mathcal{F}_{n-1}] \\ &= M_{n-1} - \sum_{i \in \Lambda} \mathbf{P}[I_n = i] X_{n-1}(i) + \sum_{i \in \Lambda} \mathbf{P}[I_n + N_n = i] X_{n-1}(i) \\ &= M_{n-1} \end{aligned}$$

since $\mathbf{P}[I_n = i] = \mathbf{P}[I_n + N_n = i] = L^{-d}$ for all $i \in \Lambda$. Hence M is a bounded $\mathbb{F}$ martingale and thus converges almost surely and in L^1 to a random variable M_∞ .




Fig. 11.1 Snapshot of a voter model on an 800×800 torus. The black dots are the Ones

Since M takes only integer values, there is a (random) n_0 such that $M_n = M_{n_0}$ for all $n \geq n_0$. However, then also $X_n = X_{n_0}$ for all $n \geq n_0$. Manifestly, no state x with $x \not\equiv 0$ and $x \not\equiv 1$ is stable. In fact, if x is not constant and if $i, j \in \Lambda$ are neighbors with $x(i) \neq x(j)$, then

$$\mathbf{P}[X_n \neq X_{n-1} \mid X_{n-1} = x] \geq \mathbf{P}[I_{n-1} = i, N_{n-1} = j - i] = L^{-d}(2d)^{-1}.$$

This implies $M_\infty \in \{0, L^d\}$. Now $\mathbf{E}[M_\infty] = M_0$; hence we have

$$\mathbf{P}[M_\infty = L^d] = \frac{M_0}{L^d} \quad \text{and} \quad \mathbf{P}[M_\infty = 0] = 1 - \frac{M_0}{L^d}.$$

Thus, eventually there will be a consensus of all individuals, and the probability that the surviving opinion is $e \in \{0, 1\}$ is the initial frequency of opinion e .

We could argue more formally to show that only the constant states are stable: Let $\langle M \rangle$ be the square variation process of M . Then

$$\begin{aligned} \langle M \rangle_n &= \sum_{k=1}^n \mathbf{E}[\mathbb{1}_{\{M_k \neq M_{k-1}\}} \mid \mathcal{F}_{k-1}] \\ &= \sum_{k=1}^n \mathbf{P}[X_{k-1}(I_k) \neq X_{k-1}(I_k + N_k) \mid \mathcal{F}_{k-1}]. \end{aligned}$$

Hence

$$\begin{aligned} L^{2d} &\geq \text{Var}[M_n] = \mathbf{E}[\langle M \rangle_n] \\ &= \sum_{k=1}^n \mathbf{P}[X_{k-1}(I_k) \neq X_{k-1}(I_k + N_k)] \\ &\geq (2d)^{-1} L^{-d} \sum_{k=1}^n \mathbf{P}[M_{k-1} \notin \{0, L^d\}] \end{aligned}$$

Therefore, $\sum_{k=1}^{\infty} \mathbf{P}[M_{k-1} \notin \{0, L^d\}] \leq 2dL^{3d} < \infty$, and so, by the Borel-Cantelli lemma, $M_\infty \in \{0, L^d\}$.

Example 11.17 (Radon-Nikodym theorem) With the aid of the martingale convergence theorem, we give an alternative proof of the Radon-Nikodym theorem (Corollary 7.34).

Let $(\Omega, \mathcal{F}, \mathbf{P})$ be a probability space and let Q be another probability measure on $(\Omega, \mathcal{A})$. We assume that $\mathcal{F}$ is countably generated; that is, there exist countably many sets $A_1, A_2, \dots \in \mathcal{F}$ such that $\mathcal{F} = \sigma(\{A_1, A_2, \dots\})$. For example, this is the case if $\mathcal{F}$ is the Borel σ -algebra on a Polish space. For the case $\Omega = \mathbb{R}^d$, one could take the open balls with rational radii, centered at points with rational coordinates (compare Remark 1.24).

We construct a filtration $\mathbb{F} = (\mathcal{F}_n)_{n \in \mathbb{N}}$ by letting $\mathcal{F}_n := \sigma(\{A_1, \dots, A_n\})$. Evidently, $\#\mathcal{F}_n < \infty$ for all $n \in \mathbb{N}$. More precisely, there exists a unique finite subset $Z_n \subset \mathcal{F}_n \setminus \{\emptyset\}$ such that $B = \bigsqcup_{\substack{C \in Z_n \\ C \subset B}} C$ for any $B \in \mathcal{F}_n$. Z_n decomposes $\mathcal{F}_n$ into its "atoms". Finally, define a stochastic process $(X_n)_{n \in \mathbb{N}}$ by

$$X_n := \sum_{C \in Z_n: \mathbf{P}[C] > 0} \frac{Q(C)}{\mathbf{P}[C]} \mathbb{1}_C$$

Clearly, X is adapted to $\mathbb{F}$. Let $B \in \mathcal{F}_n$ and $m \geq n$. For any $C \in Z_m$, either $C \cap B = \emptyset$ or $C \subset B$. Hence

$$\mathbf{E}[X_m \mathbb{1}_B] = \sum_{C \in Z_m: \mathbf{P}[C] > 0} \frac{Q(C)}{\mathbf{P}[C]} \mathbf{P}[C \cap B] = \sum_{C \in Z_m: C \subset B} Q(C) = Q(B). \quad (11.1)$$

In particular, X is an $\mathbb{F}$ -martingale.

Now assume that Q is absolutely continuous with respect to P . By Example 7.39, this implies that X is uniformly integrable. By the martingale convergence theorem, X converges $\mathbf{P}$ -almost surely and in $L^1(\mathbf{P})$ to a random variable X_∞ . By (11.1), we have $\mathbf{E}[X_\infty \mathbb{1}_B] = Q(B)$ for all $B \in \bigcup_{n \in \mathbb{N}} \mathcal{F}_n$ and thus also for all $B \in \mathcal{F}$. Therefore, X_∞ is the Radon-Nikodym density of Q with respect to $\mathbf{P}$.

Note that for this proof we did not presume the existence of conditional expectations (rather we constructed them explicitly for finite σ -algebras); that is, we did not resort to the Radon-Nikodym theorem in a hidden way.

It could be objected that this argument works only for probability measures. However, this flaw can easily be remedied. Let μ and ν be arbitrary (but nonzero) σ -finite measures. Then there exist measurable functions $g, h : \Omega \rightarrow (0, \infty)$ with $\int g d\mu = 1$ and $\int h d\nu = 1$. Define $\mathbf{P} = g\mu$ and $Q = h\nu$. Clearly, $Q \ll \mathbf{P}$ if $\nu \ll \mu$. In this case, $\frac{g}{h} X_\infty$ is a version of the Radon-Nikodym derivative $\frac{d\nu}{d\mu}$.

The restriction that $\mathcal{F}$ is countably generated can also be dropped. Using the approximation theorems for measures, it can be shown that there is always a countably generated σ -algebra $\mathcal{G} \subset \mathcal{F}$ such that for any $A \in \mathcal{F}$, there is a $B \in \mathcal{G}$ with $\mathbf{P}[A \Delta B] = 0$. This can be employed to prove the general case. We do not give the details but refer to [169, Chapter 14.13].

Exercise 11.2.1 For $p = 1$, the statement of Theorem 11.10 may fail. Give an example of a nonnegative martingale X with $\mathbf{E}[X_n] = 1$ for all $n \in \mathbb{N}$ but such that $X_n \xrightarrow{n \rightarrow \infty} 0$ almost surely.

Exercise 11.2.2 Let $X_1, X_2, \dots$ be independent, square integrable random variables with $\sum_{n=1}^{\infty} \frac{1}{n^2} \mathbf{Var}[X_n] < \infty$. Use the martingale convergence theorem to show the strong law of large numbers for $(X_n)_{n \in \mathbb{N}}$.

Exercise 11.2.3 Give an example of a square integrable martingale that converges almost surely but not in L^2 .

Exercise 11.2.4 Show that in Theorem 11.14 the converse implication may fail. That is, there exists a square integrable martingale X that converges almost surely but without $\lim_{n \rightarrow \infty} \langle X \rangle_n < \infty$ almost surely.

Exercise 11.2.5 Show the following converse of Theorem 11.14. Let $L > 0$ and let $(X_n)_{n \in \mathbb{N}}$ be a martingale with the property

$$|X_{n+1} - X_n| \leq L \quad \text{a.s.} \quad (11.2)$$

Define the events

$$\begin{aligned} C &:= \{(X_n)_{n \in \mathbb{N}} \text{ converges as } n \rightarrow \infty\}, \\ A^+ &:= \left\{ \limsup_{n \rightarrow \infty} X_n < \infty \right\}, \\ A^- &:= \left\{ \liminf_{n \rightarrow \infty} X_n > -\infty \right\}, \\ F &:= \left\{ \sup_{n \in \mathbb{N}} \langle X \rangle_n < \infty \right\}. \end{aligned}$$

Show that

$$C = A^+ = A^- = F(\text{mod } \mathbf{P}).$$

Here equality of events (mod $\mathbf{P}$) means that the events differ at most by a $\mathbf{P}$ -null set (see Definition 1.68(iii)).

Hint: Use the stopping times $\sigma_K = \inf \{n \in \mathbb{N} : |X_n| \geq K\}$; $\sigma_K^\pm = \inf \{n \in \mathbb{N} : \pm X_n \geq K\}$ and τ_K as in the proof of Theorem 11.14.

Exercise 11.2.6 Let the notation be as in Exercise 11.2.5. However, instead of (11.2) we make the weaker assumption

$$\mathbf{E} \left[\sup_{n \in \mathbb{N}} |X_{n+1} - X_n| \right] < \infty \quad (11.3)$$

Show that

$$C = A^+ = A^-(\text{mod } \mathbf{P})$$

Hint: Use suitable stopping times ϱ_K and apply the martingale convergence theorem (Theorem 11.4) to the stopped process X^{ϱ_K} .

Exercise 11.2.7 (Conditional Borel-Cantelli lemma) Let $(\mathcal{F}_n)_{n \in \mathbb{N}_0}$ be a filtration and let $(A_n)_{n \in \mathbb{N}}$ be events with $A_n \in \mathcal{F}_n$ for all $n \in \mathbb{N}$. Define $A_\infty = \{\sum_{n=1}^\infty \mathbf{P}[A_n | \mathcal{F}_{n-1}] = \infty\}$ and $A^* = \limsup_{n \rightarrow \infty} A_n$. Show the conditional Borel-Cantelli lemma: $\mathbf{P}[A_\infty \triangle A^*] = 0$.

Hint: Apply Exercise 11.2.5 to $X_n = \sum_{i=1}^n (\mathbb{1}_{A_i} - \mathbf{P}[A_i | \mathcal{F}_{i-1}])$.

Exercise 11.2.8 Let $p \in [0, 1]$ and let $X = (X_n)_{n \in \mathbb{N}_0}$ be a stochastic process with values in $[0, 1]$. Assume that for all $n \in \mathbb{N}_0$, given $X_0, \dots, X_n$, we have

$$X_{n+1} = \begin{cases} 1 - p + pX_n & \text{with probability } X_n \\ pX_n & \text{with probability } 1 - X_n \end{cases}$$

Show that X is a martingale that converges almost surely. Compute the distribution of the almost sure limit $\lim_{n \rightarrow \infty} X_n$.

Exercise 11.2.9 Let $f \in \mathcal{L}^1(\lambda)$, where λ is the restriction of the Lebesgue measure to $[0, 1]$. Let $I_{n,k} = [k2^{-n}, (k+1)2^{-n})$ for $n \in \mathbb{N}$ and $k = 0, \dots, 2^n - 1$. Define $f_n : [0, 1] \rightarrow \mathbb{R}$ by

$$f_n(x) = 2^n \int_{I_{k,n}} f d\lambda, \quad \text{if } k \text{ is chosen such that } x \in I_{k,n}$$

Show that $f_n(x) \xrightarrow{n \rightarrow \infty} f(x)$ for λ -almost all $x \in [0, 1]$.

Exercise 11.2.10 Assume that $\mathbb{F} = (\mathcal{F}_n)_{n \in \mathbb{N}}$ is a filtration on the probability space $(\Omega, \mathcal{A}, \mathbf{P})$. Let

$\mathcal{F}_\infty := \sigma(\mathcal{F}_n : n \in \mathbb{N})$, and let $\mathcal{M}$ be the vector space of uniformly integrable $\mathbb{F}$ -martingales. Show that the map $\Phi : \mathcal{L}^1(\mathcal{F}_\infty) \rightarrow \mathcal{M}, X_\infty \mapsto (\mathbf{E}[X_\infty | \mathcal{F}_n])_{n \in \mathbb{N}}$ is an isomorphism of vector spaces.

5.11.3 11.3 Example: Branching Process

Let $p = (p_k)_{k \in \mathbb{N}_0}$ be a probability vector on $\mathbb{N}_0$ and let $(Z_n)_{n \in \mathbb{N}_0}$ be the Galton-Watson process with one ancestor and offspring distribution p (see Definition 3.9). For convenience, we recall the construction of Z . Let $(X_{n,i})_{n \in \mathbb{N}_0, i \in \mathbb{N}}$ be i.i.d. random variables with $\mathbf{P}[X_{1,1} = k] = p_k$ for $k \in \mathbb{N}_0$. Let $Z_0 = 1$ and inductively define

$$Z_{n+1} = \sum_{i=1}^{Z_n} X_{n,i} \quad \text{for } n \in \mathbb{N}_0$$

We interpret Z_n as the size of a population at time n and $X_{n,i}$ as the number of offspring of the i th individual of the n th generation.

Let $m := \mathbf{E}[X_{1,1}] < \infty$ be the expected number of offspring of an individual and let $\sigma^2 := \text{Var}[X_{1,1}] \in (0, \infty)$ be its variance. Let $\mathcal{F}_n := \sigma(X_{k,i} : k < n, i \in \mathbb{N})$. Then Z is adapted to $\mathbb{F}$. Define $W_n = m^{-n} Z_n$.

Lemma 11.18 W is a martingale. In particular, $\mathbf{E}[Z_n] = m^n$ for all $n \in \mathbb{N}$.

Proof We compute the conditional expectation for $n \in \mathbb{N}_0$:

$$\begin{aligned} \mathbf{E}[W_{n+1} | \mathcal{F}_n] &= m^{-(n+1)} \mathbf{E}[Z_{n+1} | \mathcal{F}_n] \\ &= m^{-(n+1)} \mathbf{E}\left[\sum_{i=1}^{Z_n} X_{n,i} \mid \mathcal{F}_n\right] \\ &= m^{-(n+1)} \sum_{k=1}^{\infty} \mathbf{E}[\mathbb{1}_{\{Z_n=k\}} k \cdot X_{n,i} \mid \mathcal{F}_n] \\ &= m^{-n} \sum_{k=1}^{\infty} \mathbf{E}[k \cdot \mathbb{1}_{\{Z_n=k\}} \mid \mathcal{F}_n] \\ &= m^{-n} Z_n = W_n \end{aligned}$$

Theorem 11.19 Let $\text{Var}[X_{1,1}] \in (0, \infty)$. The a.s. limit $W_\infty = \lim_{n \rightarrow \infty} W_n$ exists and

$$m > 1 \iff \mathbf{E}[W_\infty] = 1 \iff \mathbf{E}[W_\infty] > 0.$$

Proof W_∞ exists since $W \geq 0$ is a martingale. If $m \leq 1$, then $(Z_n)_{n \in \mathbb{N}}$ converges a.s. to some random variable Z_∞ . Note that Z_∞ is the only choice since $\sigma^2 > 0$.

Now let $m > 1$. Since $\mathbf{E}[Z_{n-1}] = m^{n-1}$ (Lemma 11.18), by the Blackwell-Girshick formula (Theorem 5.10),

$$\begin{aligned} \text{Var}[W_n] &= m^{-2n} (\sigma^2 \mathbf{E}[Z_{n-1}] + m^2 \text{Var}[Z_{n-1}]) \\ &= \sigma^2 m^{-(n+1)} + \text{Var}[W_{n-1}] \end{aligned}$$

Inductively, we get $\text{Var}[W_n] = \sigma^2 \sum_{k=2}^{n+1} m^{-k} \leq \frac{\sigma^2 m}{m-1} < \infty$. Hence W is bounded in L^2 and Theorem 11.10 yields $W_n \rightarrow W_\infty$ in L^2 and thus in L^1 . In particular, $\mathbf{E}[W_\infty] = \mathbf{E}[W_0] = 1$.

The proof of Theorem 11.19 was simple due to the assumption of finite variance of the offspring distribution. However, there is a much stronger statement that here we can only quote (see [96], and see [110] for a modern proof).

Theorem 11.20 (Kesten-Stigum (1966)) Let $m > 1$. Then

$$\mathbf{E}[W_\infty] = 1 \iff \mathbf{E}[W_\infty] > 0 \iff \mathbf{E}[X_{1,1} \log(X_{1,1})^+] < \infty.$$

5.12 12 Backwards Martingales and Exchangeability

With many data acquisitions, such as telephone surveys, the order in which the data come does not matter. Mathematically, we say that a family of random variables is exchangeable if the joint distribution does not change under finite permutations. De Finetti's structural theorem says that an infinite family of E -valued exchangeable random variables can be described by a two-stage experiment. At the first stage, a probability distribution Ξ on E is drawn at random. At the second stage, i.i.d. random variables with distribution Ξ are implemented.

We first define the notion of exchangeability. Then we consider backwards martingales and prove the convergence theorem for them. This is the cornerstone for the proof of de Finetti's theorem.

5.12.1 12.1 Exchangeable Families of Random Variables

Definition 12.1 Let I be an arbitrary index set and let E be a Polish space. A family $(X_i)_{i \in I}$ of random variables with values in E is called exchangeable if

$$\mathcal{L}[(X_{\varrho(i)})_{i \in I}] = \mathcal{L}[(X_i)_{i \in I}]$$

for any finite permutation $\varrho : I \rightarrow I$.

Recall that a finite permutation is a bijection $\varrho : I \rightarrow I$ that leaves all but finitely many points unchanged.

Remark 12.2 Clearly, the following are equivalent.

- (i) $(X_i)_{i \in I}$ is exchangeable.
- (ii) Let $n \in \mathbb{N}$ and assume $i_1, \dots, i_n \in I$ are pairwise distinct and $j_1, \dots, j_n \in I$ are pairwise distinct. Then we have $\mathcal{L}[(X_{i_1}, \dots, X_{i_n})] = \mathcal{L}[(X_{j_1}, \dots, X_{j_n})]$.

In particular ($n = 1$), exchangeable random variables are identically distributed.

Example 12.3

- (i) If $(X_i)_{i \in I}$ is i.i.d., then $(X_i)_{i \in I}$ is exchangeable.
- (ii) Consider an urn with N balls, M of which are black. Successively draw without replacement all of the balls and define

$$X_n := \begin{cases} 1, & \text{if the } n\text{th ball is black} \\ 0, & \text{else} \end{cases}$$

Then $(X_n)_{n=1, \dots, N}$ is exchangeable. Indeed, this follows by elementary combinatorics since for any choice $x_1, \dots, x_N \in \{0, 1\}$ with $x_1 + \dots + x_N = M$, we have

$$\mathbf{P}[X_1 = x_1, \dots, X_N = x_N] = \frac{1}{\binom{N}{M}}$$

This formula can be derived formally via a small computation with conditional probabilities. As we will need a similar computation for Pólya's urn model in Example 12.29, we give the details here. Let $s_k = x_1 + \dots + x_k$ for $k = 0, \dots, N$ and let

$$g_k(x) = \begin{cases} M - s_k, & \text{if } x = 1, \\ N - M + s_k - k, & \text{if } x = 0 \end{cases}$$

Then $\mathbf{P}[X_1 = x_1] = g_0(x_1)/N$ and
 $\mathbf{P}[X_{k+1} = x_{k+1} \mid X_1 = x_1, \dots, X_k = x_k] = \frac{g_k(x_{k+1})}{N-k}$ for $k = 1, \dots, N-1$.
 Clearly, $g_k(0) = N - M - l$, where $l = \#\{i \leq k : x_i = 0\}$. Therefore,

$$\begin{aligned} \mathbf{P}[X_1 = x_1, \dots, X_N = x_N] &= \mathbf{P}[X_1 = x_1] \prod_{k=1}^{N-1} \mathbf{P}[X_{k+1} = x_{k+1} \mid X_1 = x_1, \dots, X_k = x_k] \\ &= \frac{1}{N!} \prod_{k=0}^{N-1} g_k(x_{k+1}) = \frac{1}{N!} \prod_{k:x_k=1} g_k(1) \prod_{k:x_k=0} g_k(0) \\ &= \frac{1}{N!} \prod_{l=0}^{M-1} (M-l) \prod_{l=0}^{N-1} (N-M-l) = \frac{M!(N-M)!}{N!}. \end{aligned}$$

(iii) Let Y be a random variable with values in $[0, 1]$. Assume that, given Y , the random variables $(X_i)_{i \in I}$ are independent and Ber_Y -distributed. That is, for any finite $J \subset I$,

$$\mathbf{P}[X_j = 1 \text{ for all } j \in J \mid Y] = Y^{\#J}.$$

Then $(X_i)_{i \in I}$ is exchangeable.

Let $X = (X_n)_{n \in \mathbb{N}}$ be a stochastic process with values in a Polish space E . Let $S(n)$ be the set of permutations $\varrho : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$. We consider ϱ also as a map $\mathbb{N} \rightarrow \mathbb{N}$ by defining $\varrho(k) = k$ for $k > n$. For $\varrho \in S(n)$ and $x = (x_1, \dots, x_n) \in E^n$, denote $x^\varrho = (x_{\varrho(1)}, \dots, x_{\varrho(n)})$. Similarly, for $x \in E^\mathbb{N}$, denote $x^\varrho = (x_{\varrho(1)}, x_{\varrho(2)}, \dots) \in E^\mathbb{N}$. Let E' be another Polish space. For measurable maps $f : E^n \rightarrow E'$ and $F : E^\mathbb{N} \rightarrow E'$, define the maps f^ϱ and F^ϱ by $f^\varrho(x) = f(x^\varrho)$ and $F^\varrho(x) = F(x^\varrho)$. Further, we write $f(x) = f(x_1, \dots, x_n)$ for $x \in E^n$ and for $x \in E^\mathbb{N}$.

Definition 12.4 (i) A map $f : E^n \rightarrow E'$ is called symmetric if $f^\varrho = f$ for all $\varrho \in S(n)$.

(ii) A map $F : E^\mathbb{N} \rightarrow E'$ is called n -symmetric if $F^\varrho = F$ for all $\varrho \in S(n)$. F is called symmetric if F is n -symmetric for all $n \in \mathbb{N}$.

Example 12.5

(i) For $x \in \mathbb{R}^\mathbb{N}$, define the n th arithmetic mean by $a_n(x) = \frac{1}{n} \sum_{i=1}^n x_i$. Clearly, a_n is an n -symmetric map (but not m -symmetric for any $m > n$). Furthermore, $\bar{a}(x) := \limsup_{n \rightarrow \infty} a_n(x)$ defines a symmetric map $\mathbb{R}^\mathbb{N} \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$.

(ii) The map $s : \mathbb{R}^\mathbb{N} \rightarrow [0, \infty]$, $x \mapsto \sum_{i=1}^\infty |x_i|$ is symmetric. Unlike $\bar{a}$, the value of s depends on every coordinate if it is finite.

(iii) For $x \in E^\mathbb{N}$, define the n th empirical distribution by $\xi_n(x) = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ (recall that δ_{x_i} is the Dirac measure at the point x_i). Clearly, ξ_n is an n -symmetric map.

(iv) Let $k \in \mathbb{N}$ and let $\varphi : E^k \rightarrow \mathbb{R}$ be a map. The n th symmetrized average

$$A_n(\varphi) : E^\mathbb{N} \rightarrow \mathbb{R}, \quad x \mapsto \frac{1}{n!} \sum_{\varrho \in S(n)} \varphi(x^\varrho) \quad (12.1)$$

is an n -symmetric map.

Definition 12.6 Let $X = (X_n)_{n \in \mathbb{N}}$ be a stochastic process with values in E . For $n \in \mathbb{N}$, define

$$\mathcal{E}'_n := \sigma(F : F : E^\mathbb{N} \rightarrow \mathbb{R} \text{ is measurable and } n\text{-symmetric})$$

and let $\mathcal{E}_n := X^{-1}(\mathcal{E}'_n)$ be the σ -algebra of events that are invariant under all permutations $\varrho \in S(n)$. Further, let

$$\mathcal{E}' := \bigcap_{n=1}^{\infty} \mathcal{E}'_n = \sigma(F : F : E^{\mathbb{N}} \rightarrow \mathbb{R} \text{ is measurable and symmetric})$$

and let $\mathcal{E}_n := \bigcap_{m=1}^{\infty} \mathcal{E}'_m = X^{-1}(\mathcal{E}')$ be the σ -algebra of exchangeable events for X , or briefly the exchangeable σ -algebra.

Remark 12.7 If $A \in \sigma(X_n, n \in \mathbb{N})$ is an event, then there is a measurable $B \subset E^{\mathbb{N}}$ with $A = \{X \in B\}$. If we denote $A^\varrho = \{X^\varrho \in B\}$ for $\varrho \in S(n)$, then $\mathcal{E}_n = \{A : A^\varrho = A \text{ for all } \varrho \in S(n)\}$. This justifies the name "exchangeable event".

Remark 12.8 If we write $\Xi_n(\omega) := \xi_n(X(\omega)) = \frac{1}{n} \sum_{i=1}^n \delta_{X_i(\omega)}$ for the n thempirical distribution, then, by Exercise 12.1.1, $\mathcal{E}_n = \sigma(\Xi_n)$.

Remark 12.9 Denote by $\mathcal{T} = \bigcap_{n \in \mathbb{N}} \sigma(X_{n+1}, X_{n+2}, \dots)$ the tail σ -algebra. Then $\mathcal{T} \subset \mathcal{E}$, and strict inclusion is possible.

Indeed, evidently $\sigma(X_{n+1}, X_{n+2}, \dots) \subset \mathcal{E}_n$ for $n \in \mathbb{N}$; hence $\mathcal{T} \subset \mathcal{E}$. Now let $E = \{0, 1\}$ and let $X_1, X_2, \dots$ be independent random variables with $\mathbf{P}[X_n = 1] \in (0, 1)$ for all $n \in \mathbb{N}$. The random variable $S := \sum_{n=1}^{\infty} X_n$ is measurable with respect to $\mathcal{E}$ but not with respect to $\mathcal{T}$.

Theorem 12.10 Let $X = (X_n)_{n \in \mathbb{N}}$ be exchangeable. If $\varphi : E^{\mathbb{N}} \rightarrow \mathbb{R}$ is measurable and if $\mathbf{E}[|\varphi(X)|] < \infty$, then for all $n \in \mathbb{N}$ and all $\varrho \in S(n)$,

$$\mathbf{E}[\varphi(X) \mid \mathcal{E}_n] = \mathbf{E}[\varphi(X^\varrho) \mid \mathcal{E}_n]. \quad (12.2)$$

In particular,

$$\mathbf{E}[\varphi(X) \mid \mathcal{E}_n] = A_n(\varphi) := \frac{1}{n!} \sum_{\varrho \in S(n)} \varphi(X^\varrho). \quad (12.3)$$

Proof Let $A \in \mathcal{E}_n$. Then there exists a $B \in \mathcal{E}'_n$ such that $A = X^{-1}(B)$. Let $F = \mathbb{1}_B$. Then $F \circ X = \mathbb{1}_A$. By the definition of $\mathcal{E}_n$, the map $F : E^{\mathbb{N}} \rightarrow \mathbb{R}$ is measurable, n -symmetric and bounded. Therefore,

$$\mathbf{E}[\varphi(X)F(X)] = \mathbf{E}[\varphi(X^\varrho)F(X^\varrho)] = \mathbf{E}[\varphi(X^\varrho)F(X)].$$

Here we used the exchangeability of X in the first equality and the symmetry of F in the second equality. From this (12.2) follows. However, $A_n(\varphi)$ is $\mathcal{E}_n$ -measurable and hence

$$\mathbf{E}[\varphi(X) \mid \mathcal{E}_n] = \mathbf{E}\left[\frac{1}{n!} \sum_{\varrho \in S(n)} \varphi(X^\varrho) \mid \mathcal{E}_n\right] = \frac{1}{n!} \sum_{\varrho \in S(n)} \varphi(X^\varrho).$$

Heuristic for the Structure of Exchangeable Families Consider a finite exchangeable family $X_1, \dots, X_N$ of E -valued random variables. For $n \leq N$, what is the conditional distribution of $(X_1, \dots, X_n)$ given Ξ_N ? For any measurable $A \subset E$, $\{X_i \in A\}$ occurs for exactly $N\Xi_N(A)$ of the $i \in \{1, \dots, N\}$, where the order does not change the probability. Hence we are in the situation of

drawing colored balls without replacement. More precisely, let the pairwise distinct points $e_1, \dots, e_k \in E$ be the atoms of Ξ_N and let $N_1, \dots, N_k$ be the corresponding absolute frequencies. Hence $\Xi_N = \sum_{i=1}^k (N_i/N) \delta_{e_i}$. We thus deal with balls of k different colors and with N_i balls of the i th color. We draw n of these balls without replacement but respecting the order. Up to the order, the resulting distribution is thus the generalized hypergeometric distribution (see (1.19) on page 44). Hence, for pairwise disjoint, measurable sets $A_1, \dots, A_k$ with $\bigsqcup_{l=1}^k A_l = E$, for $i_1, \dots, i_n \in \{1, \dots, k\}$, pairwise distinct $j_1, \dots, j_n \in \{1, \dots, N\}$ and with the convention $m_l := \#\{r \in \{1, \dots, n\} : i_r = l\}$ for $l \in \{1, \dots, k\}$, we have

$$\mathbf{P}[X_{j_r} \in A_{i_r} \text{ for all } r = 1, \dots, n \mid \Xi_N] = \frac{1}{(N)_n} \prod_{l=1}^k (N \Xi_N(A_l))^{m_l}. \quad (12.4)$$

Here we defined $(n)_l := n(n-1)\dots(n-l+1)$.

What happens if we let $N \rightarrow \infty$? For simplicity, assume that for all $l = 1, \dots, k$, the limit $\Xi_\infty(A_l) = \lim_{N \rightarrow \infty} \Xi_N(A_l)$ exists in a suitable sense. Then (12.4) formally becomes

$$\mathbf{P}[X_{j_r} \in A_{i_r} \text{ for all } r = 1, \dots, n \mid \Xi_\infty] = \prod_{l=1}^k \Xi_\infty(A_l)^{m_l} \quad (12.5)$$

Drawing without replacements thus asymptotically turns into drawing with replacements. Hence the random variables $X_1, X_2, \dots$ are independent with distribution Ξ_∞ given Ξ_∞ .

For a formal proof along the lines of this heuristic, see Section 13.4.

In order to formulate (and prove) this statement (de Finetti's theorem) rigorously in Section 12.3, we need some more technical tools (e.g., the notion of conditional independence). A further tool will be the convergence theorem for backwards martingales that will be formulated in Section 12.2. For further reading on exchangeable random variables, we refer to [4, 33, 98, 105].

Exercise 12.1.1 Let $n \in \mathbb{N}$. Show that every symmetric function $f : E^n \rightarrow \mathbb{R}$ can be written in the form $f(x) = g(\frac{1}{n} \sum_{i=1}^n \delta_{x_i})$, where g has to be chosen appropriately (depending on f).

Exercise 12.1.2 Derive equation (12.4) formally.

Exercise 12.1.3 Let $X_1, \dots, X_n$ be exchangeable, square integrable random variables. Show that

$$\text{Cov}[X_1, X_2] \geq -\frac{1}{n-1} \text{Var}[X_1] \quad (12.6)$$

For $n \geq 2$, give a nontrivial example for equality in (12.6).

Exercise 12.1.4 Let $X_1, X_2, X_3, \dots$ be exchangeable, square integrable random variables. Show that $\text{Cov}[X_1, X_2] \geq 0$.

Exercise 12.1.5 Show that for all $n \in \mathbb{N} \setminus \{1\}$, there is an exchangeable family of random variables $X_1, \dots, X_n$ that cannot be extended to an infinite exchangeable family $X_1, X_2, \dots$.

5.12.2 12.2 Backwards Martingales

The concepts of filtration and martingale do not require the index set I (interpreted as time) to be a subset of $[0, \infty)$. Hence we can consider the case $I = -\mathbb{N}_0$.

Definition 12.11 (Backwards martingale) Let $\mathbb{F} = (\mathcal{F}_n)_{n \in -\mathbb{N}_0}$ be a filtration and let $X = (X_n)_{n \in -\mathbb{N}_0}$ be an $\mathbb{F}$ -martingale. Then $X = (X_{-n})_{n \in \mathbb{N}_0}$ is called a backwards martingale.

Remark 12.12 A backwards martingale is always uniformly integrable. This follows from Corollary 8.22 and the fact that $X_{-n} = \mathbf{E}[X_0 \mid \mathcal{F}_{-n}]$ for any $n \in \mathbb{N}_0$.

Example 12.13 Let $X_1, X_2, \dots$ be exchangeable real random variables. For $n \in \mathbb{N}$, let $\mathcal{F}_{-n} = \mathcal{E}_n$ and

$$Y_{-n} = \frac{1}{n} \sum_{i=1}^n X_i.$$

We show that $(Y_{-n})_{n \in \mathbb{N}}$ is an $\mathbb{F}$ -backwards martingale. Clearly, Y is adapted. Furthermore, by Theorem 12.10 (with $k = n$ and $\varphi(X_1, \dots, X_n) = \frac{1}{n-1} (X_1 + \dots + X_{n-1})$),

$$\mathbf{E}[Y_{-n+1} \mid \mathcal{F}_{-n}] = \frac{1}{n!} \sum_{\varrho \in S(n)} \frac{1}{n-1} (X_{\varrho(1)} + \dots + X_{\varrho(n-1)}) = Y_{-n}$$

Now replace $\mathbb{F}$ by the smaller filtration $\mathbb{G} = (\mathcal{G}_n)_{n \in -\mathbb{N}}$ that is defined by $\mathcal{G}_{-n} = \sigma(Y_{-n}, X_{n+1}, X_{n+2}, \dots) = \sigma(Y_{-n}, Y_{-n-1}, Y_{-n-2}, \dots)$ for $n \in \mathbb{N}$. This is the filtration generated by Y ; thus Y is also a $\mathbb{G}$ -backwards martingale (see Remark 9.29). $\diamond$

Let $a < b$ and $n \in \mathbb{N}$. Let $U_{-n}^{a,b}$ be the number of upcrossings of X over $[a, b]$ between times $-n$ and 0 . Further, let $U^{a,b} = \lim_{n \rightarrow \infty} U_{-n}^{a,b}$. By the upcrossing inequality (Lemma 11.3), we have $\mathbf{E}[U_{-n}^{a,b}] \leq \frac{1}{b-a} \mathbf{E}[(X_0 - a)^+]$; hence $\mathbf{P}[U^{a,b} < \infty] = 1$. As in the proof of the martingale convergence theorem (Theorem 11.4), we infer the following.

Theorem 12.14 (Convergence theorem for backwards martingales) Let $(X_n)_{n \in -\mathbb{N}_0}$ be a martingale with respect to $\mathbb{F} = (\mathcal{F}_n)_{n \in -\mathbb{N}_0}$. Then there exists $X_{-\infty} = \lim_{n \rightarrow \infty} X_{-n}$ almost surely and in L^1 . Furthermore, $X_{-\infty} = \mathbf{E}[X_0 | \mathcal{F}_{-\infty}]$, where $\mathcal{F}_{-\infty} = \bigcap_{n=1}^{\infty} \mathcal{F}_{-n}$.

Example 12.15 Let $X_1, X_2, \dots$ be exchangeable, integrable random variables. Further, let $\mathcal{T} = \bigcap_{n=1}^{\infty} \sigma(X_m, m \geq n)$ be the tail σ -algebra of $X_1, X_2, \dots$ and let $\mathcal{E}$ be the exchangeable σ -algebra. Then $\mathbf{E}[X_1 | \mathcal{T}] = \mathbf{E}[X_1 | \mathcal{E}]$ a.s. and

$$\frac{1}{n} \sum_{i=1}^n X_i \xrightarrow{n \rightarrow \infty} \mathbf{E}[X_1 | \mathcal{E}] \quad \text{a.s. and in } L^1$$

Indeed, if we let $Y_{-n} := \frac{1}{n} \sum_{i=1}^n X_i$, then (by Example 12.13) $(Y_{-n})_{n \in \mathbb{N}}$ is a backwards martingale with respect to $(\mathcal{F}_n)_{n \in -\mathbb{N}} = (\mathcal{E}_{-n})_{n \in -\mathbb{N}}$ and thus

$$Y_{-n} \xrightarrow{n \rightarrow \infty} Y_{-\infty} = \mathbf{E}[X_1 | \mathcal{E}] \quad \text{a.s. and in } L^1.$$

Now, by Example 2.36(ii), $Y_{-\infty}$ is $\mathcal{T}$ -measurable; hence (since $\mathcal{T} \subset \mathcal{E}$ and by virtue of the tower property of conditional expectation) $Y_{-\infty} = \mathbf{E}[X_1 | \mathcal{T}]$.

Example 12.16 (Strong law of large numbers) If $Z_1, Z_2, \dots$ are real and i.i.d. with $\mathbf{E}[|Z_1|] < \infty$, then

$$\frac{1}{n} \sum_{i=1}^n Z_i \xrightarrow{n \rightarrow \infty} \mathbf{E}[Z_1] \quad \text{almost surely.}$$

By Kolmogorov's 0-1 law (Theorem 2.37), the tail σ -algebra $\mathcal{T}$ is trivial; hence we have

$$\mathbf{E}[Z_1 | \mathcal{T}] = \mathbf{E}[Z_1] \quad \text{almost surely.}$$

In Corollary 12.19, we will see that in the case of independent random variables, $\mathcal{E}$ is also $\mathbf{P}$ -trivial. This implies $\mathbf{E}[Z_1 | \mathcal{E}] = \mathbf{E}[Z_1]$.

We close this section with a generalization of Example 12.15 to mean values of functions of $k \in \mathbb{N}$ variables. This conclusion from the convergence theorem for backwards martingales will be used in an essential way in the next section.

Theorem 12.17 Let $X = (X_n)_{n \in \mathbb{N}}$ be an exchangeable family of random variables with values in E . Assume that $k \in \mathbb{N}$ and let $\varphi : E^k \rightarrow \mathbb{R}$ be measurable with $\mathbf{E}[|\varphi(X_1, \dots, X_k)|] < \infty$. Denote $\varphi(X) = \varphi(X_1, \dots, X_k)$ and let $A_n(\varphi) := \frac{1}{n!} \sum_{\varrho \in S(n)} \varphi(X^{\varrho})$. Then

$$\mathbf{E}[\varphi(X) | \mathcal{E}] = \mathbf{E}[\varphi(X) | \mathcal{T}] = \lim_{n \rightarrow \infty} A_n(\varphi) \quad \text{a.s. and in } L^1. \quad (12.7)$$

Proof By Theorem 12.10, $A_n(\varphi) = \mathbf{E}[\varphi(X) | \mathcal{E}_n]$. Hence $(A_{-n}(\varphi))_{n \geq k}$ is a backwards martingale with respect to $(\mathcal{E}_{-n})_{n \in -\mathbb{N}}$. Hence, by Theorem 12.14,

$$A_n(\varphi) \xrightarrow{n \rightarrow \infty} \mathbf{E}[\varphi(X) | \mathcal{E}] \quad \text{a.s. and in } L^1. \quad (12.8)$$

As for the arithmetic mean (Example 12.16), we can argue that $\lim_{n \rightarrow \infty} A_n(\varphi)$ is $\mathcal{T}$ -measurable. Indeed,

$$\limsup_{n \rightarrow \infty} \frac{\#\{\varrho \in S(n) : \varrho^{-1}(i) \leq l \text{ for some } i \in \{1, \dots, k\}\}}{n!} = 0 \quad \text{for all } l \in \mathbb{N}.$$

Thus, for large n , the dependence of $A_n(\varphi)$ on the first l coordinates is negligible. Together with (12.8), we get (12.7).

Corollary 12.18 Let $X = (X_n)_{n \in \mathbb{N}}$ be exchangeable. Then, for any $A \in \mathcal{E}$ there exists a $B \in \mathcal{T}$ with $\mathbf{P}[A \Delta B] = 0$.

Note that $\mathcal{T} \subset \mathcal{E}$; hence the statement is trivially true if the roles of $\mathcal{E}$ and $\mathcal{T}$ are interchanged.

Proof Since $\mathcal{E} \subset \sigma(X_1, X_2, \dots)$, by the approximation theorem for measures, there exists a sequence of measurable sets $(A_k)_{k \in \mathbb{N}}$ with $A_k \in \sigma(X_1, \dots, X_k)$ and such that $\mathbf{P}[A \Delta A_k] \xrightarrow{k \rightarrow \infty} 0$. Let $C_k \in \mathcal{E}^k$ be measurable with

$$A_k = \{(X_1, \dots, X_k) \in C_k\}$$

for all $k \in \mathbb{N}$. Letting $\varphi_k := \mathbb{1}_{C_k}$, Theorem 12.17 implies that

$$\begin{aligned} \mathbb{1}_A &= \mathbf{E}[\mathbb{1}_A \mid \mathcal{E}] = \mathbf{E}\left[\lim_{k \rightarrow \infty} \varphi_k(X) \mid \mathcal{E}\right] = \lim_{k \rightarrow \infty} \mathbf{E}[\varphi_k(X) \mid \mathcal{E}] \\ &= \lim_{k \rightarrow \infty} \mathbf{E}[\varphi_k(X) \mid \mathcal{T}] =: \psi \quad \text{almost surely} \end{aligned}$$

Hence there is a $\mathcal{T}$ -measurable function ψ with $\psi = \mathbb{1}_A$ almost surely. We can assume that $\psi = \mathbb{1}_B$ for some $B \in \mathcal{T}$.

As a further application, we get the 0-1 law of Hewitt and Savage [72].

Corollary 12.19 (0 – 1 law of Hewitt-Savage) Let $X_1, X_2, \dots$ be i.i.d. random variables. Then the exchangeable σ -algebra is $\mathbf{P}$ -trivial; that is, $\mathbf{P}[A] \in \{0, 1\}$ for all $A \in \mathcal{E}$.

Proof By Kolmogorov's 0-1 law (Theorem 2.37), $\mathcal{T}$ is trivial. Hence the claim follows immediately from Corollary 12.18.

5.12.3 12.3 De Finetti's Theorem

In this section, we show the structural theorem for countably infinite exchangeable families that was heuristically motivated at the end of Section 12.1. Hence we shall show that a countably infinite exchangeable family of random variables is an i.i.d.

family given the exchangeable σ -algebra $\mathcal{E}$. Furthermore, we compute the conditional distribution of the individual random variables. As a first step, we define conditional independence formally (see [25, Chapter 7.3]).

Definition 12.20 (Conditional independence) Let $(\Omega, \mathcal{F}, \mathbf{P})$ be a probability space, let $\mathcal{A} \subset \mathcal{F}$ be a sub- σ -algebra and let $(\mathcal{A}_i)_{i \in I}$ be an arbitrary family of sub- σ -algebras of $\mathcal{F}$. Assume that for any finite $J \subset I$, any choice of $A_j \in \mathcal{A}_j$ and for all $j \in J$,

$$\mathbf{P}\left[\bigcap_{j \in J} A_j \mid \mathcal{A}\right] = \prod_{j \in J} \mathbf{P}[A_j \mid \mathcal{A}] \quad \text{almost surely.} \quad (12.9)$$

Then the family $(\mathcal{A}_i)_{i \in I}$ is called independent given $\mathcal{A}$.

A family $(X_i)_{i \in I}$ of random variables on $(\Omega, \mathcal{F}, \mathbf{P})$ is called independent (and identically distributed) given $\mathcal{A}$ if the generated σ -algebras $(\sigma(X_i))_{i \in I}$ are independent given $\mathcal{A}$ (and the conditional distributions $\mathbf{P}[X_i \in \cdot \mid \mathcal{A}]$ are equal).

Example 12.21 Any family $(\mathcal{A}_i)_{i \in I}$ of sub- σ -algebras of $\mathcal{F}$ is independent given $\mathcal{F}$. Indeed, letting $A = \bigcap_{j \in J} A_j$,

$$\mathbf{P}[A \mid \mathcal{F}] = \mathbb{1}_A = \prod_{j \in J} \mathbb{1}_{A_j} = \prod_{j \in J} \mathbf{P}[A_j \mid \mathcal{F}] \quad \text{almost surely.}$$

□

Example 12.22 If $(\mathcal{A}_i)_{i \in I}$ is an independent family of σ -algebras and if $\mathcal{A}$ is trivial, then $(\mathcal{A}_i)_{i \in I}$ is independent given $\mathcal{A}$. □

Example 12.23 There is no "monotonicity" for conditional independence in the following sense: If $\mathcal{F}_1, \mathcal{F}_2$ and $\mathcal{F}_3$ are σ -algebras with $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \mathcal{F}_3$ and such that $(\mathcal{A}_i)_{i \in I}$ is independent given $\mathcal{F}_1$ as well as given $\mathcal{F}_3$, then this does not imply independence given $\mathcal{F}_2$.

In order to illustrate this, assume that X and Y are nontrivial independent real random variables. Let $\mathcal{F}_1 = \{\emptyset, \Omega\}$, $\mathcal{F}_2 = \sigma(X + Y)$ and $\mathcal{F}_3 = \sigma(X, Y)$. Then $\sigma(X)$ and $\sigma(Y)$ are independent given $\mathcal{F}_1$ as well as given $\mathcal{F}_3$ but not given $\mathcal{F}_2$. □

Let $X = (X_n)_{n \in \mathbb{N}}$ be a stochastic process on a probability space $(\Omega, \mathcal{F}, \mathbf{P})$ with values in a Polish space E . Let $\mathcal{E}$ be the exchangeable σ -algebra and let $\mathcal{T}$ be the tail σ -algebra.

Theorem 12.24 (de Finetti) The family $X = (X_n)_{n \in \mathbb{N}}$ is exchangeable if and only if there exists a σ -algebra $\mathcal{A} \subset \mathcal{F}$ such that $(X_n)_{n \in \mathbb{N}}$ is i.i.d. given $\mathcal{A}$. In this case, $\mathcal{A}$ can be chosen to equal the exchangeable σ -algebra $\mathcal{E}$ or the tail- σ -algebra $\mathcal{T}$.

Proof " $\implies$ " Let X be exchangeable and let $\mathcal{A} = \mathcal{E}$ or $\mathcal{A} = \mathcal{T}$. For any $n \in \mathbb{N}$, let $f_n : E \rightarrow \mathbb{R}$ be a bounded measurable map. Let

$$\varphi_k(x_1, \dots, x_k) = \prod_{i=1}^k f_i(x_i) \quad \text{for any } k \in \mathbb{N}.$$

Let $A_n(\varphi)$ be the symmetrized average from Theorem 12.17. Then

$$\begin{aligned} A_n(\varphi_{k-1}) A_n(f_k) &= \frac{1}{n!} \sum_{\varrho \in S(n)} \varphi_{k-1}(X^{\varrho}) \frac{1}{n} \sum_{i=1}^n f_k(X_i) \\ &= \frac{1}{n!} \sum_{\varrho \in S(n)} \varphi_k(X^{\varrho}) + R_{n,k} = A_n(\varphi_k) + R_{n,k} \end{aligned}$$

where

$$\begin{aligned} |R_{n,k}| &\leq 2 \|\varphi_{k-1}\|_{\infty} \cdot \|f_k\|_{\infty} \cdot \frac{1}{n!} \frac{1}{n} \sum_{\varrho \in S(n)} \sum_{i=1}^n \mathbb{1}_{\{i \in \{\varrho(1), \dots, \varrho(k-1)\}\}} \\ &= 2 \|\varphi_{k-1}\|_{\infty} \cdot \|f_k\|_{\infty} \cdot \frac{k-1}{n} \xrightarrow{n \rightarrow \infty} 0 \end{aligned}$$

Together with Theorem 12.17, we conclude that

$$A_n(\varphi_{k-1}) A_n(f_k) \xrightarrow{n \rightarrow \infty} \mathbf{E}[\varphi_k(X_1, \dots, X_k) \mid \mathcal{A}] \quad \text{a.s. and in } L^1.$$

On the other hand, again by Theorem 12.17,

$$A_n(\varphi_{k-1}) \xrightarrow{n \rightarrow \infty} \mathbf{E}[\varphi_{k-1}(X_1, \dots, X_{k-1}) \mid \mathcal{A}]$$

and

$$A_n(f_k) \xrightarrow{n \rightarrow \infty} \mathbf{E}[f_k(X_1) \mid \mathcal{A}].$$

Hence

$$\mathbf{E}[\varphi_k(X_1, \dots, X_k) \mid \mathcal{A}] = \mathbf{E}[\varphi_{k-1}(X_1, \dots, X_{k-1}) \mid \mathcal{A}] \mathbf{E}[f_k(X_1) \mid \mathcal{A}]$$

Thus we get inductively

$$\mathbf{E} \left[\prod_{i=1}^k f_i(X_i) \mid \mathcal{A} \right] = \prod_{i=1}^k \mathbf{E} [f_i(X_1) \mid \mathcal{A}]$$

Therefore, X is i.i.d. given $\mathcal{A}$.

" $\Leftarrow$ " Now let X be i.i.d. given $\mathcal{A}$ for a suitable σ -algebra $\mathcal{A} \subset \mathcal{F}$. For any bounded measurable function $\varphi : E^n \rightarrow \mathbb{R}$ and for any $\varrho \in S(n)$, we have $\mathbf{E}[\varphi(X) \mid \mathcal{A}] = \mathbf{E}[\varphi(X^\varrho) \mid \mathcal{A}]$. Hence

$$\mathbf{E}[\varphi(X)] = \mathbf{E}[\mathbf{E}[\varphi(X) \mid \mathcal{A}]] = \mathbf{E}[\mathbf{E}[\varphi(X^\varrho) \mid \mathcal{A}]] = \mathbf{E}[\varphi(X^\varrho)],$$

whence X is exchangeable.

Denote by $\mathcal{M}_1(E)$ the set of probability measures on E equipped with the topology of weak convergence (see Definition 13.12 and Remark 13.14). That is, a sequence $(\mu_n)_{n \in \mathbb{N}}$ in $\mathcal{M}_1(E)$ converges weakly to a $\mu \in \mathcal{M}_1(E)$ if and only if $\int f d\mu_n \xrightarrow{n \rightarrow \infty} \int f d\mu$ for any bounded continuous function $f : E \rightarrow \mathbb{R}$. We will study weak convergence in Chapter 13 in greater detail. At this point, we use the topology only to make $\mathcal{M}_1(E)$ a measurable space, namely with the Borel σ -algebra $\mathcal{B}(\mathcal{M}_1(E))$. Now we can study random variables with values in $\mathcal{M}_1(E)$, so-called random measures (compare also Section 24.1). For $x \in E^{\mathbb{N}}$, let $\xi_n(x) = \frac{1}{n} \sum_{i=1}^n \delta_{x_i} \in \mathcal{M}_1(E)$.

Definition 12.25 The random measure

$$\Xi_n := \xi_n(X) := \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$$

is called the empirical distribution of $X_1, \dots, X_n$.

Assume the conditions of Theorem 12.24 are in force.

Theorem 12.26 (de Finetti representation theorem) The family $X = (X_n)_{n \in \mathbb{N}}$ is exchangeable if and only if there is a σ -algebra $\mathcal{A} \subset \mathcal{F}$ and an $\mathcal{A}$ -measurable random variable $\Xi_\infty : \Omega \rightarrow \mathcal{M}_1(E)$ with the property that given Ξ_∞ , $(X_n)_{n \in \mathbb{N}}$ is i.i.d. with $\mathcal{L}[X_1 \mid \Xi_\infty] = \Xi_\infty$. In this case, we can choose $\mathcal{A} = \mathcal{E}$ or $\mathcal{A} = \mathcal{T}$.

Proof " $\Rightarrow$ " This follows as in the proof of Theorem 12.24.

" $\implies$ " Let X be exchangeable. Then, by Theorem 12.24, there exists a σ -algebra $\mathcal{A} \subset \mathcal{F}$ such that $(X_n)_{n \in \mathbb{N}}$ is i.i.d. given $\mathcal{A}$. As E is Polish, there exists a regular conditional distribution (see Theorem 8.37) $\Xi_\infty := \mathcal{L}[X_1 \mid \mathcal{A}]$. For measurable $A_1, \dots, A_n \subset E$, we have $\mathbf{P}[X_i \in A_i \mid \mathcal{A}] = \Xi_\infty(A_i)$ for all $i = 1, \dots, n$; hence

$$\begin{aligned} \mathbf{P} \left[\bigcap_{i=1}^n \{X_i \in A_i\} \mid \Xi_\infty \right] &= \mathbf{E} \left[\mathbf{P} \left[\bigcap_{i=1}^n \{X_i \in A_i\} \mid \mathcal{A} \right] \mid \Xi_\infty \right] \\ &= \mathbf{E} \left[\prod_{i=1}^n \Xi_\infty(A_i) \mid \Xi_\infty \right] = \prod_{i=1}^n \Xi_\infty(A_i) \end{aligned}$$

Therefore, $\mathcal{L}[X \mid \Xi_\infty] = \Xi_\infty^{\otimes \mathbb{N}}$.

Remark 12.27

(i) In the case considered in the previous theorem, by the strong law of large numbers, for any bounded continuous function $f : E \rightarrow \mathbb{R}$,

$$\int f d\Xi_n \xrightarrow{n \rightarrow \infty} \int f d\Xi_\infty \quad \text{almost surely.}$$

If in addition E is locally compact (e.g., $E = \mathbb{R}^d$), then one can even show that

$$\Xi_n \xrightarrow{n \rightarrow \infty} \Xi_\infty \quad \text{almost surely.}$$

(ii) For finite families of random variables there is no perfect analog of de Finetti's theorem. See [33] for a detailed treatment of finite exchangeable families.

Example 12.28 Let $(X_n)_{n \in \mathbb{N}}$ be exchangeable and assume $X_n \in \{0, 1\}$. Then there exists a random variable $Y : \Omega \rightarrow [0, 1]$ such that, for all finite $J \subset \mathbb{N}$,

$$\mathbf{P}[X_j = 1 \text{ for all } j \in J \mid Y] = Y^{\#J}.$$

In other words, $(X_n)_{n \in \mathbb{N}}$ is independent given Y and Ber_Y -distributed. Compare Example 12.3(iii).

Example 12.29 (Pólya's urn model) (See Example 14.38, compare also [17, 58, 135].) Consider an urn with a total of N balls among which M are black and $M - N$ are white. At each step, a ball is drawn and is returned to the urn together with an additional ball of the same color. Let

$$X_n := \begin{cases} 1, & \text{if the } n\text{th ball is black} \\ 0, & \text{else} \end{cases}$$

and let $S_n = \sum_{i=1}^n X_i$. Then

$$\mathbf{P}[X_n = 1 \mid X_1, X_2, \dots, X_{n-1}] = \frac{S_{n-1} + M}{N + n - 1}$$

Inductively, for $x_1, \dots, x_n \in \{0, 1\}$ and $s_k = \sum_{i=1}^k x_i$, we get

$$\begin{aligned} & \mathbf{P}[X_i = x_i \text{ for any } i = 1, \dots, n] \\ &= \prod_{i \leq n: x_i = 1} \frac{M + s_{i-1}}{N + i - 1} \prod_{i \leq n: x_i = 0} \frac{N + i - 1 - M - s_{i-1}}{N + i - 1} \\ &= \frac{(N-1)!}{(N-1+n)!} \cdot \frac{(M + s_n - 1)!}{(M-1)!} \frac{(N - M - 1 + (n - s_n))!}{(N - M - 1)!} \end{aligned}$$

The right-hand side depends on s_n only and not on the order of the $x_1, \dots, x_n$. Hence $(X_n)_{n \in \mathbb{N}}$ is exchangeable. Let

$$Z = \lim_{n \rightarrow \infty} \frac{1}{n} S_n$$

Then $(X_n)_{n \in \mathbb{N}}$ is i.i.d. Ber_Z -distributed given Z . Hence (see Example 12.28)

$$\begin{aligned} \mathbf{E}[Z^n] &= \mathbf{E}[\mathbf{P}[X_1 = \dots = X_n = 1 \mid Z]] \\ &= \mathbf{P}[S_n = n] \\ &= \frac{(N-1)!}{(M-1)!} \frac{(M+n-1)!}{(N+n-1)!} \quad \text{for all } n \in \mathbb{N} \end{aligned}$$

By Exercise 5.1.2, these are the moments of the Beta distribution $\beta_{M, N-M}$ on $[0, 1]$ with parameters $(M, N-M)$ (see Example 1.107(ii)). A distribution on $[0, 1]$ is uniquely characterized by its moments (see Theorem 15.4). Hence $Z \sim \beta_{M, N-M}$.

5.13 13 Convergence of Measures

One focus of probability theory is distributions that are the result of an interplay of a large number of random impacts. Often a useful approximation can be obtained by taking a limit of such distributions, for example, a limit where the number of impacts goes to infinity. With the Poisson distribution, we have encountered such a limit distribution that occurs as the number of very rare events when the number of possibilities goes to infinity (see Theorem 3.7). In many cases, it is necessary to rescale the original distributions in order to capture the behavior of the essential fluctuations, e.g., in the central limit theorem. While these theorems work with real random variables, we will also see limit theorems where the random variables take values in more general spaces such as the space of continuous functions when we model the path of the random motion of a particle.

In this chapter, we provide the abstract framework for the investigation of convergence of measures. We introduce the notion of weak convergence of probability measures on general (mostly Polish) spaces and derive the fundamental properties. The reader will profit from a solid knowledge of point set topology. Thus we start with a short overview of some topological definitions and theorems.

We do not strive for the greatest generality but rather content ourselves with the key theorems for probability theory. For further reading, we recommend [14, 82].

At first reading, the reader might wish to skip this rather analytically flavored chapter. In this case, for the time being it suffices to get acquainted with the definitions of weak convergence and tightness (Definitions 13.12 and 13.26), as well as with the statements of the Portemanteau theorem (Theorem 13.16) and Prohorov's theorem (Theorem 13.29).

5.13.1 13.1 A Topology Primer

Excursively, we present some definitions and facts from point set topology. For details, see, e.g., [90].

In the following, let (E, τ) be a topological space with the Borel σ -algebra $\mathcal{E} = \mathcal{B}(E)$ (compare Definitions 1.20 and 1.21). We will also assume that (E, τ) is a

Hausdorff space; that is, for any two points $x, y \in E$ with $x \neq y$, there exist disjoint open sets U, V such that $x \in U$ and $y \in V$.

For $A \subset E$, we denote by $\bar{A}$ the closure of A , by A° the interior and by ∂A the boundary of A . A set $A \subset E$ is called dense if $\bar{A} = E$.

(E, τ) is called metrizable if there exists a metric d on E such that τ is induced by the open balls $B_\varepsilon(x) := \{y \in E : d(x, y) < \varepsilon\}$. A metric d on E is called complete if any Cauchy sequence with respect to d converges in E . (E, τ) is called completely metrizable if there exists a complete metric on E that induces τ . If (E, d) is a metric space and $A, B \subset E$, then we write $d(A, B) = \inf\{d(x, y) : x \in A, y \in B\}$ and $d(x, B) := d(\{x\}, B)$ for $x \in E$.

A metrizable space (E, τ) is called separable if there exists a countable dense subset of E . Separability in metrizable spaces is equivalent to the existence of a countable base of the topology; that is, a countable set $\mathcal{U} \subset \tau$ with $A = \bigcup_{U \in \mathcal{U}: U \subset A} U$ for all $A \in \tau$. (For example, choose the ε -balls centered at the points of a countable subset and let ε run through the positive rational numbers.) A compact metric space is always separable (simply choose for each $n \in \mathbb{N}$ a finite cover $\mathcal{U}_n \subset \tau$ comprising balls of radius $\frac{1}{n}$ and then take $\mathcal{U} := \bigcup_{n \in \mathbb{N}} \mathcal{U}_n$).

A set $A \subset E$ is called compact if each open cover $\mathcal{U} \subset \tau$ of A (that is, $A \subset \bigcup_{U \in \mathcal{U}} U$) has a finite subcover; that is, a finite $\mathcal{U}' \subset \mathcal{U}$ with $A \subset \bigcup_{U \in \mathcal{U}'} U$. Compact sets are closed. By the Heine-Borel theorem, a subset of $\mathbb{R}^d$ is compact if and only if it is bounded and closed. $A \subset E$ is called relatively compact if $\bar{A}$ is compact. On the other hand, A is called sequentially compact (respectively relatively sequentially compact) if any sequence $(x_n)_{n \in \mathbb{N}}$ with values in A has a subsequence $(x_{n_k})_{k \in \mathbb{N}}$ that converges to some $x \in A$ (respectively $x \in \bar{A}$). In metrizable spaces,

the notions compact and sequentially compact coincide. A set $A \subset E$ is called σ -compact if A is a countable union of compact sets. E is called locally compact if any point $x \in E$ has an open neighborhood whose closure is compact. A locally compact, separable metric space is manifestly σ -compact. If E is a locally compact metric space and if $U \subset E$ is open and $K \subset U$ is compact, then there exists a compact set L with $K \subset L^\circ \subset L \subset U$. (For example, for any $x \in K$, take an open ball $B_{\varepsilon_x}(x)$ of radius $\varepsilon_x > 0$ that is contained in U and that is relatively compact. By making ε_x smaller (if necessary), one can assume that the closure of this ball is contained in U . As K is compact, there are finitely many points $x_1, \dots, x_n \in K$ with $K \subset V := \bigcup_{i=1}^n B_{\varepsilon_{x_i}}(x_i)$. By construction, $L = \bar{V} \subset U$ is compact.)

We present one type of topological space that is of particular importance in probability theory in a separate definition.

Definition 13.1 A topological space (E, τ) is called a Polish space if it is separable and if there exists a complete metric that induces the topology τ .

Examples of Polish spaces are countable discrete spaces (however, not $\mathbb{Q}$ with the usual topology), the Euclidean spaces $\mathbb{R}^n$, and the space $C([0, 1])$ of continuous functions $[0, 1] \rightarrow \mathbb{R}$, equipped with the supremum norm $\|\cdot\|_\infty$. In practice, all spaces that are of importance in probability theory are Polish spaces.

Let (E, d) be a metric space. A set $A \subset E$ is called totally bounded if, for any $\varepsilon > 0$, there exist finitely many points $x_1, \dots, x_n \in A$ such that $A \subset \bigcup_{i=1}^n B_\varepsilon(x_i)$. Evidently, compact sets are totally bounded. In Polish spaces, a partial converse is true.

Lemma 13.2 Let (E, τ) be a Polish space with complete metric d . A subset $A \subset E$ is totally bounded with respect to d if and only if A is relatively compact.

Proof This is left as an exercise. $\square$

In the following, let (E, τ) be a topological space with Borel σ -algebra $\mathcal{E} = \mathcal{B}(E) := \sigma(\tau)$ and with complete metric d . For measures on $(E, \mathcal{E})$, we introduce the following notions of regularity.

Definition 13.3 A σ -finite measure μ on $(E, \mathcal{E})$ is called

- (i) locally finite or a Borel measure if, for any point $x \in E$, there exists an open neighborhood $U \ni x$ such that $\mu(U) < \infty$,
- (ii) inner regular if

$$\mu(A) = \sup\{\mu(K) : K \subset A \text{ is compact}\} \quad \text{for all } A \in \mathcal{E}, \quad (13.1)$$

- (iii) outer regular if

$$\mu(A) = \inf\{\mu(U) : U \supset A \text{ is open}\} \quad \text{for all } A \in \mathcal{E}, \quad (13.2)$$

- (iv) regular if μ is inner and outer regular, and
- (v) a Radon measure if μ is an inner regular Borel measure.

Definition 13.4 We introduce the following spaces of measures on E :

$$\begin{aligned} \mathcal{M}(E) &:= \{ \text{Radon measures on } (E, \mathcal{E}) \}, \\ \mathcal{M}_f(E) &:= \{ \text{finite measures on } (E, \mathcal{E}) \}, \\ \mathcal{M}_1(E) &:= \{ \mu \in \mathcal{M}_f(E) : \mu(E) = 1 \}, \\ \mathcal{M}_{\leq 1}(E) &:= \{ \mu \in \mathcal{M}_f(E) : \mu(E) \leq 1 \}. \end{aligned}$$

The elements of $\mathcal{M}_{\leq 1}(E)$ are called sub-probability measures on E . Further, we agree on the following notation for spaces of continuous functions:

$$\begin{aligned} C(E) &:= \{ f : E \rightarrow \mathbb{R} \text{ is continuous} \}, \\ C_b(E) &:= \{ f \in C(E) \text{ is bounded} \}, \\ C_c(E) &:= \{ f \in C(E) \text{ has compact support} \} \subset C_b(E). \end{aligned}$$

Recall that the support of a real function f is $\overline{f^{-1}(\mathbb{R} \setminus \{0\})}$.

Unless otherwise stated, the vector spaces $C(E)$, $C_b(E)$ and $C_c(E)$ are equipped with the supremum norm.

Lemma 13.5 If E is Polish and $\mu \in \mathcal{M}_f(E)$, then for any $\varepsilon > 0$, there is a compact set $K \subset E$ with $\mu(E \setminus K) < \varepsilon$.

Proof Let $\varepsilon > 0$. For each $n \in \mathbb{N}$, there exists a sequence $x_1^n, x_2^n, \dots \in E$ with $E = \bigcup_{i=1}^{\infty} B_{1/n}(x_i^n)$. Fix $N_n \in \mathbb{N}$ such that $\mu\left(E \setminus \bigcup_{i=1}^{N_n} B_{1/n}(x_i^n)\right) < \frac{\varepsilon}{2^n}$. Define

$$A := \bigcap_{n=1}^{\infty} \bigcup_{i=1}^{N_n} B_{1/n}(x_i^n).$$

By construction, A is totally bounded. Since E is Polish, $\bar{A}$ is compact. Furthermore, it follows that $\mu(E \setminus \bar{A}) \leq \mu(E \setminus A) < \sum_{n=1}^{\infty} \varepsilon 2^{-n} = \varepsilon$.

Theorem 13.6 If E is Polish and if $\mu \in \mathcal{M}_f(E)$, then μ is regular. In particular, in this case, $\mathcal{M}_f(E) \subset \mathcal{M}(E)$.

Proof (Outer regularity) Step 1. Let $B \subset E$ be closed and let $\varepsilon > 0$. Let d be a complete metric on E . For $\delta > 0$, let

$$B_\delta := \{x \in E : d(x, B) < \delta\}$$

be the open δ -neighborhood of B . As B is closed, we have $\bigcap_{\delta>0} B_\delta = B$. Since μ is upper semicontinuous (Theorem 1.36), there is a $\delta > 0$ such that $\mu(B_\delta) \leq \mu(B) + \varepsilon$.

Step 2. Let $B \in \mathcal{E}$ and $\varepsilon > 0$. Consider the class of sets

$$\mathcal{A} := \{V \cap C : V \subset E \text{ open}, C \subset E \text{ closed}\}$$

Clearly, we have $\mathcal{E} = \sigma(\mathcal{A})$. It is easy to check that $\mathcal{A}$ is a semiring. Hence by the approximation theorem for measures (Theorem 1.65), there are mutually disjoint sets $A_n = V_n \cap C_n \in \mathcal{A}$, $n \in \mathbb{N}$, such that $B \subset A := \bigcup_{n=1}^{\infty} A_n$ and $\mu(A) \leq \mu(B) + \varepsilon/2$. As shown in the first step, for any $n \in \mathbb{N}$, there is an open set $W_n \supset C_n$ such that $\mu(W_n) \leq \mu(C_n) + \varepsilon 2^{-n-1}$. Hence also $U_n := V_n \cap W_n$ is open. Let $B \subset U := \bigcup_{n=1}^{\infty} U_n$. We conclude that $\mu(U) \leq \mu(A) + \sum_{n=1}^{\infty} \varepsilon 2^{-n-1} \leq \mu(B) + \varepsilon$.

(Inner regularity) Replacing B by B^c , the outer regularity yields the existence of a closed set $D \subset B$ with $\mu(B \setminus D) < \varepsilon/2$. By Lemma 13.5, there exists a compact set K with $\mu(K^c) < \varepsilon/2$. Define $C = D \cap K$. Then $C \subset B$ is compact and $\mu(B \setminus C) < \varepsilon$. Hence μ is also inner regular.

Corollary 13.7 The Lebesgue measure λ on $\mathbb{R}^d$ is a regular Radon measure. However, not all σ -finite measures on $\mathbb{R}^d$ are regular.

Proof Clearly, $\mathbb{R}^d$ is Polish and λ is locally finite. Let $A \in \mathcal{B}(\mathbb{R}^d)$ and $\varepsilon > 0$. There is an increasing sequence $(K_n)_{n \in \mathbb{N}}$ of compact sets with $K_n \uparrow \mathbb{R}^d$. Since any K_n is bounded, we have $\lambda(K_n) < \infty$. Hence, by the preceding theorem, for any $n \in \mathbb{N}$, there exists an open set $U_n \supset A \cap K_n$ with $\lambda(U_n \setminus A) < \varepsilon/2^n$. Thus $\lambda(U \setminus A) < \varepsilon$ for the open set $U := \bigcup_{n \in \mathbb{N}} U_n$.

If $\lambda(A) < \infty$, then there exists an $n \in \mathbb{N}$ with $\lambda(A \setminus K_n) < \varepsilon/2$. By the preceding theorem, there exists a compact set $C \subset A \cap K_n$ with $\lambda((A \cap K_n) \setminus C) < \varepsilon/2$. Therefore, $\lambda(A \setminus C) < \varepsilon$.

If, on the other hand, $\lambda(A) = \infty$, then for any $L > 0$, we have to find a compact set $C \subset A$ with $\lambda(C) > L$. However, $\lambda(A \cap K_n) \xrightarrow{n \rightarrow \infty} \infty$; hence there exists an $n \in \mathbb{N}$ with $\lambda(A \cap K_n) > L + 1$. By what we have shown already, there exists a compact set $C \subset A \cap K_n$ with $\lambda((A \cap K_n) \setminus C) < 1$; hence $\lambda(C) > L$.

Finally, consider the measure $\mu = \sum_{q \in \mathbb{Q}} \delta_q$. Clearly, this measure is σ -finite; however, it is neither locally finite nor outer regular. $\square$

Definition 13.8 Let (E, d_E) and (F, d_F) be metric spaces. A function $f : E \rightarrow F$ is called Lipschitz continuous if there exists a constant $K < \infty$, the so-called Lipschitz constant, with

$d_F(f(x), f(y)) \leq K \cdot d_E(x, y)$ for all $x, y \in E$. Denote by $\text{Lip}_K(E; F)$ the space of Lipschitz continuous functions with constant K and by $\text{Lip}(E; F) = \bigcup_{K>0} \text{Lip}_K(E; F)$ the space of Lipschitz continuous functions on E .

We abbreviate $\text{Lip}_K(E) := \text{Lip}_K(E; \mathbb{R})$ and $\text{Lip}(E) := \text{Lip}(E; \mathbb{R})$.

Definition 13.9 Let $\mathcal{F} \subset \mathcal{M}(E)$ be a family of Radon measures. A family $\mathcal{C}$ of measurable maps $E \rightarrow \mathbb{R}$ is called a separating family for $\mathcal{F}$ if, for any two measures $\mu, \nu \in \mathcal{F}$, the following holds:

$$\left(\int f d\mu = \int f d\nu \text{ for all } f \in \mathcal{C} \cap \mathcal{L}^1(\mu) \cap \mathcal{L}^1(\nu) \right) \implies \mu = \nu.$$

Lemma 13.10 Let (E, d) be a metric space. For any closed set $A \subset E$ and any $\varepsilon > 0$, there is a Lipschitz continuous map $\rho_{A,\varepsilon} : E \rightarrow [0, 1]$ with

$$\rho_{A,\varepsilon}(x) = \begin{cases} 1, & \text{if } x \in A \\ 0, & \text{if } d(x, A) \geq \varepsilon. \end{cases}$$

Proof Let $\varphi : \mathbb{R} \rightarrow [0, 1]$, $t \mapsto (t \vee 0) \wedge 1$. For $x \in E$, define $\rho_{A,\varepsilon}(x) = 1 - \varphi(\varepsilon^{-1}d(x, A))$. $\square$

Theorem 13.11 Let (E, d) be a metric space.

(i) $\text{Lip}_1(E; [0, 1])$ is separating for $\mathcal{M}(E)$.

(ii) If, in addition, E is locally compact, then $C_c(E) \cap \text{Lip}_1(E; [0, 1])$ is separating for $\mathcal{M}(E)$.

Proof (i) Assume $\mu_1, \mu_2 \in \mathcal{M}(E)$ are measures with $\int f d\mu_1 = \int f d\mu_2$ for all $f \in \text{Lip}_1(E; [0, 1])$. If $A \in \mathcal{E}$, then $\mu_i(A) = \sup \{ \mu_i(K) : K \subset A \text{ is compact} \}$ since the Radon measure μ_i is inner regular ($i = 1, 2$). Hence, it is enough to show that $\mu_1(K) = \mu_2(K)$ for any compact set K .

Now let $K \subset E$ be compact. Since μ_1 and μ_2 are locally finite, for every $x \in K$, there exists an open set $U_x \ni x$ with $\mu_1(U_x) < \infty$ and $\mu_2(U_x) < \infty$. Since K is compact, we can find finitely many points $x_1, \dots, x_n \in K$ such that $K \subset U := \bigcup_{j=1}^n U_{x_j}$. By construction, $\mu_i(U) < \infty$; hence $\mathbb{1}_U \in L^1(\mu_i)$ for $i = 1, 2$. Since U^c is closed and since $U^c \cap K = \emptyset$, we get $\delta := d(U^c, K) > 0$. Let $\rho_{K,\varepsilon}$ be the map from Lemma 13.10. Hence $\mathbb{1}_K \leq \rho_{K,\varepsilon} \leq \mathbb{1}_U \in L^1(\mu_i)$ if $\varepsilon \in (0, \delta)$. Since $\rho_{K,\varepsilon} \xrightarrow{\varepsilon \rightarrow 0} \mathbb{1}_K$, we get by dominated convergence (Corollary 6.26) that $\mu_i(K) = \lim_{\varepsilon \rightarrow 0} \int \rho_{K,\varepsilon} d\mu_i$. However, $\varepsilon \rho_{K,\varepsilon} \in \text{Lip}_1(E; [0, 1])$ for all $\varepsilon > 0$; hence, by assumption,

$$\int \rho_{K,\varepsilon} d\mu_1 = \varepsilon^{-1} \int (\varepsilon \rho_{K,\varepsilon}) d\mu_1 = \varepsilon^{-1} \int (\varepsilon \rho_{K,\varepsilon}) d\mu_2 = \int \rho_{K,\varepsilon} d\mu_2$$

This implies $\mu_1(K) = \mu_2(K)$; hence $\mu_1 = \mu_2$.

(ii) If E is locally compact, then in (i) we can choose the neighborhoods U_x to be relatively compact. Hence U is relatively compact; thus $\rho_{K,\varepsilon}$ has compact support and is thus in $C_c(E)$ for all $\varepsilon \in (0, \delta)$.

Exercise 13.1.1 (i) Show that $C([0, 1])$ has a separable dense subset.

(ii) Show that the space $(C_b([0, \infty)), \|\cdot\|_\infty)$ of bounded continuous functions, equipped with the supremum norm, is not separable.

(iii) Show that the space $C_c([0, \infty))$ of continuous functions with compact support, equipped with the supremum norm, is separable.

Exercise 13.1.2 Let μ be a locally finite measure. Show that $\mu(K) < \infty$ for any compact set K .

Exercise 13.1.3 (Lusin's theorem) Let Ω be a Polish space, let μ be a σ -finite measure on $(\Omega, \mathcal{B}(\Omega))$ and let $f : \Omega \rightarrow \mathbb{R}$ be a map. Show that the following two statements are equivalent:

(i) There is a Borel measurable map $g : \Omega \rightarrow \mathbb{R}$ with $f = g\mu$ -almost everywhere.

(ii) For any $\varepsilon > 0$, there is a compact set K_ε with $\mu(\Omega \setminus K_\varepsilon) < \varepsilon$ such that the restricted function $f|_{K_\varepsilon}$ is continuous.

Exercise 13.1.4 Let $\mathcal{U}$ be a family of intervals in $\mathbb{R}$ such that $W := \bigcup_{U \in \mathcal{U}} U$ has finite Lebesgue measure $\lambda(W)$. Show that for any $\varepsilon > 0$, there exist finitely many pairwise disjoint sets $U_1, \dots, U_n \in \mathcal{U}$ with

$$\sum_{i=1}^n \lambda(U_i) > \frac{1-\varepsilon}{3} \lambda(W).$$

Hint: Choose a finite family $\mathcal{U}' \subset \mathcal{U}$ such that $\bigcup_{U \in \mathcal{U}'} U$ has Lebesgue measure at least $(1-\varepsilon)\lambda(W)$. Choose a maximal sequence $\mathcal{U}''$ (sorted by decreasing lengths) of disjoint intervals and show that each $U \in \mathcal{U}'$ is in $(x-3a, x+3a)$ for some $(x-a, x+a) \in \mathcal{U}''$.

Exercise 13.1.5 Let $C \subset \mathbb{R}^d$ be an open, bounded and convex set and assume that $\mathcal{U} \subset \{x + rC : x \in \mathbb{R}^d, r > 0\}$ is such that $W := \bigcup_{U \in \mathcal{U}} U$ has finite Lebesgue measure $\lambda^d(W)$. Show that for any $\varepsilon > 0$, there exist finitely many pairwise disjoint sets $U_1, \dots, U_n \in \mathcal{U}$ such that

$$\sum_{i=1}^n \lambda^d(U_i) > \frac{1-\varepsilon}{3^d} \lambda(W).$$

Show by a counterexample that the condition of similarity of the open sets in $\mathcal{U}$ is essential.

Exercise 13.1.6 Let μ be a Radon measure on $\mathbb{R}^d$ and let $A \in \mathcal{B}(\mathbb{R}^d)$ be a μ -null set. Let $C \subset \mathbb{R}^d$ be bounded, convex and open with $0 \in C$. Use Exercise 13.1.5 to show that

$$\lim_{r \downarrow 0} \frac{\mu(x + rC)}{r^d} = 0 \quad \text{for } \lambda^d\text{-almost all } x \in A.$$

Conclude that if F is the distribution function of a Stieltjes measure μ on $\mathbb{R}$ and if $A \in \mathcal{B}(\mathbb{R})$ is a μ -null set, then $\frac{d}{dx}F(x) = 0$ for λ -almost all $x \in A$.

Exercise 13.1.7 (Fundamental theorem of calculus) (Compare [37].) Let $f \in \mathcal{L}^1(\mathbb{R}^d)$, $\mu = f\lambda^d$ and let $C \subset \mathbb{R}^d$ be open, convex and bounded with $0 \in C$. Show that

$$\lim_{r \downarrow 0} \frac{\mu(x + rC)}{r^d \lambda^d(C)} = f(x) \quad \text{for } \lambda^d\text{-almost all } x \in \mathbb{R}^d.$$

For the case $d = 1$, conclude the fundamental theorem of calculus:

$$\frac{d}{dx} \int_{[0,x]} f d\lambda = f(x) \quad \text{for } \lambda\text{-almost all } x \in \mathbb{R}$$

Hint: Use Exercise 13.1.6 with

$$\mu_q(dx) = (f(x) - q)^+ \lambda^d(dx) \quad \text{for } q \in \mathbb{Q}$$

as well as the inequality

$$\frac{\mu(x + rC)}{r^d \lambda^d(C)} \leq q + \frac{\mu_q(x + rC)}{r^d \lambda^d(C)}$$

Exercise 13.1.8 Similarly as in Corollary 13.7, show the following: Let E be a σ compact polish space and let μ be a measure on E . Then μ is a Radon measure if and only if $\mu(K) < \infty$ for any compact $K \subset E$.

5.13.2 13.2 Weak and Vague Convergence

In Theorem 13.11, we saw that integrals of bounded continuous functions f determine a Radon measure on a metric space (E, d) . If E is locally compact, it is enough to consider f with compact support. This suggests that we can use $C_b(E)$ and $C_c(E)$ as classes of test functions in order to define the convergence of measures.

Definition 13.12 (Weak and vague convergence) Let E be a metric space.

(i) Let $\mu, \mu_1, \mu_2, \dots \in \mathcal{M}_f(E)$. We say that $(\mu_n)_{n \in \mathbb{N}}$ converges weakly to μ , formally $\mu_n \xrightarrow{n \rightarrow \infty} \mu$ (weakly) or $\mu = w - \lim_{n \rightarrow \infty} \mu_n$, if

$$\int f d\mu_n \xrightarrow{n \rightarrow \infty} \int f d\mu \quad \text{for all } f \in C_b(E)$$

(ii) Let $\mu, \mu_1, \mu_2, \dots \in \mathcal{M}(E)$. We say that $(\mu_n)_{n \in \mathbb{N}}$ converges vaguely to μ , formally $\mu_n \xrightarrow{n \rightarrow \infty} \mu$ (vaguely) or $\mu = v - \lim_{n \rightarrow \infty} \mu_n$, if

$$\int f d\mu_n \xrightarrow{n \rightarrow \infty} \int f d\mu \quad \text{for any } f \in C_c(E)$$

Remark 13.13 If E is Polish, then by Theorems 13.6 and 13.11, the weak limit is unique. The same holds for the vague limit if E is locally compact.

Remark 13.14

(i) In functional analysis the notion of weak convergence is somewhat different. Starting from a normed vector space X (here the space of finite signed measures with the total variation norm), consider the space X' of continuous linear functionals $X \rightarrow \mathbb{R}$. The sequence (μ_n) in X converges weakly to $\mu \in X$, if $\Phi(\mu_n) \xrightarrow{n \rightarrow \infty} \Phi(\mu)$ for every $\Phi \in X'$. In the case of finite signed measures this is equivalent to: (μ_n) is bounded and $\mu_n(A) \xrightarrow{n \rightarrow \infty} \mu(A)$ for any measurable A (see [38, Theorem IV.9.5]). Comparing this to Theorem 13.16(vi), we see that the functional analysis notion of weak convergence is stronger than ours in Definition 13.12.

(ii) Weak convergence (as introduced in Definition 13.12) induces on $\mathcal{M}_f(E)$ the weak topology τ_w . This is the coarsest topology such that for all $f \in C_b(E)$, the map $\mathcal{M}_f(E) \rightarrow \mathbb{R}, \mu \mapsto \int f d\mu$ is continuous. In functional analysis, τ_w corresponds to the so-called weak*-topology. Starting from a normed vector space X (here $X = C_b(E)$ with the norm $\|\cdot\|_\infty$), we define the weak*-topology on the dual space X' by writing $\mu_n \xrightarrow{n \rightarrow \infty} \mu$ if and only if $\mu_n(x) \xrightarrow{n \rightarrow \infty} \mu(x)$ for all $x \in X$. Clearly, each μ defines a continuous linear form on $C_b(E)$ by $f \mapsto \mu(f) := \int f d\mu$. Hence $\mathcal{M}_f(E) \subset C_b(E)'$. This implies that τ_w is the trace of the weak*-topology on $\mathcal{M}_f(E)$.

(iii) If E is separable, then it can be shown that $(\mathcal{M}_f(E), \tau_w)$ is metrizable; for example, by virtue of the so-called Prohorov metric. This is defined by

$$d_P(\mu, \nu) := \max \{d'_P(\mu, \nu), d'_P(\nu, \mu)\} \quad (13.3)$$

where

$$d'_P(\mu, \nu) := \inf \{ \varepsilon > 0 : \mu(B) \leq \nu(B^\varepsilon) + \varepsilon \text{ for any } B \in \mathcal{B}(E) \}, \quad (13.4)$$

and where $B^\varepsilon = \{x : d(x, B) < \varepsilon\}$; see, e.g., [14, Appendix III, Theorem 5]. (It can be shown that $d'_P(\mu, \nu) = d'_P(\nu, \mu)$ if $\mu, \nu \in \mathcal{M}_1(E)$.) If E is locally compact and Polish, then $(\mathcal{M}_f(E), \tau_w)$ is again Polish (see [136, p. 167]).

(iv) Similarly, the vague topology τ_v on $\mathcal{M}(E)$ is the coarsest topology such that for all $f \in C_c(E)$, the map $\mathcal{M}(E) \rightarrow \mathbb{R}, \mu \mapsto \int f d\mu$ is continuous. If E is locally compact, then $(\mathcal{M}(E), \tau_v)$ is a Hausdorff space. If, in addition, E is Polish, then $(\mathcal{M}(E), \tau_v)$ is again Polish (see, e.g., [82, Section 15.7]).

While weak convergence implies convergence of the total masses (since $1 \in C_b(E)$), with vague convergence a mass defect (but not a mass gain) can be experienced in the limit.

Lemma 13.15 Let E be a locally compact Polish space and let $\mu, \mu_1, \mu_2, \dots \in \mathcal{M}(E)$ be measures such that $\mu_n \xrightarrow{n \rightarrow \infty} \mu$ vaguely. Then

$$\mu(E) \leq \liminf_{n \rightarrow \infty} \mu_n(E).$$

Proof Let $(f_N)_{N \in \mathbb{N}}$ be a sequence in $C_c(E; [0, 1])$ with $f_N \uparrow 1$. Then

$$\begin{aligned} \mu(E) &= \sup_{N \in \mathbb{N}} \int f_N d\mu \\ &= \sup_{N \in \mathbb{N}} \lim_{n \rightarrow \infty} \int f_N d\mu_n \\ &\leq \liminf_{n \rightarrow \infty} \sup_{N \in \mathbb{N}} \int f_N d\mu_n \\ &= \liminf_{n \rightarrow \infty} \mu_n(E) \end{aligned}$$

Clearly, the sequence $(\delta_{1/n})_{n \in \mathbb{N}}$ of probability measures on $\mathbb{R}$ converges weakly to δ_0 ; however, not in total variation norm. Indeed, for the closed set $(-\infty, 0]$, we have $\lim_{n \rightarrow \infty} \delta_{1/n}((-\infty, 0]) = 0 < 1 = \delta_0((-\infty, 0])$. Loosely speaking, at the boundaries of closed sets, mass can immigrate but not emigrate. The opposite is true for open sets: $\lim_{n \rightarrow \infty} \delta_{1/n}((0, \infty)) = 1 > 0 = \delta_0((0, \infty))$. Here mass can emigrate but not immigrate. In fact, weak convergence can be characterized by this property. In the following theorem, a whole bunch of such statements will be hung on a coat hanger (French: portemanteau).

For measurable $g : \Omega \rightarrow \mathbb{R}$, let U_g be the set of points of discontinuity of g . Recall from Exercise 1.1.3 that U_g is Borel measurable.

Theorem 13.16 (Portemanteau) Let E be a metric space and let $\mu, \mu_1, \mu_2, \dots \in \mathcal{M}_{\leq 1}(E)$. The following are equivalent.

- (i) $\mu = w - \lim_{n \rightarrow \infty} \mu_n$.
- (ii) $\int f d\mu_n \xrightarrow{n \rightarrow \infty} \int f d\mu$ for all bounded Lipschitz continuous f .
- (iii) $\int f d\mu_n \xrightarrow{n \rightarrow \infty} \int f d\mu$ for all bounded measurable f with $\mu(U_f) = 0$.
- (iv) $\liminf_{n \rightarrow \infty} \mu_n(E) \geq \mu(E)$ and $\limsup_{n \rightarrow \infty} \mu_n(F) \leq \mu(F)$ for all closed $F \subset E$.
- (v) $\limsup_{n \rightarrow \infty} \mu_n(E) \leq \mu(E)$ and $\liminf_{n \rightarrow \infty} \mu_n(G) \geq \mu(G)$ for all open $G \subset E$.
- (vi) $\lim_{n \rightarrow \infty} \mu_n(A) = \mu(A)$ for all measurable A with $\mu(\partial A) = 0$.

If E is locally compact and Polish, then in addition each of the following is equivalent to the previous statements.

- (vii) $\mu = v - \lim_{n \rightarrow \infty} \mu_n$ and $\mu(E) = \lim_{n \rightarrow \infty} \mu_n(E)$.
- (viii) $\mu = v - \lim_{n \rightarrow \infty} \mu_n$ and $\mu(E) \geq \limsup_{n \rightarrow \infty} \mu_n(E)$.

Proof "(iv) $\iff$ (v) $\implies$ (vi)" This is trivial.

"(iii) $\implies$ (i) $\implies$ (ii)" This is trivial.

"(ii) $\implies$ (iv)" Convergence of the total masses follows by using the test function $1 \in \text{Lip}(E; [0, 1])$. Let F be closed and let $\rho_{F, \varepsilon}$ be as in Lemma 13.10. Then

$$\limsup_{n \rightarrow \infty} \mu_n(F) \leq \inf_{\varepsilon > 0} \lim_{n \rightarrow \infty} \int \rho_{F, \varepsilon} d\mu_n = \inf_{\varepsilon > 0} \int \rho_{F, \varepsilon} d\mu = \mu(F)$$

since $\rho_{F, \varepsilon}(x) \xrightarrow{\varepsilon \rightarrow 0} \mathbb{1}_F(x)$ for all $x \in E$.

"(viii) $\implies$ (vii)" This is obvious by Lemma 13.15.

"(i) $\implies$ (vii)" This is clear since $C_c(E) \subset C_b(E)$ and $1 \in C_b(E)$.

"(vii) $\implies$ (v)" Let G be open and $\varepsilon > 0$. Since μ is inner regular (Theorem 13.6), there is a compact set $K \subset G$ with $\mu(G) - \mu(K) < \varepsilon$. As E is locally compact, there is a compact set L with $K \subset L^\circ \subset L \subset G$. Let $\delta := d(K, L^c) > 0$ and let $\rho_{K, \delta}$ be as in Lemma 13.10. Then $\mathbb{1}_K \leq \rho_{K, \delta} \leq \mathbb{1}_L$; hence $\rho_{K, \delta} \in C_c(E)$ and thus

$$\liminf_{n \rightarrow \infty} \mu_n(G) \geq \liminf_{n \rightarrow \infty} \int \rho_{K, \delta} d\mu_n = \int \rho_{K, \delta} d\mu \geq \mu(K) \geq \mu(G) - \varepsilon$$

Letting $\varepsilon \rightarrow 0$, we get (v).

"(vi) $\implies$ (iii)" Let $f : E \rightarrow \mathbb{R}$ be bounded and measurable with $\mu(U_f) = 0$. We make the

elementary observation that for all $D \subset \mathbb{R}$,

$$\partial f^{-1}(D) \subset f^{-1}(\partial D) \cup U_f \quad (13.5)$$

Indeed, if f is continuous at $x \in E$, then for any $\delta > 0$, there is an $\varepsilon(\delta) > 0$ with $f(B_{\varepsilon(\delta)}(x)) \subset B_\delta(f(x))$. If $x \in \partial f^{-1}(D)$, then there are $y \in f^{-1}(D) \cap B_{\varepsilon(\delta)}(x)$ and $z \in f^{-1}(D^c) \cap B_{\varepsilon(\delta)}(x)$. Therefore, $f(y) \in B_\delta(f(x)) \cap D \neq \emptyset$ and $f(z) \in B_\delta(f(x)) \cap D^c \neq \emptyset$; hence $f(x) \in \partial D$.

Let $\varepsilon > 0$. Evidently, the set $A := \{y \in \mathbb{R} : \mu(f^{-1}(\{y\})) > 0\}$ of atoms of the finite measure $\mu \circ f^{-1}$ is at most countable. Hence, there exist $N \in \mathbb{N}$ and $y_0 \leq -\|f\|_\infty < y_1 < \dots < y_{N-1} < \|f\|_\infty$ such that

$$y_i \in \mathbb{R} \setminus A \quad \text{and} \quad |y_{i+1} - y_i| < \varepsilon \quad \text{for all } i.$$

Let $E_i = f^{-1}([y_{i-1}, y_i])$ for $i = 1, \dots, N$. Then $E = \bigsqcup_{i=1}^N E_i$ and by (13.5),

$$\mu(\partial E_i) \leq \mu(f^{-1}(\{y_{i-1}\})) + \mu(f^{-1}(\{y_i\})) + \mu(U_f) = 0.$$

Therefore,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \int f d\mu_n &\leq \limsup_{n \rightarrow \infty} \sum_{i=1}^N \mu_n(E_i) \cdot y_i = \sum_{i=1}^N \mu(E_i) \cdot y_i \\ &\leq \varepsilon + \int f d\mu \end{aligned}$$

We let $\varepsilon \rightarrow 0$ and obtain $\limsup_{n \rightarrow \infty} \int f d\mu_n \leq \int f d\mu$. Finally, consider $(-f)$ to obtain the reverse inequality $\liminf_{n \rightarrow \infty} \int f d\mu_n \geq \int f d\mu$.

Definition 13.17 Let $X, X_1, X_2, \dots$ be random variables with values in E . We say that $(X_n)_{n \in \mathbb{N}}$ converges in distribution to X , formally $X_n \xrightarrow{\mathcal{D}} X$ or $X_n \xrightarrow{n \rightarrow \infty} X$, if the distributions converge weakly and hence if $P_X = \text{w-lim}_{n \rightarrow \infty} \mathbf{P}_{X_n}$. Sometimes we write $X_n \xrightarrow{\mathcal{D}} \mathbf{P}_X$ or $X_n \xrightarrow{n \rightarrow \infty} \mathbf{P}_X$ if we want to specify only the distribution $\mathbf{P}_X$ but not the random variable X .

Theorem 13.18 (Slutzky's theorem) Let $X, X_1, X_2, \dots$ and $Y_1, Y_2, \dots$ be random variables with values in E . Assume $X_n \xrightarrow{\mathcal{D}} X$ and $d(X_n, Y_n) \xrightarrow{n \rightarrow \infty} 0$ in probability. Then $Y_n \xrightarrow{\mathcal{D}} X$.

Proof Let $f : E \rightarrow \mathbb{R}$ be bounded and Lipschitz continuous with constant K . Then

$$|f(x) - f(y)| \leq Kd(x, y) \wedge 2\|f\|_\infty \quad \text{for all } x, y \in E.$$

Dominated convergence yields $\limsup_{n \rightarrow \infty} \mathbf{E}[|f(X_n) - f(Y_n)|] = 0$. Hence we have

$$\begin{aligned} \limsup_{n \rightarrow \infty} |\mathbf{E}[f(Y_n)] - \mathbf{E}[f(X)]| \\ \leq \limsup_{n \rightarrow \infty} |\mathbf{E}[f(X)] - \mathbf{E}[f(X_n)]| + \limsup_{n \rightarrow \infty} |\mathbf{E}[f(X_n) - f(Y_n)]| = 0. \end{aligned}$$

Corollary 13.19 If $X_n \xrightarrow{n \rightarrow \infty} X$ in probability, then $X_n \xrightarrow{\mathcal{D}} X, n \rightarrow \infty$. The converse is false in general.

Example 13.20 If $X, X_1, X_2, \dots$ are i.i.d. (with nontrivial distribution), then trivially $X_n \xrightarrow{\mathcal{D}} X$ but not $X_n \xrightarrow{n \rightarrow \infty} X$ in probability.

Recall the definition of a distribution function of a probability measure from Definition 1.59.

Definition 13.21 Let $F, F_1, F_2, \dots$ be distribution functions of probability measures on $\mathbb{R}$. We say that $(F_n)_{n \in \mathbb{N}}$ converges weakly to F , formally $F_n \xrightarrow{n \rightarrow \infty} F, F_n \xrightarrow{\mathcal{D}} F$ or $F = \text{w-lim}_{n \rightarrow \infty} F_n$, if

$$F(x) = \lim_{n \rightarrow \infty} F_n(x) \quad \text{for all points of continuity } x \text{ of } F. \quad (13.6)$$

If $F, F_1, F_2, \dots$ are distribution functions of sub-probability measures, then we define $F(\infty) := \lim_{x \rightarrow \infty} F(x)$ and for weak convergence require in addition $F(\infty) \geq \limsup_{n \rightarrow \infty} F_n(\infty)$.

Note that (13.6) implies $F(\infty) \leq \liminf_{n \rightarrow \infty} F_n(\infty)$. Hence, if $F_n \xrightarrow{\mathcal{D}} F$, then $F(\infty) = \lim_{n \rightarrow \infty} F_n(\infty)$.

Example 13.22 If F is the distribution function of a probability measure on $\mathbb{R}$ and $F_n(x) := F(x+n)$ for $x \in \mathbb{R}$, then $(F_n)_{n \in \mathbb{N}}$ converges pointwise to 1. However, this is not a distribution function, as 1 does not converge to 0 for $x \rightarrow -\infty$. On the other hand, if $G_n(x) = F(x-n)$, then $(G_n)_{n \in \mathbb{N}}$ converges pointwise to $G \equiv 0$. However, $G(\infty) = 0 < \limsup_{n \rightarrow \infty} G_n(\infty) = 1$; hence we do not have weak convergence here either. Indeed, in each case, there is a mass defect in the limit (in the case of the F_n on the left and in the case of the G_n on the right). However, the definition of weak convergence of distribution functions is constructed so that no mass defect occurs in the limit.

Theorem 13.23 Let $\mu, \mu_1, \mu_2, \dots \in \mathcal{M}_{\leq 1}(\mathbb{R})$ with corresponding distribution functions $F, F_1, F_2, \dots$. The following are equivalent.

(i) $\mu = \text{w-}\lim_{n \rightarrow \infty} \mu_n$.

(ii) $F_n \xrightarrow{\mathcal{D}} F$.

Proof "(i) $\implies$ (ii)" Let F be continuous at x . Then $\mu(\partial(-\infty, x]) = \mu(\{x\}) = 0$. By Theorem 13.16, $F_n(x) = \mu_n((-\infty, x]) \xrightarrow{n \rightarrow \infty} \mu((-\infty, x]) = F(x)$.

"(ii) $\implies$ (i)" Let $f \in \text{Lip}_1(\mathbb{R}; [0, 1])$. By Theorem 13.16, it is enough to show that

$$\int f d\mu_n \xrightarrow{n \rightarrow \infty} \int f d\mu \quad (13.7)$$

Let $\varepsilon > 0$. Fix $N \in \mathbb{N}$ and choose $N+1$ points of continuity $y_0 < y_1 < \dots < y_N$ of F such that $F(y_0) < \varepsilon$, $F(y_N) > F(\infty) - \varepsilon$ and $y_i - y_{i-1} < \varepsilon$ for all i . Then

$$\begin{aligned} \int f d\mu_n &\leq (F_n(y_0) + F_n(\infty) - F_n(y_N)) \\ &\quad + \sum_{i=1}^N (f(y_i) + \varepsilon) (F_n(y_i) - F_n(y_{i-1})). \end{aligned}$$

By assumption, $\lim_{n \rightarrow \infty} F_n(\infty) = F(\infty)$ and $F_n(y_i) \xrightarrow{n \rightarrow \infty} F(y_i)$ for every $i = 0, \dots, N$; hence

$$\limsup_{n \rightarrow \infty} \int f d\mu_n \leq 3\varepsilon + \sum_{i=1}^N f(y_i) (F(y_i) - F(y_{i-1})) \leq 4\varepsilon + \int f d\mu.$$

Therefore,

$$\limsup_{n \rightarrow \infty} \int f d\mu_n \leq \int f d\mu$$

Replacing f by $(1 - f)$, we get (13.7).

Corollary 13.24 Let $X, X_1, X_2, \dots$ be real random variables with distribution functions $F, F_1, F_2, \dots$. Then the following are equivalent.

(i) $X_n \xrightarrow{\mathcal{D}} X$.

(ii) $\mathbf{E}[f(X_n)] \xrightarrow{n \rightarrow \infty} \mathbf{E}[f(X)]$ for all $f \in C_b(\mathbb{R})$.

(iii) $F_n \xrightarrow{\mathcal{D}} F$.

How stable is weak convergence if we pass to image measures under some map φ ? Clearly, we need a certain continuity of φ at least at those points where the limit measure puts mass. The following theorem formalizes this idea and will come in handy in many applications.

Theorem 13.25 (Continuous mapping theorem) Let (E_1, d_1) and (E_2, d_2) be metric spaces and let $\varphi : E_1 \rightarrow E_2$ be measurable. Denote by U_φ the set of points of discontinuity of φ .

(i) If $\mu, \mu_1, \mu_2, \dots \in \mathcal{M}_{\leq 1}(E_1)$ with $\mu(U_\varphi) = 0$ and $\mu_n \xrightarrow{n \rightarrow \infty} \mu$ weakly, then $\mu_n \circ \varphi^{-1} \xrightarrow{n \rightarrow \infty} \mu \circ \varphi^{-1}$ weakly.

(ii) If $X, X_1, X_2, \dots$ are E_1 -valued random variables with $\mathbf{P}[X \in U_\varphi] = 0$ and $X_n \xrightarrow{\mathcal{D}} X$, then $\varphi(X_n) \xrightarrow{\mathcal{D}} \varphi(X)$.

Proof First note that $U_\varphi \subset E_1$ is Borel measurable by Exercise 1.1.3. Hence the conditions make sense.

(i) Let $f \in C_b(E_2)$. Then $f \circ \varphi$ is bounded and measurable and $U_{f \circ \varphi} \subset U_\varphi$; hence $\mu(U_{f \circ \varphi}) = 0$. By Theorem 13.16,

$$\begin{aligned} \lim_{n \rightarrow \infty} \int f d(\mu_n \circ \varphi^{-1}) &= \lim_{n \rightarrow \infty} \int (f \circ \varphi) d\mu_n \\ &= \int (f \circ \varphi) d\mu = \int f d(\mu \circ \varphi^{-1}) \end{aligned}$$

(ii) This is obvious since $\mathbf{P}_{\varphi(X)} = \mathbf{P}_X \circ \varphi^{-1}$.

Exercise 13.2.1 Recall d'_P from (13.4). Show that $d_P(\mu, \nu) = d'_P(\mu, \nu) = d'_P(\nu, \mu)$ for all $\mu, \nu \in \mathcal{M}_1(E)$.

Exercise 13.2.2 Show that the topology of weak convergence on $\mathcal{M}_f(E)$ is coarser than the topology induced on $\mathcal{M}_f(E)$ by the total variation norm (see Corollary 7.45). That is, $\|\mu_n - \mu\|_{TV} \xrightarrow{n \rightarrow \infty} 0$ implies $\mu_n \xrightarrow{n \rightarrow \infty} \mu$ weakly.

Exercise 13.2.3 Let $E = \mathbb{R}$ and $\mu_n = \frac{1}{n} \sum_{k=0}^n \delta_{k/n}$. Let $\mu = \lambda|_{[0,1]}$ be the Lebesgue measure restricted to $[0, 1]$. Show that $\mu = \text{w-}\lim_{n \rightarrow \infty} \mu_n$.

Exercise 13.2.4 Let $E = \mathbb{R}$ and λ be the Lebesgue measure on $\mathbb{R}$. For $n \in \mathbb{N}$, let $\mu_n = \lambda|_{[-n, n]}$. Show that $\lambda = \text{v-}\lim_{n \rightarrow \infty} \mu_n$ but that $(\mu_n)_{n \in \mathbb{N}}$ does not converge weakly.

Exercise 13.2.5 Let $E = \mathbb{R}$ and $\mu_n = \delta_n$ for $n \in \mathbb{N}$. Show that $\text{v-}\lim_{n \rightarrow \infty} \mu_n = 0$ but that $(\mu_n)_{n \in \mathbb{N}}$ does not converge weakly.

Exercise 13.2.6 (Lévy metric) For two probability distribution functions F and G on $\mathbb{R}$, define the Lévy distance by

$$d(F, G) = \inf\{\varepsilon \geq 0 : G(x - \varepsilon) - \varepsilon \leq F(x) \leq G(x + \varepsilon) + \varepsilon \text{ for all } x \in \mathbb{R}\}.$$

Show the following:

(i) d is a metric on the set of distribution functions.

(ii) $F_n \xrightarrow{n \rightarrow \infty} F$ if and only if $d(F_n, F) \xrightarrow{n \rightarrow \infty} 0$.

(iii) For every $P \in \mathcal{M}_1(\mathbb{R})$, there is a sequence $(P_n)_{n \in \mathbb{N}}$ in $\mathcal{M}_1(\mathbb{R})$ such that each P_n has finite support and such that $P_n \xrightarrow{n \rightarrow \infty} P$.

Exercise 13.2.7 We can extend the notions of weak convergence and vague convergence to signed measures; that is, to differences $\varphi := \mu^+ - \mu^-$ of measures from $\mathcal{M}_f(E)$ and $\mathcal{M}(E)$, respectively, by repeating the words of Definition 13.12 for these classes. Show that the topology of weak convergence is not metrizable in general.

Hint: Consider $E = [0, 1]$.

(i) For $n \in \mathbb{N}$, define $\varphi_n = \delta_{1/n} - \delta_{2/n}$. Show that, for any $C > 0$, $(C\varphi_n)_{n \in \mathbb{N}}$ converges weakly to the zero measure.

(ii) Assume there is a metric that induces weak convergence. Show that then there would be a sequence $(C_n)_{n \in \mathbb{N}}$ with $C_n \uparrow \infty$ and $0 = \text{w-}\lim_{n \rightarrow \infty} (C_n \varphi_n)$.

(iii) Choose an $f \in C([0, 1])$ with $f(2^{-n}) = (-1)^n C_n^{-1/2}$ for any $n \in \mathbb{N}$, and show that

$(\int f d(C_n \varphi_n))_{n \in \mathbb{N}}$ does not converge to zero.

(iv) Use this construction to contradict the assumption of metrizable.

Exercise 13.2.8 Show that (13.3) defines a metric on $\mathcal{M}_1(E)$ and that this metric induces the topology of weak convergence.

Exercise 13.2.9 Show the implication "(vi) $\implies$ (iv)" of Theorem 13.16 directly.

Exercise 13.2.10 Let $X, X_1, X_2, \dots$ and $Y_1, Y_2, \dots$ be real random variables. Assume $\mathbf{P}_{Y_n} = \mathcal{N}_{0,1/n}$ for all $n \in \mathbb{N}$. Show that $X_n \xrightarrow{\mathcal{D}} X$ if and only if $X_n + Y_n \xrightarrow{\mathcal{D}} X$.

Exercise 13.2.11 For each $n \in \mathbb{N}$, let X_n be a geometrically distributed random variable with parameter $p_n \in (0, 1)$. How must we choose the sequence $(p_n)_{n \in \mathbb{N}}$ in order that $\mathbf{P}_{X_n/n}$ converges weakly to the exponential distribution with parameter $\alpha > 0$?

Exercise 13.2.12 Let $X, X_1, X_2, \dots$ be real random variables with $X_n \xrightarrow{n \rightarrow \infty} X$. Show the following.

(i) $\mathbf{E}[|X|] \leq \liminf_{n \rightarrow \infty} \mathbf{E}[|X_n|]$.

(ii) Let $r > p > 0$. If $\sup_{n \in \mathbb{N}} \mathbf{E}[|X_n|^r] < \infty$, then $\mathbf{E}[|X|^p] = \lim_{n \rightarrow \infty} \mathbf{E}[|X_n|^p]$.

Exercise 13.2.13 Let $F, F_1, F_2, \dots$ be probability distribution functions on $\mathbb{R}$, and assume $F_n \xrightarrow{n \rightarrow \infty} F$. Let $F^{-1}(u) = \inf\{x \in \mathbb{R} : F(x) \geq u\}$, $u \in (0, 1)$, be the left continuous inverse of F (see the proof of Theorem 1.104). Show that

$$F_n^{-1}(u) \xrightarrow{n \rightarrow \infty} F^{-1}(u) \text{ at every point of continuity } u \text{ of } F^{-1}.$$

Conclude that $F^{-1}(u) \xrightarrow{n \rightarrow \infty} F^{-1}(u)$ for Lebesgue almost all $u \in (0, 1)$.

Exercise 13.2.14 Let $\mu, \mu_1, \mu_2, \dots \in \mathcal{M}_1(\mathbb{R})$ with $\mu_n \xrightarrow{n \rightarrow \infty} \mu$ weakly. Show that there exists a probability space $(\Omega, \mathcal{A}, \mathbf{P})$ and real random variables $X, X_1, X_2, \dots$ on $(\Omega, \mathcal{A}, \mathbf{P})$ with distributions $\mathbf{P}_X = \mu$ and $\mathbf{P}_{X_n} = \mu_n$, $n \in \mathbb{N}$, such that

$$X_n \xrightarrow{n \rightarrow \infty} X \quad \mathbf{P}\text{-a.s.}$$

Hint: Use Exercise 13.2.13.

Exercise 13.2.15 Let (E, d) be a metric space and let $\mu, \mu_1, \mu_2, \dots$ be probability measures on E . A measurable map $f : E \rightarrow \mathbb{R}$ is called uniformly integrable with respect to $(\mu_n)_{n \in \mathbb{N}}$, if

$$\inf_{a > 0} \sup_{n \in \mathbb{N}} \int_{\{|f| > a\}} |f| d\mu_n = 0$$

Let f be continuous and uniformly integrable with respect to $(\mu_n)_{n \in \mathbb{N}}$ and assume that $\mu_n \xrightarrow{n \rightarrow \infty} \mu$ weakly. Show that $\int |f| d\mu < \infty$ and that

$$\int f d\mu_n \xrightarrow{n \rightarrow \infty} \int f d\mu$$

Hint: Apply Exercise 13.2.14 to the image measures $\mu_n \circ f^{-1}$.

5.13.3 13.3 Prohorov's Theorem

In the following, let E be a Polish space with Borel σ -algebra $\mathcal{E}$. A fundamental question is: When does a sequence $(\mu_n)_{n \in \mathbb{N}}$ of measures on $(E, \mathcal{E})$ converge weakly or does at least have a weak limit point? Evidently, a necessary condition is that $(\mu_n(E))_{n \in \mathbb{N}}$ is bounded. Hence, without loss of generality, we will consider only sequences in $\mathcal{M}_{\leq 1}(E)$. However, this condition is not sufficient for the existence of weak limit points, as for example the sequence $(\delta_n)_{n \in \mathbb{N}}$ of probability measures on $\mathbb{R}$ does not have a weak limit point (although it converges vaguely to the zero measure). This example suggests that we also have to make sure that no mass "vanishes at infinity". The idea will be made precise by the notion of tightness.

We start this section by presenting as the main result Prohorov's theorem [136]. We give the proof first for the special case $E = \mathbb{R}$ and then come to a couple of applications. The full proof of the general case is deferred to the end of the section.

Definition 13.26 (Tightness) A family $\mathcal{F} \subset \mathcal{M}_f(E)$ is called tight if, for any $\varepsilon > 0$, there exists a compact set $K \subset E$ such that

$$\sup\{\mu(E \setminus K) : \mu \in \mathcal{F}\} < \varepsilon$$

Remark 13.27 If E is Polish, then by Lemma 13.5, every singleton $\{\mu\} \subset \mathcal{M}_f(E)$ is tight and thus so is every finite family.

Example 13.28

- (i) If E is compact, then $\mathcal{M}_1(E)$ and $\mathcal{M}_{\leq 1}(E)$ are tight.
- (ii) If $(X_i)_{i \in I}$ is an arbitrary family of random variables with

$$C := \sup\{\mathbf{E}[|X_i|] : i \in I\} < \infty,$$

then $\{\mathbf{P}_{X_i} : i \in I\}$ is tight. Indeed, for $\varepsilon > 0$ and $K = [-C/\varepsilon, C/\varepsilon]$, by Markov's inequality, $\mathbf{P}_{X_i}(\mathbb{R} \setminus K) = \mathbf{P}[|X_i| > C/\varepsilon] \leq \varepsilon$.

- (iii) The family $(\delta_n)_{n \in \mathbb{N}}$ of probability measures on $\mathbb{R}$ is not tight.
- (iv) The family $(\mathcal{U}_{[-n, n]})_{n \in \mathbb{N}}$ of uniform distributions on the intervals $[-n, n]$, regarded as measures on $\mathbb{R}$, is not tight.

Recall that a family $\mathcal{F}$ of measures is called weakly relatively sequentially compact if every sequence in $\mathcal{F}$ has a weak limit point (in the closure of $\mathcal{F}$).

Theorem 13.29 (Prohorov's theorem (1956)) Let (E, d) be a metric space and $\mathcal{F} \subset \mathcal{M}_{\leq 1}(E)$. Then:

- (i) $\mathcal{F}$ is tight $\implies \mathcal{F}$ is weakly relatively sequentially compact.
- (ii) If E is Polish, then also the converse holds:

$$\mathcal{F} \text{ is tight} \iff \mathcal{F} \text{ is weakly relatively sequentially compact.}$$

Corollary 13.30 Let E be a compact metric space. Then the sets $\mathcal{M}_{\leq 1}(E)$ and $\mathcal{M}_1(E)$ are weakly sequentially compact.

Corollary 13.31 If E is a locally compact separable metric space, then $\mathcal{M}_{\leq 1}(E)$ is vaguely sequentially compact.

Proof Let $x_1, x_2, \dots$ be dense in E . As E is locally compact, for each $n \in \mathbb{N}$, there exists an open neighborhood $U_n \ni x_n$ whose closure $\bar{U}_n$ is compact. Hence, also $V_n := \bigcup_{i=1}^n V_i$ is relatively compact for each $n \in \mathbb{N}$. This implies that $N(K) := \min\{m : K \subset V_m\}$ is finite for any compact $K \subset E$. Inductively define $W_1 := V_1$ and $W_{n+1} := W_{N(\bar{V}_n)}, n \in \mathbb{N}$. Then W_n is open, $\bar{W}_n$ is compact, and $\bar{W}_n \subset W_{n+1}$ for all $n \in \mathbb{N}$. Furthermore, we have $W_n \uparrow E$.

Applying Prohorov's theorem (i.e., Corollary 13.30) to the measures $(\mu_k \mathbb{1}_{\bar{W}_n})_{k \in \mathbb{N}}$, for each $n \in \mathbb{N}$, we can choose a sequence $(k_l^n)_{l \in \mathbb{N}}$ and a measure $\tilde{\mu}_n := w - \lim_{l \rightarrow \infty} \mu_{k_l^n} \mathbb{1}_{\bar{W}_n}$ whose support lies in $\bar{W}_n$. We may assume that the sequences $(k_l^n)_{l \in \mathbb{N}}$ were chosen successively such that (k_l^{n+1}) is a subsequence of (k_l^n) .

Note that we have $\tilde{\mu}_n(\bar{W}_n) \leq \tilde{\mu}_{n+1}(\bar{W}_n)$, but equality does not hold in general. For $f \in C_c(E)$, there exists an $n_0 \in \mathbb{N}$ such that the support of f is contained in W_{n_0} . Hence, for $m \geq n \geq n_0$, we have

$$\begin{aligned}
\int f d\tilde{\mu}_n &= \lim_{l \rightarrow \infty} \int f \mathbb{1}_{\bar{W}_n} d\mu_{k_l^n} \\
&= \lim_{l \rightarrow \infty} \int f \mathbb{1}_{\bar{W}_n} d\mu_{k_l^m} \\
&= \lim_{l \rightarrow \infty} \int f \mathbb{1}_{\bar{W}_m} d\mu_{k_l^m} = \int f d\tilde{\mu}_m
\end{aligned}$$

and thus

$$\int f d\tilde{\mu}_n = \lim_{m \rightarrow \infty} \int f d\mu_{k_m^m}.$$

This implies that for any measurable relatively compact set $A \subset E$, we have

$$\tilde{\mu}_m(A) = \tilde{\mu}_{N(\bar{A})} \quad \text{for any } m \geq N(\bar{A})$$

For any measurable set $A \subset E$, define

$$\mu(A) := \sup_{n \in \mathbb{N}} \sup_{m > n} \tilde{\mu}_m(A \cap W_n) = \sup_{n \in \mathbb{N}} \tilde{\mu}_{n+1}(A \cap W_n)$$

It is easy to check that μ is a lower semicontinuous content and is hence a measure (see Theorem 1.36). By construction, for any $f \in C_c(E)$, we infer

$$\int f d\mu = \lim_{n \rightarrow \infty} \int f d\mu_{k_n^n}$$

Concluding, we have $\mu = \nu - \lim_{n \rightarrow \infty} \mu_{k_n^n}$.

Remark 13.32 The implication (ii) in Theorem 13.29 is less useful but a lot simpler to prove. Here we need that E is Polish since clearly every singleton is weakly compact but is tight only under additional assumptions; for example, if E is Polish (see Lemma 13.5).

Proof of Theorem 13.29(ii) We start as in the proof of Lemma 13.5. Let $\{x_1, x_2, \dots\} \subset E$ be dense. For $n \in \mathbb{N}$, define $A_{n,N} := \bigcup_{i=1}^N B_{1/n}(x_i)$. Then $A_{n,N} \uparrow E$ for $N \rightarrow \infty$ for all $n \in \mathbb{N}$. Let

$$\delta := \sup_{n \in \mathbb{N}} \inf_{N \in \mathbb{N}} \sup_{\mu \in \mathcal{F}} \mu(A_{n,N}^c)$$

Then there is an $n \in \mathbb{N}$ such that for any $N \in \mathbb{N}$, there is a $\mu_N \in \mathcal{F}$ with $\mu_N(A_{n,N}^c) \geq \delta/2$. As $\mathcal{F}$ is weakly relatively sequentially compact, $(\mu_N)_{N \in \mathbb{N}}$ has a weakly convergent subsequence $(\mu_{N_k})_{k \in \mathbb{N}}$ whose weak limit will be denoted by $\mu \in \mathcal{M}_{\leq 1}(E)$. By the Portemanteau theorem (Theorem 13.16(iv)), for any $N \in \mathbb{N}$,

$$\mu(A_{n,N}^c) \geq \liminf_{k \rightarrow \infty} \mu_{N_k}(A_{n,N}^c) \geq \liminf_{k \rightarrow \infty} \mu_{N_k}(A_{n,N_k}^c) \geq \delta/2.$$

On the other hand, $A_{n,N}^c \downarrow \emptyset$ for $N \rightarrow \infty$; hence $\mu(A_{n,N}^c) \xrightarrow{N \rightarrow \infty} 0$. Thus $\delta = 0$.

Now fix $\varepsilon > 0$. By the above, for any $n \in \mathbb{N}$, we can choose an $N'_n \in \mathbb{N}$ such that $\mu(A_{n,N'_n}^c) < \varepsilon/2^n$ for all $\mu \in \mathcal{F}$. By construction, the set $A := \bigcap_{n=1}^{\infty} A_{n,N'_n}$ is totally bounded and hence relatively compact. Further, for every $\mu \in \mathcal{F}$,

$$\mu((\bar{A})^c) \leq \mu(A^c) \leq \sum_{n=1}^{\infty} \mu(A_{n,N'_n}^c) \leq \varepsilon.$$

Hence $\mathcal{F}$ is tight.

The other implication in Prohorov's theorem is more difficult to prove, especially in the case of a general metric space. For this reason, we first give a proof only for

the case $E = \mathbb{R}$ and come to applications before proving the difficult implication in the general situation.

The problem consists in finding a candidate for a weak limit point. For distributions on $\mathbb{R}$, the problem is equivalent to finding a weak limit point for a sequence of distribution functions. Here Helly's theorem is the tool. It is based on a diagonal sequence argument that will be recycled later in the proof of Prohorov's theorem in the general case.

Let

$$V = \{F : \mathbb{R} \rightarrow \mathbb{R} \text{ is right continuous, monotone increasing and bounded} \}$$

be the set of distribution functions of finite measures on $\mathbb{R}$.

Theorem 13.33 (Helly's theorem) Let $(F_n)_{n \in \mathbb{N}}$ be a uniformly bounded sequence in V . Then there exists an $F \in V$ and a subsequence $(F_{n_k})_{k \in \mathbb{N}}$ with

$$F_{n_k}(x) \xrightarrow{k \rightarrow \infty} F(x) \text{ at all points of continuity of } F.$$

Proof We use a diagonal sequence argument. Choose an enumeration of the rational numbers $\mathbb{Q} = \{q_1, q_2, q_3, \dots\}$. By the Bolzano-Weierstraß theorem, the sequence $(F_n(q_1))_{n \in \mathbb{N}}$ has a convergent subsequence $(F_{n_k^1}(q_1))_{k \in \mathbb{N}}$. Analogously, we find a subsequence $(n_k^2)_{k \in \mathbb{N}}$ of $(n_k^1)_{k \in \mathbb{N}}$ such that $(F_{n_k^2}(q_2))_{k \in \mathbb{N}}$ converges. Inductively, we obtain subsequences $(n_k^1) \supset (n_k^2) \supset (n_k^3) \supset \dots$ such that $(F_{n_k^l}(q_l))_{k \in \mathbb{N}}$ converges for all $l \in \mathbb{N}$. Now define $n_k := n_k^k$. Then $(F_{n_k}(q))_{k \in \mathbb{N}}$ converges for all $q \in \mathbb{Q}$. Define $\tilde{F}(q) = \lim_{k \rightarrow \infty} F_{n_k}(q)$ and

$$F(x) = \inf\{\tilde{F}(q) : q \in \mathbb{Q} \text{ with } q > x\}$$

As $\tilde{F}$ is monotone increasing, F is right continuous and monotone increasing.

If F is continuous at x , then for every $\varepsilon > 0$, there exist numbers $q^-, q^+ \in \mathbb{Q}$, $q^- < x < q^+$ with $\tilde{F}(q^-) \geq F(x) - \varepsilon$ and $\tilde{F}(q^+) \leq F(x) + \varepsilon$. By construction,

$$\limsup_{k \rightarrow \infty} F_{n_k}(x) \leq \lim_{k \rightarrow \infty} F_{n_k}(q^+) = \tilde{F}(q^+) \leq F(x) + \varepsilon$$

Hence $\limsup_{k \rightarrow \infty} F_{n_k}(x) \leq F(x)$. A similar argument for q^- yields $\liminf_{k \rightarrow \infty} F_{n_k}(x) \geq F(x)$.

Proof of Theorem 13.29(i) for the case $E = \mathbb{R}$ Assume $\mathcal{F}$ is tight and $(\mu_n)_{n \in \mathbb{N}}$ is a sequence in $\mathcal{F}$ with distribution functions $F_n : x \mapsto \mu_n((-\infty, x])$. By Helly's theorem, there is a monotone right continuous function $F : \mathbb{R} \rightarrow [0, 1]$ and a subsequence $(F_{n_k})_{k \in \mathbb{N}}$ of $(F_n)_{n \in \mathbb{N}}$ with $F_{n_k}(x) \xrightarrow{k \rightarrow \infty} F(x)$ at all points of continuity x of F . By Theorem 13.23, it is enough to show that $F(\infty) \geq \limsup_{k \rightarrow \infty} F_{n_k}(\infty)$.

As $\mathcal{F}$ is tight, for every $\varepsilon > 0$, there is a $K < \infty$ with $F_n(\infty) - F_n(x) < \varepsilon$ for all $n \in \mathbb{N}$ and $x > K$. If $x > K$ is a point of continuity of F , then $\limsup_{k \rightarrow \infty} F_{n_k}(\infty) \leq \limsup_{k \rightarrow \infty} F_{n_k}(x) + \varepsilon = F(x) + \varepsilon \leq F(\infty) + \varepsilon$.

We come to a first application of Prohorov's theorem. The full strength of that theorem will become manifest when suitable separating classes of functions are at our disposal. We come back to this point in more detail in Chapter 15.

Theorem 13.34 Let E be Polish and let $\mu, \mu_1, \mu_2, \dots \in \mathcal{M}_{\leq 1}(E)$. Then the following are equivalent.

- (i) $\mu = \text{w-lim}_{n \rightarrow \infty} \mu_n$.
- (ii) $(\mu_n)_{n \in \mathbb{N}}$ is tight, and there is a separating family $\mathcal{C} \subset C_b(E)$ such that

$$\int f d\mu = \lim_{n \rightarrow \infty} \int f d\mu_n \quad \text{for all } f \in \mathcal{C} \quad (13.8)$$

Proof "(i) $\implies$ (ii)" By the simple implication in Prohorov's theorem (Theorem 13.29(ii)), weak convergence implies tightness.

"(ii) $\implies$ (i)" Let $(\mu_n)_{n \in \mathbb{N}}$ be tight and let $\mathcal{C} \subset C_b(E)$ be a separating class with (13.8). Assume that $(\mu_n)_{n \in \mathbb{N}}$ does not converge weakly to μ . Then there are $\varepsilon > 0$, $f \in C_b(E)$ and $(n_k)_{k \in \mathbb{N}}$ with $n_k \uparrow \infty$ and such that

$$\left| \int f d\mu_{n_k} - \int f d\mu \right| > \varepsilon \quad \text{for all } k \in \mathbb{N}. \quad (13.9)$$

By Prohorov's theorem, there exists a $\nu \in \mathcal{M}_{\leq 1}(E)$ and a subsequence $(n'_k)_{k \in \mathbb{N}}$ of $(n_k)_{k \in \mathbb{N}}$ with $\mu_{n'_k} \rightarrow \nu$ weakly. Due to (13.9), we have $|\int f d\mu - \int f d\nu| \geq \varepsilon$; hence $\mu \neq \nu$. On the other hand,

$$\int h d\mu = \lim_{k \rightarrow \infty} \int h d\mu_{n'_k} = \int h d\nu \quad \text{for all } h \in \mathcal{C}$$

hence $\mu = \nu$. This contradicts the assumption and thus (i) holds.

We want to shed some more light on the connection between weak and vague convergence.

Theorem 13.35 Let E be a locally compact Polish space and let $\mu, \mu_1, \mu_2, \dots \in \mathcal{M}_f(E)$. Then the following are equivalent.

- (i) $\mu = w - \lim_{n \rightarrow \infty} \mu_n$.
- (ii) $\mu = v - \lim_{n \rightarrow \infty} \mu_n$ and $\mu(E) = \lim_{n \rightarrow \infty} \mu_n(E)$.
- (iii) $\mu = v - \lim_{n \rightarrow \infty} \mu_n$ and $\mu(E) \geq \limsup_{n \rightarrow \infty} \mu_n(E)$.
- (iv) $\mu = v - \lim_{n \rightarrow \infty} \mu_n$ and $\{\mu_n, n \in \mathbb{N}\}$ is tight.

Proof "(i) $\iff$ (ii) $\iff$ (iii)" This follows by the Portemanteau theorem.

"(ii) $\implies$ (iv)" It is enough to show that for any $\varepsilon > 0$, there is a compact set $K \subset E$ with $\limsup_{n \rightarrow \infty} \mu_n(E \setminus K) \leq \varepsilon$. As μ is regular (Theorem 13.6), there is a compact set $L \subset E$ with $\mu(E \setminus L) < \varepsilon$. Since E is locally compact, there exists a compact set $K \subset E$ with $K^\circ \supset L$ and a $\rho_{L,K} \in C_c(E)$ with $\mathbb{1}_L \leq \rho_{L,K}(x) \leq \mathbb{1}_K$. Therefore,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \mu_n(E \setminus K) &\leq \limsup_{n \rightarrow \infty} \left(\mu_n(E) - \int \rho_{L,K} d\mu_n \right) \\ &= \mu(E) - \int \rho_{L,K} d\mu \leq \mu(E \setminus L) < \varepsilon \end{aligned}$$

"(iv) $\implies$ (i)" Let $L \subset E$ be compact with $\mu_n(E \setminus L) \leq 1$ for all $n \in \mathbb{N}$. Let $\rho \in C_c(E)$ with $\rho \geq \mathbb{1}_L$. Since $\int \rho d\mu_n$ converges by assumption, we thus have

$$\sup_{n \in \mathbb{N}} \mu_n(E) \leq 1 + \sup_{n \in \mathbb{N}} \mu_n(L) \leq 1 + \sup_{n \in \mathbb{N}} \int \rho d\mu_n < \infty.$$

Hence also

$$C := \max(\mu(E), \sup \{\mu_n(E) : n \in \mathbb{N}\}) < \infty,$$

and we can pass to μ/C and μ_n/C . Thus, without loss of generality assume that all measures are in $\mathcal{M}_{\leq 1}(E)$. As $C_c(E)$ is a separating class for $\mathcal{M}_{\leq 1}(E)$ (see Theorem 13.11), (i) follows by Theorem 13.34.

Proof of Prohorov's Theorem, Part (i), General Case There are two main routes for proving Prohorov's theorem in the general situation. One possibility is to show the claim first for measures on $\mathbb{R}^d$. (We have done this already for $d = 1$, see Exercise 13.3.4 for $d \geq 2$.) In a second step, the statement is lifted to sequence spaces $\mathbb{R}^{\mathbb{N}}$. Finally, in the third step, an embedding of E into $\mathbb{R}^{\mathbb{N}}$ is constructed. For a detailed description, see [12] or [83].

Here we follow the alternative route as described in [13] (and later [14]) or [44]. The main point of this proof consists in finding a candidate for a weak limit point for the family $\mathcal{F}$. This

candidate will be constructed first as a content on a countable class of sets. From this an outer measure will be derived. Finally, we show that closed sets are measurable with respect to this outer measure. As you see, the argument follows a pattern similar to the proof of Carathéodory's theorem.

Let (E, d) be a metric space and let $\mathcal{F} \subset \mathcal{M}_{\leq 1}(E)$ be tight. Then there exists an increasing sequence $K_1 \subset K_2 \subset K_3 \subset \dots$ of compact sets in E such that $\mu(K_n^c) < \frac{1}{n}$ for all $\mu \in \mathcal{F}$ and all $n \in \mathbb{N}$. Define $E' := \bigcup_{n=1}^{\infty} K_n$. Then E' is a σ -compact metric space and therefore in particular, separable. By construction, $\mu(E \setminus E') = 0$ for all $\mu \in \mathcal{F}$. Thus, any μ can be regarded as a measure on E' . Without loss of generality, we may hence assume that E is σ -compact and thus separable. Hence there exists a countable base $\mathcal{U}$ of the topology $\tau|_E$ on E ; that is, a countable set $\mathcal{U}$ of open sets such that $A = \bigcup_{U \in \mathcal{U}, U \subset A} U$ for any open $A \subset E$. Define

$$\mathcal{C}' := \{\bar{U} \cap K_n : U \in \mathcal{U}, n \in \mathbb{N}\}$$

and

$$\mathcal{C} := \left\{ \bigcup_{n=1}^N C_n : N \in \mathbb{N} \text{ and } C_1, \dots, C_N \in \mathcal{C}' \right\}$$

Clearly, $\mathcal{C}$ is a countable set of compact sets in E , and $\mathcal{C}$ is stable under formation of unions. Any K_n possesses a finite covering with sets from $\mathcal{U}$; hence $K_n \in \mathcal{C}$.

Now let $(\mu_n)_{n \in \mathbb{N}}$ be a sequence in $\mathcal{F}$. By virtue of the diagonal sequence argument (see the proof of Helly's theorem, Theorem 13.33), we can find a subsequence $(\mu_{n_k})_{k \in \mathbb{N}}$ such that for all $C \in \mathcal{C}$, there exists the limit

$$\alpha(C) := \lim_{k \rightarrow \infty} \mu_{n_k}(C). \quad (13.10)$$

Assume that we can show that there is a measure μ on the Borel σ -algebra $\mathcal{E}$ of E such that

$$\mu(A) = \sup\{\alpha(C) : C \in \mathcal{C} \text{ with } C \subset A\} \quad \text{for all } A \subset E \text{ open.} \quad (13.11)$$

Then

$$\begin{aligned} \mu(E) &\geq \sup_{n \in \mathbb{N}} \alpha(K_n) = \sup_{n \in \mathbb{N}} \lim_{k \rightarrow \infty} \mu_{n_k}(K_n) \\ &\geq \sup_{n \in \mathbb{N}} \limsup_{k \rightarrow \infty} \left(\mu_{n_k}(E) - \frac{1}{n} \right) \\ &= \limsup_{k \rightarrow \infty} \mu_{n_k}(E) \end{aligned}$$

Furthermore, for open A and for $C \in \mathcal{C}$ with $C \subset A$,

$$\alpha(C) = \lim_{k \rightarrow \infty} \mu_{n_k}(C) \leq \liminf_{k \rightarrow \infty} \mu_{n_k}(A),$$

hence $\mu(A) \leq \liminf_{k \rightarrow \infty} \mu_{n_k}(A)$. By the Portemanteau theorem (Theorem 13.16), $\mu = w - \lim_{k \rightarrow \infty} \mu_{n_k}$; hence $\mathcal{F}$ is recognized as weakly relatively sequentially compact. It remains to show that there exists a measure μ on $(E, \mathcal{E})$ that satisfies (13.11).

Clearly, the set function α on $\mathcal{C}$ is monotone, additive and subadditive:

$$\begin{aligned} \alpha(C_1) &\leq \alpha(C_2), & \text{if } C_1 \subset C_2, \\ \alpha(C_1 \cup C_2) &= \alpha(C_1) + \alpha(C_2), & \text{if } C_1 \cap C_2 = \emptyset, \\ \alpha(C_1 \cup C_2) &\leq \alpha(C_1) + \alpha(C_2). \end{aligned} \quad (13.12)$$

We define

$$\beta(A) := \sup\{\alpha(C) : C \in \mathcal{C} \text{ with } C \subset A\} \quad \text{for } A \subset E \text{ open}$$

and


$$\mu^*(G) := \inf\{\beta(A) : A \supset G \text{ is open}\} \quad \text{for } G \in 2^E.$$

Manifestly, $\beta(A) = \mu^*(A)$ for any open A . It is enough to show (Steps 1-3 below) that μ^* is an outer measure (see Definition 1.46) and that (Step 4) the σ -algebra of μ^* -measurable sets (see Definition 1.48 and Lemma 1.52) contains the closed sets and thus $\mathcal{E}$. Indeed, Lemma 1.52 would then imply that μ^* is a measure on the σ -algebra of μ^* -measurable sets and the restricted measure $\mu := \mu^*|_{\mathcal{E}}$ fulfills $\mu(A) = \mu^*(A) = \beta(A)$ for all open A . Hence Eq. (13.11) holds.

Evidently, $\mu^*(\emptyset) = 0$ and μ^* is monotone. In order to show that μ^* is an outer measure, it only remains to check that μ^* is σ -subadditive.

Step 1 (Finite subadditivity of β) Let $A_1, A_2 \subset E$ be open and let $C \in \mathcal{C}$ with $C \subset A_1 \cup A_2$. Let $n \in \mathbb{N}$ with $C \subset K_n$. Define two sets

$$\begin{aligned} B_1 &:= \{x \in C : d(x, A_1^c) \geq d(x, A_2^c)\}, \\ B_2 &:= \{x \in C : d(x, A_1^c) \leq d(x, A_2^c)\}. \end{aligned}$$



Evidently, $B_1 \subset A_1$ and $B_2 \subset A_2$. As $x \mapsto d(x, A_i^c)$ is continuous for $i = 1, 2$, the closed subsets B_1 and B_2 of C are compact. Hence $d(B_1, A_1^c) > 0$. Thus there exists an open set D_1 with $B_1 \subset D_1 \subset \bar{D}_1 \subset A_1$. (One could choose D_1 as the union of the sets of a finite covering of B_1 with balls of radius $d(B_1, A_1^c)/2$. These balls, as well as their closures, are subsets of A_1 .) Let $\mathcal{U}_{D_1} := \{U \in \mathcal{U} : U \subset D_1\}$. Then $B_1 \subset D_1 = \bigcup_{U \in \mathcal{U}_{D_1}} U$. Now choose a finite subcovering $\{U_1, \dots, U_N\} \subset \mathcal{U}_{D_1}$ of B_1 and define $C_1 := \bigcup_{i=1}^N \bar{U}_i \cap K_n$. Then $B_1 \subset C_1 \subset A_1$ and $C_1 \in \mathcal{C}$. Similarly, choose $C_2 \in \mathcal{C}$ with $B_2 \subset C_2 \subset A_2$. Thus

$$\alpha(C) \leq \alpha(C_1 \cup C_2) \leq \alpha(C_1) + \alpha(C_2) \leq \beta(A_1) + \beta(A_2).$$

Hence also

$$\begin{aligned} \beta(A_1 \cup A_2) &= \sup\{\alpha(C) : C \in \mathcal{C} \text{ with } C \subset A_1 \cup A_2\} \\ &\leq \beta(A_1) + \beta(A_2). \end{aligned}$$

Step 2 (σ -subadditivity of β) Let $A_1, A_2, \dots$ be open sets and let $C \in \mathcal{C}$ with $C \subset \bigcup_{i=1}^{\infty} A_i$. As C is compact, there exists an $n \in \mathbb{N}$ with $C \subset \bigcup_{i=1}^n A_i$. As shown above, β is subadditive; thus

$$\alpha(C) \leq \beta\left(\bigcup_{i=1}^n A_i\right) \leq \sum_{i=1}^{\infty} \beta(A_i).$$

Taking the supremum over such C yields

$$\beta\left(\bigcup_{i=1}^{\infty} A_i\right) = \sup\left\{\alpha(C) : C \in \mathcal{C} \text{ with } C \subset \bigcup_{i=1}^{\infty} A_i\right\} \leq \sum_{i=1}^{\infty} \beta(A_i).$$

Step 3 (σ -subadditivity of μ^*) Let $G_1, G_2, \dots \in 2^E$. Let $\varepsilon > 0$. For any $n \in \mathbb{N}$ choose an open set $A_n \supset G_n$ with $\beta(A_n) < \mu^*(G_n) + \varepsilon/2^n$. By the σ -subadditivity of β ,

$$\mu^*\left(\bigcup_{n=1}^{\infty} G_n\right) \leq \beta\left(\bigcup_{n=1}^{\infty} A_n\right) \leq \sum_{n=1}^{\infty} \beta(A_n) \leq \varepsilon + \sum_{n=1}^{\infty} \mu^*(G_n).$$

Letting $\varepsilon \downarrow 0$ yields $\mu^*\left(\bigcup_{n=1}^{\infty} G_n\right) \leq \sum_{n=1}^{\infty} \mu^*(G_n)$. Hence μ^* is an outer measure.

Step 4 (Closed sets are μ^* -measurable) By Lemma 1.49, a set $B \subset E$ is μ^* measurable if and only if

$$\mu^*(B \cap G) + \mu^*(B^c \cap G) \leq \mu^*(G) \quad \text{for all } G \in 2^E.$$

Taking the infimum over all open sets $A \supset G$, it is enough to show that for every open B and every open $A \subset E$,

$$\mu^*(B \cap A) + \mu^*(B^c \cap A) \leq \beta(A). \quad (13.13)$$

Let $\varepsilon > 0$. Choose $C_1 \in \mathcal{C}$ with $C_1 \subset A \cap B^c$ and $\alpha(C_1) > \beta(A \cap B^c) - \varepsilon$. Further, let $C_2 \in \mathcal{C}$ with $C_2 \subset A \cap C_1^c$ and $\alpha(C_2) > \beta(A \cap C_1^c) - \varepsilon$. Since $C_1 \cap C_2 = \emptyset$ and $C_1 \cup C_2 \subset A$, we get

$$\begin{aligned} \beta(A) &\geq \alpha(C_1 \cup C_2) \\ &= \alpha(C_1) + \alpha(C_2) \geq \beta(A \cap B^c) + \beta(A \cap C_1^c) - 2\varepsilon \\ &\geq \mu^*(A \cap B^c) + \mu^*(A \cap B) - 2\varepsilon. \end{aligned}$$

Letting $\varepsilon \rightarrow 0$, we get (13.13). This completes the proof of Prohorov's theorem.

Exercise 13.3.1 Show that a family $\mathcal{F} \subset \mathcal{M}_f(\mathbb{R})$ is tight if and only if there exists a measurable map $f : \mathbb{R} \rightarrow [0, \infty)$ such that $f(x) \rightarrow \infty$ for $|x| \rightarrow \infty$ and $\sup_{\mu \in \mathcal{F}} \int f d\mu < \infty$.

Exercise 13.3.2 Let $L \subset \mathbb{R} \times (0, \infty)$ and let $\mathcal{F} = \{\mathcal{N}_{\mu, \sigma^2} : (\mu, \sigma^2) \in L\}$ be a family of normal distributions with parameters in L . Show that $\mathcal{F}$ is tight if and only if L is bounded.

Exercise 13.3.3 If P is a probability measure on $[0, \infty)$ with $m_P := \int xP(dx) \in (0, \infty)$, then we define the size-biased distribution $\hat{P}$ on $[0, \infty)$ by

$$\hat{P}(A) = m_P^{-1} \int_A xP(dx) \quad (13.14)$$

Now let $(X_i)_{i \in I}$ be a family of random variables on $[0, \infty)$ with $\mathbf{E}[X_i] = 1$. Show that $(\widehat{\mathbf{P}}_{X_i})_{i \in I}$ is tight if and only if $(X_i)_{i \in I}$ is uniformly integrable.

Exercise 13.3.4 (Helly's theorem in $\mathbb{R}^d$) Let $x = (x^1, \dots, x^d) \in \mathbb{R}^d$ and $y = (y^1, \dots, y^d) \in \mathbb{R}^d$. Recall the notation $x \leq y$ if $x^i \leq y^i$ for all $i = 1, \dots, d$. A map $F : \mathbb{R}^d \rightarrow \mathbb{R}$ is called monotone increasing if $F(x) \leq F(y)$ whenever $x \leq y$. F is called right continuous if $F(x) = \lim_{n \rightarrow \infty} F(x_n)$ for all $x \in \mathbb{R}^d$ and every sequence $(x_n)_{n \in \mathbb{N}}$ in $\mathbb{R}^d$ with $x_1 \geq x_2 \geq x_3 \geq \dots$ and $x = \lim_{n \rightarrow \infty} x_n$. By V_d denote the set of monotone increasing, bounded right continuous functions on $\mathbb{R}^d$.

(i) Show the validity of Helly's theorem with V replaced by V_d .

(ii) Conclude that Prohorov's theorem holds for $E = \mathbb{R}^d$.

5.13.4 13.4 Application: a Fresh Look at de Finetti's Theorem

(After an idea of Götz Kersting.) Let E be a Polish space and let $X_1, X_2, \dots$ be an exchangeable sequence of random variables with values in E . As an alternative to the backwards martingale argument of Section 12.3, here we give a different proof of de Finetti's theorem (Theorem 12.26). Recall that de Finetti's theorem states that there exists a random probability measure Ξ on E such that, given Ξ , the random variables $X_1, X_2, \dots$ are independent and Ξ -distributed. For $x = (x_1, x_2, \dots) \in E^{\mathbb{N}}$, let $\xi_n(x) := \frac{1}{n} \sum_{l=1}^n \delta_{x_l}$ be the empirical distribution of $x_1, \dots, x_n$. Let

$$\mu_{n,k}(x) := \xi_n(x)^{\otimes k} = n^{-k} \sum_{i_1, \dots, i_k=1}^n \delta_{(x_{i_1}, \dots, x_{i_k})}$$

be the distribution on E^k that describes k -fold independent sampling with replacement (respecting the order) from $(x_1, \dots, x_n)$. Let

$$v_{n,k}(x) := \frac{(n-k)!}{n!} \sum_{\substack{i_1, \dots, i_k=1 \\ \#\{i_1, \dots, i_k\}=k}}^n \delta_{(x_{i_1}, \dots, x_{i_k})}$$

be the distribution on E^k that describes k -fold independent sampling without replacement (respecting the order) from $(x_1, \dots, x_n)$. For all $x \in E^{\mathbb{N}}$,

$$\|\mu_{n,k}(x) - v_{n,k}(x)\|_{TV} \leq R_{n,k} := \frac{k(k-1)}{n}.$$

Indeed, the probability $p_{n,k}$ that we do not see any ball twice when drawing k balls (with replacement) from n different balls is

$$p_{n,k} = \prod_{l=1}^{k-1} (1 - l/n)$$

and thus $R_{n,k} \geq 2(1 - p_{n,k})$. We therefore obtain the rather intuitive statement that as $n \rightarrow \infty$ the distributions of k -samples with replacement and without replacement, respectively, become the same:

$$\lim_{n \rightarrow \infty} \sup_{x \in E^{\mathbb{N}}} \|\mu_{n,k}(x) - v_{n,k}(x)\|_{TV} = 0$$

Now let $f_1, \dots, f_k \in C_b(E)$ and $F(x_1, \dots, x_k) := f_1(x_1) \dots f_k(x_k)$. As the sequence $X_1, X_2, \dots$ is exchangeable, for any choice of pairwise distinct numbers $1 \leq i_1, \dots, i_k \leq n$,

$$\mathbf{E}[F(X_1, \dots, X_k)] = \mathbf{E}[F(X_{i_1}, \dots, X_{i_k})].$$

Averaging over all choices $i_1, \dots, i_k$, we get

$$\mathbf{E}[f_1(X_1) \dots f_k(X_k)] = \mathbf{E}[F(X_1, \dots, X_k)] = \mathbf{E}\left[\int F dv_{n,k}(X)\right].$$

Hence

$$\begin{aligned} & \left| \mathbf{E}[f_1(X_1) \dots f_k(X_k)] - \mathbf{E}\left[\int f_1 d\xi_n(X) \dots \int f_k d\xi_n(X)\right] \right| \\ &= \left| \mathbf{E}\left[\int F dv_{n,k}(X)\right] - \mathbf{E}\left[\int F d\mu_{n,k}(X)\right] \right| \\ &\leq \|F\|_{\infty} R_{n,k} \xrightarrow{n \rightarrow \infty} 0 \end{aligned}$$

We will exploit the following criterion for tightness of subsets of $\mathcal{M}_1(\mathcal{M}_1(E))$.

Exercise 13.4.1 Show that a subset $\mathcal{K} \subset \mathcal{M}_1(\mathcal{M}_1(E))$ is tight if and only if, for any $\varepsilon > 0$, there exists a compact set $K \subset E$ with the property

$$\tilde{\mu}(\{\mu \in \mathcal{M}_1(E) : \mu(K^c) > \varepsilon\}) < \varepsilon \quad \text{for all } \tilde{\mu} \in \mathcal{K}.$$

Since E is Polish, $\mathbf{P}_{X_1}$ is tight. Hence, for any $\varepsilon > 0$, there exists a compact set $K \subset E$ with $\mathbf{P}[X_1 \in K^c] < \varepsilon^2$. Therefore,

$$\mathbf{P}[\xi_n(X)(K^c) > \varepsilon] \leq \varepsilon^{-1} \mathbf{E}[\xi_n(X)(K^c)] = \varepsilon^{-1} \mathbf{P}[X_1 \in K^c] \leq \varepsilon.$$

Hence the family $(\mathbf{P}_{\xi_n(X)})_{n \in \mathbb{N}}$ is tight. Let Ξ_∞ be a random variable (with values in $\mathcal{M}_1(E)$) such that $\mathbf{P}_{\Xi_\infty} = \text{w-lim}_{l \rightarrow \infty} \mathbf{P}_{\xi_{n_l}(X)}$ for a suitable subsequence $(n_l)_{l \in \mathbb{N}}$. The map $\xi \mapsto \int F d\xi = \int f_1 d\xi \dots \int f_k d\xi$ is bounded and (as a product of continuous maps) is continuous with respect to the topology of weak convergence on $\mathcal{M}_1(E)$; hence it is in $C_b(\mathcal{M}_1(E))$. Thus

$$\begin{aligned} \mathbf{E} \left[\int F d\Xi_\infty^{\otimes k} \right] &= \lim_{l \rightarrow \infty} \mathbf{E} \left[\int f_1 d\xi_{n_l}(X) \dots \int f_k d\xi_{n_l}(X) \right] \\ &= \mathbf{E}[f_1(X_1) \dots f_k(X_k)] \end{aligned}$$

Note that the limit does not depend on the choice of the subsequence and is thus unique. Summarising, we have

$$\mathbf{E}[f_1(X_1) \dots f_k(X_k)] = \mathbf{E} \left[\int f_1 d\Xi_\infty \dots \int f_k d\Xi_\infty \right]$$

Since the distribution of $(X_1, \dots, X_k)$ is uniquely determined by integrals of the above type, we conclude that $\mathbf{P}_{(X_1, \dots, X_k)} = \mathbf{P}_{\Xi_\infty^{\otimes k}}$. In other words, $(X_1, \dots, X_k) \stackrel{\mathcal{D}}{=} (Y_1, \dots, Y_k)$, where, given Ξ_∞ , the random variables $Y_1, \dots, Y_k$ are independent with distribution Ξ_∞ .

Exercise 13.4.2 Show that a family $(X_n)_{n \in \mathbb{N}}$ of random variables is exchangeable if and only if, for every choice of natural numbers $1 \leq n_1 < n_2 < n_3 < \dots$, we have

$$(X_1, X_2, \dots) \stackrel{\mathcal{D}}{=} (X_{n_1}, X_{n_2}, \dots)$$

Warning: One of the implications is rather difficult to show.

5.14 14 Probability Measures on Product Spaces

As a motivation, consider the following example. Let X be a random variable that is uniformly distributed on $[0, 1]$. As soon as we know the value of X , we toss n times a coin that has probability X for a success. Denote the results by $Y_1, \dots, Y_n$.

How can we construct a probability space on which all these random variables are defined? One possibility is to construct $n+1$ independent random variables $Z_0, \dots, Z_n$ that are uniformly distributed on $[0, 1]$ (see, e.g., Corollary 2.23 for the construction). Then define $X = Z_0$ and

$$Y_k = \begin{cases} 1, & \text{if } Z_k < X \\ 0, & \text{if } Z_k \geq X \end{cases}$$

Intuitively, this fits well with our idea that the $Y_1, \dots, Y_n$ are independent as soon as we know X and record a success with probability X .

In the above description, we have constructed by hand a two-stage experiment. At the first stage, we determine the value of X . At the second stage, depending on the value of X , the values of $Y = (Y_1, \dots, Y_n)$ are determined. Clearly, this construction makes use of the specific structure

of the problem. However, we now want to develop a systematic framework for the description and construction of multistage experiments. In contrast to Chapter 2, here the random variables need not be independent. In addition, we also want to construct systematically infinite families of random variables with given (joint) distributions.

In the first section, we start with products of measurable spaces. Then we come to finite products of measure spaces and product measures with transition kernels. Finally, we consider infinite products of probability spaces. The main result is Kolmogorov's extension theorem.

5.14.1 14.1 Product Spaces

Definition 14.1 (Product space) Let $(\Omega_i, i \in I)$ be an arbitrary family of sets. Denote by $\Omega = \prod_{i \in I} \Omega_i$ the set of maps $\omega : I \rightarrow \prod_{i \in I} \Omega_i$ such that $\omega(i) \in \Omega_i$ for all $i \in I$. Ω is called the product of the spaces $(\Omega_i, i \in I)$, or briefly the product space. If, in particular, all the Ω_i are equal, say $\Omega_i = \Omega_0$, then we write $\Omega = \times_{i \in I} \Omega_i = \Omega_0^I$.

Example 14.2

(i) If $\Omega_1 = \{1, \dots, 6\}$ and $\Omega_2 = \{1, 2, 3\}$, then

$$\Omega_1 \times \Omega_2 = \{\omega = (\omega_1, \omega_2) : \omega_1 \in \{1, \dots, 6\}, \omega_2 \in \{1, 2, 3\}\}.$$

(ii) If $\Omega_0 = \mathbb{R}$ and $I = \{1, 2, 3\}$, then $\mathbb{R}^{\{1,2,3\}}$ is isomorphic to the customary $\mathbb{R}^3$.

(iii) If $\Omega_0 = \mathbb{R}$ and $I = \mathbb{N}$, then $\mathbb{R}^{\mathbb{N}}$ is the space of sequences $(\omega(n), n \in \mathbb{N})$ in $\mathbb{R}$.

(iv) If $I = \mathbb{R}$ and $\Omega_0 = \mathbb{R}$, then $\mathbb{R}^{\mathbb{R}}$ is the set of maps $\mathbb{R} \rightarrow \mathbb{R}$.

Definition 14.3 (Coordinate maps) If $i \in I$, then $X_i : \Omega \rightarrow \Omega_i, \omega \mapsto \omega(i)$ denotes the i th coordinate map. More generally, for $J \subset J' \subset I$, the restricted map

$$X_J^{J'} : \prod_{j \in J'} \Omega_j \longrightarrow \prod_{j \in J} \Omega_j, \quad \omega' \mapsto \omega'|_J \quad (14.1)$$

is called the canonical projection. In particular, we write $X_J := X_J^I$.

Definition 14.4 (Product- σ -algebra) Let $(\Omega_i, \mathcal{A}_i), i \in I$, be measurable spaces. The product- σ -algebra

$$\mathcal{A} = \bigotimes_{i \in I} \mathcal{A}_i$$

is the smallest σ -algebra on Ω such that for every $i \in I$, the coordinate map X_i is measurable with respect to $\mathcal{A} - \mathcal{A}_i$; that is,

$$\mathcal{A} = \sigma(X_i, i \in I) := \sigma(X_i^{-1}(\mathcal{A}_i), i \in I).$$

If $(\Omega_i, \mathcal{A}_i) = (\Omega_0, \mathcal{A}_0)$ for all $i \in I$, then we also write $\mathcal{A} = \mathcal{A}_0^{\otimes I}$.

For $J \subset I$, let $\Omega_J := \prod_{j \in J} \Omega_j$ and $\mathcal{A}_J = \bigotimes_{j \in J} \mathcal{A}_j$.

Remark 14.5 The concept of the product- σ -algebra is similar to that of the product topology: If $(\Omega_i, \tau_i), i \in I$ are topological spaces, then the product topology τ on $\Omega = \prod_{i \in I} \Omega_i$ is the coarsest topology with respect to which all coordinate maps $X_i : \Omega \rightarrow \Omega_i$ are continuous.

Definition 14.6 Let $I \neq \emptyset$ be an arbitrary index set, let $(E, \mathcal{E})$ be a measurable space, let $(\Omega, \mathcal{A}) = (E^I, \mathcal{E}^{\otimes I})$ and let $X_t : \Omega \rightarrow E$ be the coordinate map for every $t \in I$. Then the family $(X_t)_{t \in I}$ is called the canonical process on $(\Omega, \mathcal{A})$.

Lemma 14.7 Let $\emptyset \neq J \subset I$. Then X_J^I is measurable with respect to $\mathcal{A}_I - \mathcal{A}_J$.

Proof For any $j \in J$, $X_j = X_j^J \circ X_J^I$ is measurable with respect to $\mathcal{A} - \mathcal{A}_j$. Thus, by Corollary 1.82, X_J^I is measurable.

Theorem 14.8 Let I be countable, and for every $i \in I$, let (Ω_i, τ_i) be Polish with Borel σ -algebra $\mathcal{B}_i = \sigma(\tau_i)$. Let τ be the product topology on $\Omega = \prod_{i \in I} \Omega_i$ and $\mathcal{B} = \sigma(\tau)$.

Then (Ω, τ) is Polish and $\mathcal{B} = \bigotimes_{i \in I} \mathcal{B}_i$. In particular, $\mathcal{B}(\mathbb{R}^d) = \mathcal{B}(\mathbb{R})^{\otimes d}$ for $d \in \mathbb{N}$.

Proof Without loss of generality, assume $I = \mathbb{N}$. For $i \in \mathbb{N}$, let d_i be a complete metric that induces τ_i . It is easy to check that

$$d(\omega, \omega') := \sum_{i=1}^{\infty} 2^{-i} \frac{d_i(\omega(i), \omega'(i))}{1 + d_i(\omega(i), \omega'(i))} \quad (14.2)$$

is a complete metric on Ω that induces τ .

Now for any $i \in \mathbb{N}$, let $D_i \subset \Omega_i$ be a countable dense subset and let $y_i \in D_i$ be an arbitrary point. It is easy to see that the set

$$D = \left\{ x \in \prod_{i \in \mathbb{N}} D_i : x_i \neq y_i \text{ only finitely often} \right\}$$

is a countable dense subset of Ω . Hence Ω is separable and thus Polish.

Now, for any $i \in I$, let $\beta_i = \{B_\varepsilon(x_i) : x_i \in D_i, \varepsilon \in \mathbb{Q}^+\}$ be a countable base of the topology of Ω_i consisting of ε -balls. Define

$$\beta := \bigcup_{N=1}^{\infty} \left\{ \bigcap_{i=1}^N X_i^{-1}(B_i) : B_1 \in \beta_1, \dots, B_N \in \beta_N \right\}.$$

Then β is a countable base of the topology τ ; hence any open set $A \subset \Omega$ is a (countable) union of sets in $\beta \subset \bigotimes_{i \in \mathbb{N}} \mathcal{B}_i$. Hence $\tau \subset \bigotimes_{i \in \mathbb{N}} \mathcal{B}_i$ and thus $\mathcal{B} \subset \bigotimes_{i \in \mathbb{N}} \mathcal{B}_i$.

On the other hand, each X_i is continuous and thus measurable with respect to $\mathcal{B} - \mathcal{B}_i$. Therefore, $\mathcal{B} \supset \bigotimes_{i \in \mathbb{N}} \mathcal{B}_i$.

Definition 14.9 (Cylinder sets) For any $i \in I$, let $\mathcal{E}_i \subset \mathcal{A}_i$ be a subclass of the class of measurable sets.

For any $A \in \mathcal{A}_J$, $X_J^{-1}(A) \subset \Omega$ is called a cylinder set with base J . The set of such cylinder sets is denoted by $\mathcal{Z}_J$. In particular, if $A = \times_{j \in J} A_j$ for certain $A_j \in \mathcal{A}_j$, then $X_J^{-1}(A)$ is called a rectangular cylinder with base J . The set of such rectangular cylinders will be denoted by $\mathcal{Z}_J^R$. The set of such rectangular cylinders for which in addition $A_j \in \mathcal{E}_j$ for all $j \in J$ holds will be denoted by $\mathcal{Z}_J^{\mathcal{E}, R}$.

Write

$$\mathcal{Z} = \bigcup_{J \subset I \text{ finite}} \mathcal{Z}_J \quad (14.3)$$

and similarly define $\mathcal{Z}^R$ and $\mathcal{Z}^{\mathcal{E}, R}$. Further, define

$$\mathcal{Z}_*^R = \bigcup_{N=1}^{\infty} \left\{ \bigcup_{n=1}^N A_n : A_1, \dots, A_N \in \mathcal{Z}^R \right\}$$

and similarly $\mathcal{Z}_*^{\mathcal{E}, R}$.

Remark 14.10 Every $\mathcal{Z}_J$ is a σ -algebra, and $\mathcal{Z}$ and $\mathcal{Z}_*^R$ are algebras. Furthermore, $\bigotimes_{i \in I} \mathcal{A}_i = \sigma(\mathcal{Z})$.

Lemma 14.11 If every $\mathcal{E}_i$ is a π -system, then $\mathcal{Z}^{\mathcal{E}, R}$ is a π -system.

Proof This is left as an exercise.

Theorem 14.12 For any $i \in I$, let $\mathcal{E}_i \subset \mathcal{A}_i$ be a generator of $\mathcal{A}_i$.

- (i) $\bigotimes_{j \in J} \mathcal{A}_j = \sigma(\times_{j \in J} E_j : E_j \in \mathcal{E}_j)$ for every finite $J \subset I$.
- (ii) $\bigotimes_{i \in I} \mathcal{A}_i = \sigma(\mathcal{Z}^R) = \sigma(\mathcal{Z}^{\mathcal{E}, R})$.
- (iii) Let μ be a σ -finite measure on $\mathcal{A}$, and assume every $\mathcal{E}_i$ is also a π -system. Furthermore,

assume there is a sequence $(E_n)_{n \in \mathbb{N}}$ in $\mathcal{Z}^{\mathcal{E}, R}$ with $E_n \uparrow \Omega$ and $\mu(E_n) < \infty$ for all $n \in \mathbb{N}$ (this condition is satisfied, for example, if μ is finite and $\Omega_i \in \mathcal{E}_i$ for all $i \in I$). Then μ is uniquely determined by the values $\mu(A)$ for all $A \in \mathcal{Z}^{\mathcal{E}, R}$.

Proof (i) Let $\mathcal{A}'_J = \sigma(X_{j \in J} E_j : E_j \in \mathcal{E}_j \text{ for every } j \in J)$. Note that

$$\times_{j \in J} E_j = \bigcap_{j \in J} (X_j^J)^{-1}(E_j) \in \mathcal{A}_J,$$

hence $\mathcal{A}'_J \subset \mathcal{A}_J$. On the other hand, $(X_j^J)^{-1}(E_j) \in \mathcal{A}'_J$ for all $j \in J$ and $E_j \in \mathcal{E}_j$. Since $\mathcal{E}_i$ is a generator of $\mathcal{A}_i$, we have $(X_j^J)^{-1}(A_j) \in \mathcal{A}'_J$ for all $A_j \in \mathcal{A}_j$, and hence $\mathcal{A}_J \subset \mathcal{A}'_J$.

(ii) Evidently, $\mathcal{Z}^{\mathcal{E}, R} \subset \mathcal{Z}^R \subset \mathcal{A}$; hence also $\sigma(\mathcal{Z}^{\mathcal{E}, R}) \subset \sigma(\mathcal{Z}^R) \subset \mathcal{A}$. By Theorem 1.81, we have $\sigma(\mathcal{Z}_{\{i\}}^{\mathcal{E}, R}) = \sigma(X_i)$ for all $i \in I$; hence $\sigma(X_i) \subset \sigma(\mathcal{Z}^{\mathcal{E}, R})$. Therefore, $\mathcal{A}_I \subset \sigma(\mathcal{Z}^{\mathcal{E}, R})$.

(iii) By (ii) and Lemma 14.11, $\mathcal{Z}^{\mathcal{E}, R}$ is a π -system that generates $\mathcal{A}$. Hence, the claim follows by Lemma 1.42.

Exercise 14.1.1 Show that

$$\bigotimes_{i \in I} \mathcal{A}_i = \bigcup_{J \subset I \text{ countable}} \mathcal{Z}_J \quad (14.4)$$

Hint: Show that the right-hand side is a σ -algebra.

5.14.2 14.2 Finite Products and Transition Kernels

Consider now the situation of finitely many σ -finite measure spaces $(\Omega_i, \mathcal{A}_i, \mu_i)$, $i = 1, \dots, n$, where $n \in \mathbb{N}$.

Lemma 14.13 Let $A \in \mathcal{A}_1 \otimes \mathcal{A}_2$ and let $f : \Omega_1 \times \Omega_2 \rightarrow \overline{\mathbb{R}}$ be an $\mathcal{A}_1 \otimes \mathcal{A}_2$ measurable map. Then, for all $\tilde{\omega}_1 \in \Omega_1$ and $\tilde{\omega}_2 \in \Omega_2$,

$$\begin{aligned} A_{\tilde{\omega}_1} &:= \{\omega_2 \in \Omega_2 : (\tilde{\omega}_1, \omega_2) \in A\} \in \mathcal{A}_2 \\ A_{\tilde{\omega}_2} &:= \{\omega_1 \in \Omega_1 : (\omega_1, \tilde{\omega}_2) \in A\} \in \mathcal{A}_1 \\ f_{\tilde{\omega}_1} : \Omega_2 &\rightarrow \overline{\mathbb{R}}, \quad \omega_2 \mapsto f(\tilde{\omega}_1, \omega_2) \quad \text{is } \mathcal{A}_2\text{-measurable,} \\ f_{\tilde{\omega}_2} : \Omega_1 &\rightarrow \overline{\mathbb{R}}, \quad \omega_1 \mapsto f(\omega_1, \tilde{\omega}_2) \quad \text{is } \mathcal{A}_1\text{-measurable.} \end{aligned}$$

Proof For $\tilde{\omega}_1$, define the embedding map $i : \Omega_2 \rightarrow \Omega_1 \times \Omega_2$ by $i(\omega_2) = (\tilde{\omega}_1, \omega_2)$. Note that $X_1 \circ i$ is constantly $\tilde{\omega}_1$ (and hence $\mathcal{A}_1$ -measurable), and $X_2 \circ i = \text{id}_{\Omega_2}$ (and hence $\mathcal{A}_2$ -measurable). Thus, by Corollary 1.82, the map i is measurable with respect to $\mathcal{A}_2 - (\mathcal{A}_1 \otimes \mathcal{A}_2)$. Hence $A_{\tilde{\omega}_1} = i^{-1}(A) \in \mathcal{A}_2$ and $f_{\tilde{\omega}_1} = f \circ i$ is measurable with respect to $\mathcal{A}_2$.

The following theorem generalizes Theorem 1.61.

Theorem 14.14 (Finite product measures) There exists a unique σ -finite measure μ on $\mathcal{A} := \bigotimes_{i=1}^n \mathcal{A}_i$ such that

$$\mu(A_1 \times \dots \times A_n) = \prod_{i=1}^n \mu_i(A_i) \quad \text{for } A_i \in \mathcal{A}_i, i = 1, \dots, n \quad (14.5)$$

$\bigotimes_{i=1}^n \mu_i := \mu_1 \otimes \dots \otimes \mu_n := \mu$ is called the product measure of the μ_i . If all spaces involved equal $(\Omega_0, \mathcal{A}_0, \mu_0)$, then we write $\mu_0^{\otimes n} := \bigotimes_{i=1}^n \mu_0$.

Proof Let $\tilde{\mu}$ be the restriction of μ to $\mathcal{Z}^R$. Evidently, $\tilde{\mu}(\emptyset) = 0$, and it is simple to check that $\tilde{\mu}$ is σ -finite. Let $A^1, A^2, \dots \in \mathcal{Z}^R$ be pairwise disjoint and let $A \in \mathcal{Z}^R$ with $A \subset \bigcup_{k=1}^{\infty} A^k$. Then, by the monotone convergence theorem,

$$\begin{aligned}\tilde{\mu}(A) &= \int \mu_1(d\omega_1) \dots \int \mu_n(d\omega_n) \mathbb{1}_A((\omega_1, \dots, \omega_n)) \\ &\leq \int \mu_1(d\omega_1) \dots \int \mu_n(d\omega_n) \sum_{k=1}^{\infty} \mathbb{1}_{A^k}((\omega_1, \dots, \omega_n)) = \sum_{k=1}^{\infty} \tilde{\mu}(A^k).\end{aligned}$$

In particular, if $A = A^1 \uplus A^2$, one similarly gets $\tilde{\mu}(A) = \tilde{\mu}(A^1) + \tilde{\mu}(A^2)$. Hence $\tilde{\mu}$ is a σ -finite, additive, σ -subadditive set function on the semiring $\mathcal{Z}^R$ with $\tilde{\mu}(\emptyset) = 0$. By the measure extension theorem (Theorem 1.53), $\tilde{\mu}$ can be uniquely extended to a σ -finite measure on $\mathcal{A} = \sigma(\mathcal{Z}^R)$.

Example 14.15 For $i = 1, \dots, n$, let $(\Omega_i, \mathcal{A}_i, \mathbf{P}_i)$ be a probability space. On the space $(\Omega, \mathcal{A}, \mathbf{P}) := (\times_{i=1}^n \Omega_i, \bigotimes_{i=1}^n \mathcal{A}_i, \bigotimes_{i=1}^n \mathbf{P}_i)$, the coordinate maps $X_i : \Omega \rightarrow \Omega_i$ are independent with distribution $\mathbf{P}_{X_i} = \mathbf{P}_i$.

Theorem 14.16 (Fubini) Let $(\Omega_i, \mathcal{A}_i, \mu_i)$ be σ -finite measure spaces, $i = 1, 2$. Let $f : \Omega_1 \times \Omega_2 \rightarrow \overline{\mathbb{R}}$ be measurable with respect to $\mathcal{A}_1 \otimes \mathcal{A}_2$. If $f \geq 0$ or $f \in \mathcal{L}^1(\mu_1 \otimes \mu_2)$, then

$$\begin{aligned}\omega_1 &\mapsto \int f(\omega_1, \omega_2) \mu_2(d\omega_2) \quad \text{is } \mathcal{A}_1\text{-measurable,} \\ \omega_2 &\mapsto \int f(\omega_1, \omega_2) \mu_1(d\omega_1) \quad \text{is } \mathcal{A}_2\text{-measurable,}\end{aligned}\tag{14.6}$$

and

$$\begin{aligned}\int_{\Omega_1 \times \Omega_2} f d(\mu_1 \otimes \mu_2) &= \int_{\Omega_1} \left(\int_{\Omega_2} f(\omega_1, \omega_2) \mu_2(d\omega_2) \right) \mu_1(d\omega_1) \\ &= \int_{\Omega_2} \left(\int_{\Omega_1} f(\omega_1, \omega_2) \mu_1(d\omega_1) \right) \mu_2(d\omega_2)\end{aligned}\tag{14.7}$$

Proof The proof follows the usual procedure of stepwise approximations, starting with an indicator function.

First let $f = \mathbb{1}_A$ for $A = A_1 \times A_2$ with $A_1 \in \mathcal{A}_1$ and $A_2 \in \mathcal{A}_2$. Then (14.6) and (14.7) hold trivially. Building finite sums, this is also true for $A \in \mathcal{Z}_*^R$ (the algebra of finite unions of rectangles).

Now let $A \in \mathcal{A}_1 \otimes \mathcal{A}_2$. By the approximation theorem (Theorem 1.65), there is a sequence of sets $(A^n)_{n \in \mathbb{N}}$ in $\mathcal{Z}_*^R$ that approximate A in $\mu_1 \otimes \mu_2$ -measure. As limits of measurable functions are again measurable, and since by construction the integrals converge, (14.6) and (14.7) hold also for $f = \mathbb{1}_A$ and $A \in \mathcal{A}_1 \otimes \mathcal{A}_2$. Building finite sums, (14.6) and (14.7) also hold if f is a simple function.

Consider now $f \geq 0$. Then, by Theorem 1.96, there exists a sequence of simple functions $(f_n)_{n \in \mathbb{N}}$ with $f_n \uparrow f$. By the monotone convergence theorem (Theorem 4.20), (14.6) and (14.7) also hold for this f .

Now let $f \in \mathcal{L}^1(\mu_1 \otimes \mu_2)$. Then $f = f^+ - f^-$ with $f^+, f^- \geq 0$ being integrable functions. Since (14.6) and (14.7) hold for f^- and f^+ , they also hold for f .

In Definition 2.32, we defined the convolution of two real probability measures μ and ν as the distribution of the sum of two independent random variables with distributions μ and ν , respectively. As a simple application of Fubini's theorem, we can give a new definition for the convolution of, more generally, finite measures on $\mathbb{R}^n$. Of course, for real probability measures, it coincides with the old definition. If the measures have Lebesgue densities, then we obtain an explicit formula for the density of the convolution.

Let X and Y be $\mathbb{R}^n$ -valued random variables with densities f_X and f_Y . That is, $f_X, f_Y : \mathbb{R}^n \rightarrow [0, \infty]$ are measurable and integrable with respect to n -dimensional Lebesgue measure λ^n and, for all $x \in \mathbb{R}^n$,

$$\mathbf{P}[X \leq x] = \int_{(-\infty, x]} f_X(t) \lambda^n(dt) \quad \text{and} \quad \mathbf{P}[Y \leq x] = \int_{(-\infty, x]} f_Y(t) \lambda^n(dt).$$

Here $(-\infty, x] = \{y \in \mathbb{R}^n : y_i \leq x_i \text{ for } i = 1, \dots, n\}$ (compare (1.5)).

Definition 14.17 Let $n \in \mathbb{N}$. For two Lebesgue integrable maps $f, g : \mathbb{R}^n \rightarrow [0, \infty]$, define the convolution $f * g : \mathbb{R}^n \rightarrow [0, \infty]$ by

$$(f * g)(x) = \int_{\mathbb{R}^n} f(y)g(x - y)\lambda^n(dy)$$

For two finite measures $\mu, \nu \in \mathcal{M}_f(\mathbb{R}^n)$, define the convolution $\mu * \nu \in \mathcal{M}_f(\mathbb{R}^n)$ by

$$(\mu * \nu)((-\infty, x]) = \iint \mathbb{1}_{A_x}(u, v)\mu(du)\nu(dv)$$

where $A_x := \{(u, v) \in \mathbb{R}^n \times \mathbb{R}^n : u + v \leq x\}$.

Lemma 14.18 The map $f * g$ is measurable and we have $f * g = g * f$ and

$$\int_{\mathbb{R}^n} (f * g)d\lambda^n = \left(\int_{\mathbb{R}^n} f d\lambda^n \right) \left(\int_{\mathbb{R}^n} g d\lambda^n \right)$$

Furthermore, $\mu * \nu = \nu * \mu$ and $(\mu * \nu)(\mathbb{R}^n) = \mu(\mathbb{R}^n)\nu(\mathbb{R}^n)$.

Proof The claims follow immediately from Fubini's theorem.

Theorem 14.19 (Convolution of n -dimensional measures)

- (i) If X and Y are independent $\mathbb{R}^n$ -valued random variables with densities f_X and f_Y , then $X + Y$ has density $f_X * f_Y$.
- (ii) If $\mu = f\lambda^n$ and $\nu = g\lambda^n$ are finite measures with Lebesgue densities f and g , then $\mu * \nu = (f * g)\lambda^n$.

Proof (i) Let $x \in \mathbb{R}^n$ and $A := \{(u, v) \in \mathbb{R}^n \times \mathbb{R}^n : u + v \leq x\}$. Repeated application of Fubini's theorem and the translation invariance of λ^n yields

$$\begin{aligned} \mathbf{P}[X + Y \leq x] &= \mathbf{P}[(X, Y) \in A] \\ &= \int_{\mathbb{R}^n \times \mathbb{R}^n} \mathbb{1}_A(u, v) f_X(u) f_Y(v) (\lambda^n)^{\otimes 2}(d(u, v)) \\ &= \int_{\mathbb{R}^n} \left(\int_{\mathbb{R}^n} \mathbb{1}_A(u, v) f_X(u) \lambda^n(du) \right) f_Y(v) \lambda^n(dv) \\ &= \int_{\mathbb{R}^n} \left(\int_{(-\infty, x-v]} f_X(u) \lambda^n(du) \right) f_Y(v) \lambda^n(dv) \\ &= \int_{\mathbb{R}^n} \left(\int_{(-\infty, x]} f_X(u - v) \lambda^n(du) \right) f_Y(v) \lambda^n(dv) \\ &= \int_{(-\infty, x]} \left(\int_{\mathbb{R}^n} f_X(u - v) f_Y(v) \lambda^n(dv) \right) \lambda^n(du) \\ &= \int_{(-\infty, x]} (f_X * f_Y) d\lambda^n \end{aligned}$$

- (ii) In (i), replace μ by $\mathbf{P}_X$ and ν by $\mathbf{P}_Y$. The claim is immediate.

We come next to a concept that generalizes the notion of product measures and points in the direction of the example from the introduction to this chapter.

Recall the definition of a transition kernel from Definition 8.25.

Lemma 14.20 Let κ be a finite transition kernel from $(\Omega_1, \mathcal{A}_1)$ to $(\Omega_2, \mathcal{A}_2)$ and let $f : \Omega_1 \times \Omega_2 \rightarrow [0, \infty]$ be measurable with respect to $\mathcal{A}_1 \otimes \mathcal{A}_2 - \mathcal{B}([0, \infty])$. Then the map

$$I_f : \Omega_1 \rightarrow [0, \infty] \\ \omega_1 \mapsto \int f(\omega_1, \omega_2) \kappa(\omega_1, d\omega_2)$$

is well-defined and $\mathcal{A}_1$ -measurable.

Proof By Lemma 14.13, for every $\omega_1 \in \Omega_1$, the map f_{ω_1} is measurable with respect to $\mathcal{A}_2$; hence $I_f(\omega_1) = \int f_{\omega_1}(\omega_2) \kappa(\omega_1, d\omega_2)$ is well-defined. Hence, it remains to show measurability of I_f .

If $g = \mathbb{1}_{A_1 \times A_2}$ for some $A_1 \in \mathcal{A}_1$ and $A_2 \in \mathcal{A}_2$, then clearly $I_g(\omega_1) = \mathbb{1}_{A_1}(\omega_1) \kappa(\omega_1, A_2)$ is measurable. Now let

$$\mathcal{D} = \{A \in \mathcal{A}_1 \otimes \mathcal{A}_2 : I_{\mathbb{1}_A} \text{ is } \mathcal{A}_1\text{-measurable}\}$$

We show that $\mathcal{D}$ is a λ -system:

- (i) Evidently, $\Omega_1 \times \Omega_2 \in \mathcal{D}$.
- (ii) If $A, B \in \mathcal{D}$ with $A \subset B$, then $I_{\mathbb{1}_{B \setminus A}} = I_{\mathbb{1}_B} - I_{\mathbb{1}_A}$ is measurable, where we used the fact that κ is finite; hence $B \setminus A \in \mathcal{D}$.
- (iii) If $A_1, A_2, \dots \in \mathcal{D}$ are pairwise disjoint and $A := \bigcup_{n=1}^{\infty} A_n$, then $I_{\mathbb{1}_A} = \sum_{n=1}^{\infty} I_{\mathbb{1}_{A_n}}$ is measurable; hence $A \in \mathcal{D}$.

Summarising, $\mathcal{D}$ is a λ -system that contains a π -system that generates $\mathcal{A}_1 \otimes \mathcal{A}_2$ (namely, the rectangles). Hence, by the $\pi - \lambda$ theorem (Theorem 1.19), $\mathcal{D} = \mathcal{A}_1 \otimes \mathcal{A}_2$. Hence $I_{\mathbb{1}_A}$ is measurable for all $A \in \mathcal{A}_1 \otimes \mathcal{A}_2$. We infer that I_g is measurable for any simple function g . Now let $(f_n)_{n \in \mathbb{N}}$ be a sequence of simple functions with $f_n \uparrow f$. For any fixed $\omega_1 \in \Omega_1$, by the monotone convergence theorem, $I_f(\omega_1) = \lim_{n \rightarrow \infty} I_{f_n}(\omega_1)$. As a limit of measurable functions, I_f is measurable.

Remark 14.21 In the following, we often write $\int \kappa(\omega_1, d\omega_2) f(\omega_1, \omega_2)$ instead of $\int f(\omega_1, \omega_2) \kappa(\omega_1, d\omega_2)$ since for multiple integrals this notation allows us to write the integrator closer to the corresponding integral sign.

Theorem 14.22 Let $(\Omega_i, \mathcal{A}_i)$, $i = 0, 1, 2$, be measurable spaces. Let κ_1 be a finite transition kernel from $(\Omega_0, \mathcal{A}_0)$ to $(\Omega_1, \mathcal{A}_1)$ and let κ_2 be a finite transition kernel from $(\Omega_0 \times \Omega_1, \mathcal{A}_0 \otimes \mathcal{A}_1)$ to $(\Omega_2, \mathcal{A}_2)$. Then the map

$$\kappa_1 \otimes \kappa_2 : \Omega_0 \times (\mathcal{A}_1 \otimes \mathcal{A}_2) \rightarrow [0, \infty) \\ (\omega_0, A) \mapsto \int_{\Omega_1} \kappa_1(\omega_0, d\omega_1) \int_{\Omega_2} \kappa_2((\omega_0, \omega_1), d\omega_2) \mathbb{1}_A((\omega_1, \omega_2))$$

is well-defined and is a σ -finite (but not necessarily a finite) transition kernel from $(\Omega_0, \mathcal{A}_0)$ to $(\Omega_1 \times \Omega_2, \mathcal{A}_1 \otimes \mathcal{A}_2)$. If κ_1 and κ_2 are (sub)stochastic, then $\kappa_1 \otimes \kappa_2$ is (sub)stochastic. We call $\kappa_1 \otimes \kappa_2$ the product of κ_1 and κ_2 .

If κ_2 is a kernel from $(\Omega_1, \mathcal{A}_1)$ to $(\Omega_2, \mathcal{A}_2)$, then we define the product $\kappa_1 \otimes \kappa_2$ similarly by formally understanding κ_2 as a kernel from $(\Omega_0 \times \Omega_1, \mathcal{A}_0 \otimes \mathcal{A}_1)$ to $(\Omega_2, \mathcal{A}_2)$ that does not depend on the Ω_0 -coordinate.

Proof Let $A \in \mathcal{A}_1 \otimes \mathcal{A}_2$. By Lemma 14.20, the map

$$g_A : (\omega_0, \omega_1) \mapsto \int \kappa_2((\omega_0, \omega_1), d\omega_2) \mathbb{1}_A(\omega_1, \omega_2)$$

is well-defined and $\mathcal{A}_0 \otimes \mathcal{A}_1$ -measurable. Thus, again by Lemma 14.20, the map

$$\omega_0 \mapsto \kappa_1 \otimes \kappa_2(\omega_0, A) = \int \kappa_1(\omega_0, d\omega_1) g_A(\omega_0, \omega_1)$$

is well-defined and $\mathcal{A}_0$ -measurable. For fixed ω_0 , by the monotone convergence theorem, the map $A \mapsto \kappa_1 \otimes \kappa_2(\omega_0, A)$ is σ -additive and thus a measure.

For $\omega_0 \in \Omega_0$ and $n \in \mathbb{N}$, let $A_{\omega_0, n} := \{\omega_1 \in \Omega_1 : \kappa_2((\omega_0, \omega_1), \Omega_2) < n\}$. Since κ_2 is finite, we have $\bigcup_{n \geq 1} A_{\omega_0, n} = \Omega_1$ for all $\omega_0 \in \Omega_0$. Furthermore, $\kappa_1 \otimes \kappa_2(\omega_0, A_n \times \Omega_2) \leq n \cdot \kappa_1(\omega_0, \bar{A}_n) < \infty$. Hence $\kappa_1 \otimes \kappa_2(\omega_0, \cdot)$ is σ -finite and is thus a transition kernel.

The supplement is trivial.

Corollary 14.23 (Products via kernels) Let $(\Omega_1, \mathcal{A}_1, \mu)$ be a finite measure space, let $(\Omega_2, \mathcal{A}_2)$ be a measurable space and let κ be a finite transition kernel from Ω_1 to Ω_2 . Then there exists a unique σ -finite measure $\mu \otimes \kappa$ on $(\Omega_1 \times \Omega_2, \mathcal{A}_1 \otimes \mathcal{A}_2)$ with

$$\mu \otimes \kappa(A_1 \times A_2) = \int_{A_1} \kappa(\omega_1, A_2) \mu(d\omega_1) \quad \text{for all } A_1 \in \mathcal{A}_1, A_2 \in \mathcal{A}_2$$

If κ is stochastic and if μ is a probability measure, then $\mu \otimes \kappa$ is a probability measure.

Proof Apply Theorem 14.22 with $\kappa_2 = \kappa$ and $\kappa_1(\omega_0, \cdot) = \mu$.

Corollary 14.24 Let $n \in \mathbb{N}$ and let $(\Omega_i, \mathcal{A}_i)$, $i = 0, \dots, n$, be measurable spaces. For $i = 1, \dots, n$, let κ_i be a substochastic kernel from $(\times_{k=0}^{i-1} \Omega_k, \bigotimes_{k=0}^{i-1} \mathcal{A}_k)$ to $(\Omega_i, \mathcal{A}_i)$ or from $(\Omega_{i-1}, \mathcal{A}_{i-1})$ to $(\Omega_i, \mathcal{A}_i)$. Then the recursion $\kappa_1 \otimes \dots \otimes \kappa_i := (\kappa_1 \otimes \dots \otimes \kappa_{i-1}) \otimes \kappa_i$ for any $i = 1, \dots, n$ defines a substochastic kernel $\bigotimes_{k=1}^i \kappa_k := \kappa_1 \otimes \dots \otimes \kappa_i$ from $(\Omega_0, \mathcal{A}_0)$ to $(\times_{k=1}^i \Omega_k, \bigotimes_{k=1}^i \mathcal{A}_k)$. If all κ_k are stochastic, then all $\bigotimes_{k=1}^i \kappa_k$ are stochastic.

If μ is a finite measure on $(\Omega_0, \mathcal{A}_0)$, then $\mu_i := \mu \otimes \bigotimes_{k=1}^i \kappa_k$ is a finite measure on $(\times_{k=0}^i \Omega_k, \bigotimes_{k=0}^i \mathcal{A}_k)$. If μ is a probability measure and if every κ_i is stochastic, then μ_i is a probability measure.

Proof The claims follow inductively by Theorem 14.22.

Definition 14.25 (Composition of kernels) Let $(\Omega_i, \mathcal{A}_i)$ be measurable spaces, $i = 0, 1, 2$, and let κ_i be a substochastic kernel from $(\Omega_{i-1}, \mathcal{A}_{i-1})$ to $(\Omega_i, \mathcal{A}_i)$, $i = 1, 2$. Define the composition of κ_1 and κ_2 by

$$\begin{aligned} \kappa_1 \cdot \kappa_2 : \Omega_0 \times \mathcal{A}_2 &\rightarrow [0, \infty) \\ (\omega_0, A_2) &\mapsto \int_{\Omega_1} \kappa_1(\omega_0, d\omega_1) \kappa_2(\omega_1, A_2) \end{aligned}$$

Theorem 14.26 If we denote by $\pi_2 : \Omega_1 \times \Omega_2 \rightarrow \Omega_2$ the projection to the second coordinate, then

$$(\kappa_1 \cdot \kappa_2)(\omega_0, A_2) = (\kappa_1 \otimes \kappa_2)(\omega_0, \pi_2^{-1}(A_2)) \quad \text{for all } A_2 \in \mathcal{A}_2$$

In particular, the composition $\kappa_1 \cdot \kappa_2$ is a (sub)stochastic kernel from $(\Omega_0, \mathcal{A}_0)$ to $(\Omega_2, \mathcal{A}_2)$.

Proof This is obvious.

Lemma 14.27 (Kernels and convolution) Let μ and ν be probability measures on $\mathbb{R}^d$ and define the kernels $\kappa_i : (\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d)) \rightarrow (\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$, $i = 1, 2$, by $\kappa_1(x, dy) = \mu(dy)$ and $\kappa_2(y, dz) = (\delta_y * \nu)(dz)$. Then $\kappa_1 \cdot \kappa_2 = \mu * \nu$.

Proof This is trivial.

Theorem 14.28 (Kernels and convolution) Assume $X_1, X_2, \dots$ are independent $\mathbb{R}^d$ -valued random variables with distributions $\mu_i := \mathbf{P}_{X_i}$, $i = 1, \dots, n$. Let $S_k := X_1 + \dots + X_k$ for $k = 1, \dots, n$, and define stochastic kernels from $\mathbb{R}^d$ to $\mathbb{R}^d$ by $\kappa_k(x, \cdot) = \delta_x * \mu_k$ for $k = 1, \dots, n$. Then

$$\left(\bigotimes_{k=1}^n \kappa_k \right) (0, \cdot) = \mathbf{P}_{(S_1, \dots, S_n)} \quad (14.8)$$

Proof For $k = 1, \dots, n$, define the measurable bijection $\varphi_k : (\mathbb{R}^d)^k \rightarrow (\mathbb{R}^d)^k$ by

$$\varphi_k(x_1, \dots, x_k) = (x_1, x_1 + x_2, \dots, x_1 + \dots + x_k).$$

Evidently, $\mathcal{B}((\mathbb{R}^d)^n) = \sigma(\varphi_n(A_1 \times \dots \times A_n) : A_1, \dots, A_n \in \mathcal{B}(\mathbb{R}^d))$. Hence, it is enough to show (14.8) for sets of this type. That is, it is enough to show that

$$\left(\bigotimes_{k=1}^n \kappa_k \right) (0, \varphi_n(A_1 \times \dots \times A_n)) = \mathbf{P}_{(S_1, \dots, S_n)}(\varphi_n(A_1 \times \dots \times A_n)) = \prod_{k=1}^n \mu_k(A_k).$$

For $n = 1$, this is clear. By definition, $\kappa_n(y_{n-1}, y_{n-1} + A_n) = \mu_n(A_n)$. Inductively, we get

$$\begin{aligned} & \left(\bigotimes_{k=1}^n \kappa_k \right) (0, \varphi_n(A_1 \times \dots \times A_n)) \\ &= \int_{\varphi_{n-1}(A_1 \times \dots \times A_{n-1})} \left(\bigotimes_{k=1}^{n-1} \kappa_k \right) (0, d(y_1, \dots, y_{n-1})) \kappa_n(y_{n-1}, y_{n-1} + A_n) \\ &= \left(\prod_{k=1}^{n-1} \mu_k(A_k) \right) \mu_n(A_n) \end{aligned}$$

□

Theorem 14.29 (Fubini for transition kernels) Let $(\Omega_i, \mathcal{A}_i)$ be measurable spaces, $i = 1, 2$. Let μ be a finite measure on $(\Omega_1, \mathcal{A}_1)$ and let κ be a finite transition kernel from Ω_1 to Ω_2 . Assume that $f : \Omega_1 \times \Omega_2 \rightarrow \overline{\mathbb{R}}$ is measurable with respect to $\mathcal{A}_1 \otimes \mathcal{A}_2$. If $f \geq 0$ or $f \in \mathcal{L}^1(\mu \otimes \kappa)$, then

$$\int_{\Omega_1 \times \Omega_2} f d(\mu \otimes \kappa) = \int_{\Omega_1} \left(\int_{\Omega_2} f(\omega_1, \omega_2) \kappa(\omega_1, d\omega_2) \right) \mu(d\omega_1) \quad (14.9)$$

Proof For $f = \mathbb{1}_{A_1 \times A_2}$ with $A_1 \in \mathcal{A}_1$ and $A_2 \in \mathcal{A}_2$, the statement is true by definition. For general f , apply the usual approximation argument as in Theorem 14.16. □

Example 14.30 We come back to the example from the beginning of this chapter. Let $n \in \mathbb{N}$ and let $(\Omega_2, \mathcal{A}_2) = (\{0, 1\}^n, (2^{\{0, 1\}})^{\otimes n})$ be the space of n -fold coin tossing. For any $p \in [0, 1]$, define

$$P_p = (\text{Ber}_p)^{\otimes n} = ((1 - p)\delta_0 + p\delta_1)^{\otimes n}.$$

P_p is that probability measure on $(\Omega_2, \mathcal{A}_2)$ under which the coordinate maps Y_i are independent Bernoulli random variables with success probability p .

Further, let $\Omega_1 = [0, 1]$, let $\mathcal{A}_1 = \mathcal{B}([0, 1])$ be the Borel σ -algebra on Ω_1 and let $\mu = \mathcal{U}_{[0, 1]}$ be the uniform distribution on $[0, 1]$. Then the identity map $X : \Omega_1 \rightarrow [0, 1]$ is a random variable on $(\Omega_1, \mathcal{A}_1, \mu)$ that is uniformly distributed on $[0, 1]$.

Finally, consider the stochastic kernel κ from Ω_1 to Ω_2 , defined by

$$\kappa(\omega_1, A_2) = P_{\omega_1}(A_2).$$

If we let $\Omega = \Omega_1 \times \Omega_2$, $\mathcal{A} = \mathcal{A}_1 \otimes \mathcal{A}_2$ and $\mathbf{P} = \mu \otimes \kappa$, then X and $Y_1, \dots, Y_n$ describe precisely the random variables on $(\Omega, \mathcal{A}, \mathbf{P})$ from the beginning of this chapter.

Remark 14.31 The procedure can be extended to n -stage experiments. Let $(\Omega_i, \mathcal{A}_i)$ be the measurable space of the i th experiment, $i = 0, \dots, n - 1$. Let P_0 be a probability measure on $(\Omega_0, \mathcal{A}_0)$. Assume that for $i = 1, \dots, n - 1$, the distribution on $(\Omega_i, \mathcal{A}_i)$ depends on $(\omega_1, \dots, \omega_{i-1})$ and is given by a stochastic kernel κ_i from $\Omega_0 \times \dots \times \Omega_{i-1}$ to Ω_i . The whole n -stage experiment is then described by the coordinate maps on the probability space

$$\left(\bigvee_{i=0}^{n-1} \Omega_i, \bigotimes_{i=0}^{n-1} \mathcal{A}_i, \mathbf{P}_0 \otimes \bigotimes_{i=1}^{n-1} \kappa_i \right)$$

Exercise 14.2.1 Show the following convolution formulas.

- (i) Normal distribution: $\mathcal{N}_{\mu_1, \sigma_1^2} * \mathcal{N}_{\mu_2, \sigma_2^2} = \mathcal{N}_{\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2}$ for all $\mu_1, \mu_2 \in \mathbb{R}$ and $\sigma_1^2, \sigma_2^2 > 0$.
- (ii) Gamma distribution: $\Gamma_{\theta, r} * \Gamma_{\theta, s} = \Gamma_{\theta, r+s}$ for all $\theta, r, s > 0$.
- (iii) Cauchy distribution: $\text{Cau}_r * \text{Cau}_s = \text{Cau}_{r+s}$ for all $r, s > 0$.

Exercise 14.2.2 (Hilbert-Schmidt operator) Let $(\Omega_i, \mathcal{A}_i, \mu_i), i = 1, 2$, be σ -finite measure spaces and let $a : \Omega_1 \times \Omega_2 \rightarrow \mathbb{R}$ be measurable with

$$\int \mu_1(dt_1) \int \mu_2(dt_2) a(t_1, t_2)^2 < \infty$$

For $f \in \mathcal{L}^2(\mu_1)$, define

$$(Af)(t_2) = \int a(t_1, t_2) f(t_1) \mu_1(dt_1)$$

Show that A is a continuous linear operator from $\mathcal{L}^2(\mu_1)$ to $\mathcal{L}^2(\mu_2)$.

Exercise 14.2.3 (Partial integration) Let F_μ and F_ν be the distribution functions of locally finite measures μ and ν on $\mathbb{R}$. For $x \in \mathbb{R}$, define the left-sided limit $F(x-) = \sup_{y < x} F(y)$ and the jump height $\Delta F(x) = F(x) - F(x-)$. Show that, for $a < b$,

$$\begin{aligned} \int_{(a,b]} F_\mu d\nu &= F_\mu(b)F_\nu(b) - F_\mu(a)F_\nu(a) - \int_{(a,b]} F_\nu(x-) \mu(dx) \\ &= F_\mu(b)F_\nu(b) - F_\mu(a)F_\nu(a) - \int_{(a,b]} F_\nu d\mu + \sum_{a < x \leq b} \Delta F_\mu(x) \Delta F_\nu(x) \end{aligned}$$

5.14.3 14.3 Kolmogorov's Extension Theorem

In the previous section, we saw how we can implement n -stage experiments on a probability space. In this section, we first show how to implement countably many successive experiments on one probability space (Ionescu-Tulcea's theorem). Thereafter we also construct probability measures on products of uncountably many spaces (Kolmogorov's extension theorem).

Let $(\Omega_i, \mathcal{A}_i), i \in \mathbb{N}_0$, be measurable spaces and let P_0 be a probability measure on $(\Omega_0, \mathcal{A}_0)$. Let $\Omega^i := \times_{k=0}^i \Omega_k$ and $\mathcal{A}^i = \bigotimes_{k=0}^i \mathcal{A}_k$ and

$$\Omega := \prod_{k=0}^{\infty} \Omega_k \quad \text{and} \quad \mathcal{A} = \bigotimes_{k=0}^{\infty} \mathcal{A}_k$$

For every $i \in \mathbb{N}$, let κ_i be a stochastic kernel from $(\Omega^{i-1}, \mathcal{A}^{i-1})$ to $(\Omega_i, \mathcal{A}_i)$. In Corollary 14.24, we defined inductively probability measures $P_i = P_0 \otimes \bigotimes_{k=1}^i \kappa_k$ on $(\Omega^i, \mathcal{A}^i)$. By construction, for $i, j \geq k$ and $A \in \mathcal{A}^k$, we had

$$P_i(A \times \Omega_{k+1} \times \dots \times \Omega_i) = P_j(A \times \Omega_{k+1} \times \dots \times \Omega_j). \quad (14.10)$$

Now we want to define a probability measure P on $(\Omega, \mathcal{A})$ such that for $k \in \mathbb{N}_0$ and $A \in \mathcal{A}^k$

$$P\left(A \times \prod_{i=k+1}^{\infty} \Omega_i\right) = P_k(A) \quad (14.11)$$

Theorem 14.32 (Ionescu-Tulcea) There is a uniquely determined probability measure on $(\Omega, \mathcal{A})$ such that (14.11) holds.

Proof Uniqueness is clear since the finite-dimensional rectangular cylinders form a π -system that generates $\mathcal{A}$. It remains to show the existence of that measure.

We use (14.11) to define a set function P on cylinder sets. Clearly, P is additive and is hence a content. If we can show that P is $\emptyset$ -continuous, then P is a premeasure (by Theorem 1.36) and thus by Carathéodory's theorem (Theorem 1.41) can be extended uniquely to a measure on $\mathcal{A}$.

Hence, let $A_0 \supset A_1 \supset A_2 \supset \dots$ be a sequence in $\mathcal{Z}$ with $\alpha := \inf_{n \in \mathbb{N}_0} P(A_n) > 0$. It is enough to show that $\bigcap_{n=0}^{\infty} A_n \neq \emptyset$. Without loss of generality, we can assume that $A_n = A'_n \times \times_{k=n+1}^{\infty} \Omega_k$ for certain $A'_n \in \mathcal{A}^n$. For $n \geq m$, define

$$h_{m,n}(\omega_0, \dots, \omega_m) := \left(\delta_{(\omega_0, \dots, \omega_m)} \otimes \bigotimes_{k=m+1}^n \kappa_k \right) (A'_n)$$

and $h_m := \inf_{n \geq m} h_{m,n}$. Inductively, we show that for every $i \in \mathbb{N}_0$, there exists a $\varrho_i \in \Omega_i$ such that

$$h_m(\varrho_0, \dots, \varrho_m) \geq \alpha \quad (14.12)$$

Since $A'_{n+1} \subset A'_n \times \Omega_{n+1}$, we have

$$\begin{aligned} h_{m,n+1}(\omega_0, \dots, \omega_m) &= \left(\delta_{(\omega_0, \dots, \omega_m)} \otimes \bigotimes_{k=m+1}^{n+1} \kappa_k \right) (A'_{n+1}) \\ &\leq \left(\delta_{(\omega_0, \dots, \omega_m)} \otimes \bigotimes_{k=m+1}^{n+1} \kappa_k \right) (A'_n \times \Omega_{n+1}) \\ &= \left(\delta_{(\omega_0, \dots, \omega_m)} \otimes \bigotimes_{k=m+1}^n \kappa_k \right) (A'_n) = h_{m,n}(\omega_0, \dots, \omega_m) \end{aligned}$$

Hence $h_{m,n} \downarrow h_m$ for $n \rightarrow \infty$ and by the monotone convergence theorem,

$$\int h_m dP_m = \inf_{n \geq m} \int h_{m,n} dP_m = \inf_{n \in \mathbb{N}_0} P_n(A'_n) = \alpha$$

whence we have (14.12) for $m = 0$. Now assume that (14.12) holds for $m \in \mathbb{N}_0$. Then

$$\begin{aligned} &\int h_{m+1}(\varrho_0, \dots, \varrho_m, \omega_{m+1}) \kappa_{m+1}((\varrho_0, \dots, \varrho_m), d\omega_{m+1}) \\ &= \inf_{n \geq m+1} \int h_{m+1,n}(\varrho_0, \dots, \varrho_m, \omega_{m+1}) \kappa_{m+1}((\varrho_0, \dots, \varrho_m), d\omega_{m+1}) \\ &= h_m(\varrho_0, \dots, \varrho_m) \geq \alpha \end{aligned}$$

Hence (14.12) holds for $m + 1$.

Let $\varrho := (\varrho_0, \varrho_1, \dots) \in \Omega$. By construction,

$$\alpha \leq h_{m,m}(\varrho_0, \dots, \varrho_m) = \mathbb{1}_{A'_m}(\varrho_0, \dots, \varrho_m)$$

hence $\varrho \in A_m$ for all $m \in \mathbb{N}_0$ and thus $\bigcap_{i=0}^{\infty} A_i \neq \emptyset$.

Corollary 14.33 (Product measure) For every $n \in \mathbb{N}_0$, let P_n be a probability measure on $(\Omega_n, \mathcal{A}_n)$. Then there exists a uniquely determined probability measure P on $(\Omega, \mathcal{A})$ with

$$P\left(A_0 \times \dots \times A_n \times \prod_{i=n+1}^{\infty} \Omega_i\right) = \prod_{k=0}^n P_k(A_k)$$

for $A_i \in \mathcal{A}_i, i = 0, \dots, n$ and $n \in \mathbb{N}_0$.

$\bigotimes_{i=0}^{\infty} P_i := P$ is called the product of the measures $P_0, P_1, \dots$. Under P , the coordinate maps $(X_i)_{i \in \mathbb{N}_0}$ are independent.

Proof This follows by Ionescu-Tulcea's theorem with

$$\kappa_i((\omega_0, \dots, \omega_{i-1}), \cdot) = P_i$$

We want to make a statement similar to that of Ionescu-Tulcea's theorem; however, without the assumption that the measures P_k are defined a priori by kernels. Before we formulate the theorem, we generalize the consistency condition (14.10). Recall that for $L \subset J \subset I, X_L^J : \Omega_J \rightarrow \Omega_L$ denotes the canonical projection.

Definition 14.34 A family $(P_J, J \subset I \text{ finite})$ of probability measures on the space $(\Omega_J, \mathcal{A}_J)$ is called consistent if

$$P_L = P_J \circ (X_L^J)^{-1} \quad \text{for all } L \subset J \subset I \text{ finite.}$$

Recall that $\Omega = \prod_{i \in I} \Omega_i$ and $\mathcal{A} = \bigotimes_{i \in I} \mathcal{A}_i$. Let P be a probability measure on $(\Omega, \mathcal{A})$. Since $X_L = X_L^J \circ X_J$, the family $(P_J := P \circ X_J^{-1}, J \subset I \text{ finite})$ is consistent. Thus, consistency is a necessary condition for the existence of a measure P on the product space with $P_J := P \circ X_J^{-1}$. If all the measurable spaces are Borel spaces (recall Definition 8.35), for example $\mathbb{R}^d, \mathbb{Z}^d, C([0, 1])$ or more general Polish spaces, then this condition is also sufficient. We formulate this statement first for a countable index set.

Theorem 14.35 Let I be countable and let $(\Omega_i, \mathcal{A}_i)$ be Borel spaces for all $i \in I$. Let $(P_J, J \subset I \text{ finite})$ be a consistent family of probability measures. Then there exists a unique probability measure P on $(\Omega, \mathcal{A})$ with $P_J = P \circ X_J^{-1}$ for all finite $J \subset I$.

Proof Without loss of generality, assume $I = \mathbb{N}_0$. Let $P_n := P_{\{0, \dots, n\}}, \Omega^n := \Omega_{\{0, \dots, n\}}$ and $\mathcal{A}^n := \mathcal{A}_{\{0, \dots, n\}}$. It is easy to check that finite products of Borel spaces are again Borel spaces; hence $(\Omega^n, \mathcal{A}^n)$ is Borel for all $n \in \mathbb{N}_0$.

Let $\mathcal{F} := \{A \times \Omega_{n+1} : A \in \mathcal{A}^n\}, Y : \Omega^{n+1} \rightarrow \Omega_{n+1}, (\omega_0, \dots, \omega_{n+1}) \mapsto \omega_{n+1}$ and $Z : \Omega^{n+1} \rightarrow \Omega^n, (\omega_0, \dots, \omega_{n+1}) \mapsto (\omega_0, \dots, \omega_n)$. By Theorem 8.37 (with $\Omega = \Omega^{n+1}, \mathcal{A} = \mathcal{A}^{n+1}$ and $E = \Omega_{n+1}$), there is a stochastic kernel κ'_{n+1} from $(\Omega^{n+1}, \mathcal{F})$ to $(\Omega_{n+1}, \mathcal{A}_{n+1})$ such that κ'_{n+1} is a regular conditional distribution of Y given $\mathcal{F}$ (under the probability measure P_{n+1}). Hence, for $A \in \mathcal{A}^n$ and $B \in \mathcal{A}_{n+1}$, we have (compare (8.10))

$$P_{n+1}(A \times B) = \int \mathbb{1}_B(Y) \mathbb{1}_{A \times \Omega_{n+1}} dP_{n+1} = \int \kappa'_{n+1}(\cdot, B) \mathbb{1}_{A \times \Omega_{n+1}} dP_{n+1}$$

Since $\kappa'_{n+1}(\cdot, B)$ is $\mathcal{F}$ -measurable, there is a stochastic kernel κ_{n+1} from $(\Omega^n, \mathcal{A}^n)$ to $(\Omega_{n+1}, \mathcal{A}_{n+1})$ such that

$$\kappa_{n+1}((\omega_0, \dots, \omega_n), \cdot) = \kappa'_{n+1}((\omega_0, \dots, \omega_{n+1}), \cdot) \quad \text{for all } \omega_0, \dots, \omega_{n+1}.$$

Hence

$$\kappa'_{n+1}(\cdot, B) = \kappa_{n+1}(Z(\cdot), B) \quad \text{and} \quad \mathbb{1}_{A \times \Omega_{n+1}} = \mathbb{1}_A(Z).$$

We infer that

$$\begin{aligned} P_{n+1}(A \times B) &= \int \kappa_{n+1}(Z, B) \mathbb{1}_A(Z) dP_{n+1} \\ &= \int \kappa_{n+1}(\cdot, B) \mathbb{1}_A d(P_{n+1} \circ Z^{-1}) \\ &= \int \kappa_{n+1}(\cdot, B) dP_n \end{aligned}$$

Note that in the last equality we used the fact that $(P_n)_{n \in \mathbb{N}_0}$ is a projective family. By Corollary 14.23, we get $P_{n+1} = P_n \otimes \kappa_{n+1}$. Recursively, we obtain $P_n = P_0 \otimes \bigotimes_{k=1}^n \kappa_k$ for all $n \in \mathbb{N}$. Using Theorem 14.32, this yields the claim.

The last step in our construction is to replace the countable index set I by an arbitrary index set.

Theorem 14.36 (Kolmogorov's extension theorem) Let I be an arbitrary index set and let $(\Omega_i, \mathcal{A}_i)$ be Borel spaces, $i \in I$. Let $(P_J, J \subset I \text{ finite})$ be a consistent family of probability measures. Then there exists a unique probability measure P on $(\Omega, \mathcal{A})$ with $P_J = P \circ X_J^{-1}$ for every finite $J \subset I$. P is called the projective limit and will be denoted by $P =: \lim_{\leftarrow J \uparrow I} P_J$.

Proof For countable $J \subset I$, by Theorem 14.35, there is a unique probability measure P_J on $(\Omega_J, \mathcal{A}_J)$ with $P_J \circ (X_K^J)^{-1} = P_K$ for finite $K \subset J$. By defining $\tilde{P}_J(X_J^{-1}(A_J)) := P_J(A_J)$ for $A_J \in \mathcal{A}_J$, we get a probability measure $\tilde{P}_J$ on $(\Omega, \sigma(X_J))$.

Let $J, J' \subset I$ be countable and let $A \in \sigma(X_J) \cap \sigma(X_{J'}) \cap \mathcal{Z}$ be a $\sigma(X_J) \cap \sigma(X_{J'})$ -measurable cylinder with a finite base. Then there exists a finite $K \subset J \cap J'$ and $A_K \in \mathcal{A}_K$ with $A = X_K^{-1}(A_K)$. Hence $\tilde{P}_J(A) = P_K(A_K) = \tilde{P}_{J'}(A)$. Moreover, by Theorem 14.12, $\tilde{P}_J(A) = P_K(A_K) = \tilde{P}_{J'}(A)$ for all $A \in \sigma(X_J) \cap \sigma(X_{J'})$. Now, by Exercise 14.1.1, for any $A \in \mathcal{A}$, there is a countable $J \subset I$ with $A \in \sigma(X_J)$. Hence, independently of the choice of J , we can uniquely define a set function P on $\mathcal{A}$ by $P(A) = \tilde{P}_J(A)$. It remains to show that P is a probability measure. Evidently, $P(\Omega) = 1$. If $A_1, A_2, \dots \in \mathcal{A}$ are pairwise disjoint and $A := \bigcup_{n=1}^{\infty} A_n$, then for any $n \in \mathbb{N}$, there is a countable $J_n \subset I$ with $A_n \in \sigma(X_{J_n})$. Define $J = \bigcup_{n \in \mathbb{N}} J_n$. Then each A_n is in $\sigma(X_J)$; thus $A \in \sigma(X_J)$. Therefore,

$$P(A) = \tilde{P}_J(A) = \sum_{n=1}^{\infty} \tilde{P}_J(A_n) = \sum_{n=1}^{\infty} P(A_n).$$

This shows that P is a probability measure.

Example 14.37 Let $((\Omega_i, \tau_i), i \in I)$ be an arbitrary family of Polish spaces (recall from Theorem 8.36 that Polish spaces are also Borel spaces). Let $\mathcal{A}_i = \sigma(\tau_i)$ and let

P_i be an arbitrary probability measure on $(\Omega_i, \mathcal{A}_i)$ for every $i \in I$. For finite $J \subset I$, let $P_J := \bigotimes_{j \in J} P_j$ be the product measure of the $P_j, j \in J$. Evidently, the family $(P_J, J \subset I \text{ finite})$ is consistent. We call

$$P = \bigotimes_{i \in I} P_i := \lim_{\leftarrow J \uparrow I} P_J$$

the product measure on $(\Omega, \mathcal{A})$. Under P , all coordinate maps X_j are independent.

Example 14.38 (Pólya's urn model) (Compare Example 12.29.) In an urn there are initially k red and $n - k$ blue balls. At each step, one ball is drawn at random and is returned to the urn with an additional ball of the same color. Hence, at time $i \in \mathbb{N}_0$ there are $n + i$ balls in the urn. The random number of red balls is denoted by X_i .

For a more formal description, let $n \in \mathbb{N}$ and $k \in \{0, \dots, n\}$. Let $I = \mathbb{N}_0, \Omega_i = \{0, \dots, n + i\}, i \in \mathbb{N}$. Let $P_0[\{k\}] = 1$, and define the stochastic kernels κ_i from Ω_i to Ω_{i+1} by

$$\kappa_i(x_i, \{x_{i+1}\}) = \begin{cases} \frac{x_i}{n+i}, & \text{if } x_{i+1} = x_i + 1 \\ 1 - \frac{x_i}{n+i}, & \text{if } x_{i+1} = x_i \\ 0, & \text{else.} \end{cases}$$

Now let $P_{i+1} = P_i \otimes \kappa_i$. Under the measure $\mathbf{P} = \lim_{\leftarrow i \rightarrow \infty} P_i$, the projections $(X_i, i \in \mathbb{N}_0)$ describe Pólya's urn model.

5.14.4 14.4 Markov Semigroups

Definition 14.39 Let E be a Polish space. Let $I \subset \mathbb{R}$ be a nonempty index set and let $(\kappa_{s,t} : s, t \in I, s < t)$ be a family of stochastic kernels from E to E . We say that the family is consistent if $\kappa_{r,s} \cdot \kappa_{s,t} = \kappa_{r,t}$ for any choice of $r, s, t \in I$ with $r < s < t$.

Definition 14.40 Let E be a Polish space. Let $I \subset [0, \infty)$ be an additive semigroup (for example, $I = \mathbb{N}_0$ or $I = [0, \infty)$). A family $(\kappa_t : t \in I)$ of stochastic kernels is called a semigroup of stochastic kernels, or a Markov semigroup, if

$$\kappa_0(\omega, \cdot) = \delta_\omega \quad \text{for all } \omega \in \Omega \quad (14.13)$$

and if it satisfies the Chapman-Kolmogorov equation:

$$\kappa_s \cdot \kappa_t = \kappa_{s+t} \quad \text{for all } s, t \in I \quad (14.14)$$

Indeed, $(\{\kappa_t : t \in I\}, \cdot)$ is a semigroup in the algebraic sense and the map $t \rightarrow \kappa_t$ is a homomorphism of semigroups. In particular, the kernels commute in the sense that $\kappa_s \cdot \kappa_t = \kappa_t \cdot \kappa_s$ for all $s, t \in I$.

Lemma 14.41 If $(\kappa_t : t \in I)$ is a Markov semigroup, then the family of kernels, defined by $\tilde{\kappa}_{s,t} := \kappa_{t-s}$ for $t > s$, is consistent.

Proof This is trivial.

Theorem 14.42 (Kernel via a consistent family of kernels) Let $I \subset [0, \infty)$ with $0 \in I$ and let $(\kappa_{s,t} : s, t \in I, s < t)$ be a consistent family of stochastic kernels on the Polish space E . Then there exists a kernel κ from $(E, \mathcal{B}(E))$ to $(E^I, \mathcal{B}(E)^{\otimes I})$ such that, for all $x \in E$ and for any choice of finitely many numbers $0 = j_0 < j_1 < j_2 < \dots < j_n$ from I , and with the notation $J := \{j_0, \dots, j_n\}$, we have

$$\kappa(x, \cdot) \circ X_J^{-1} = \left(\delta_x \otimes \bigotimes_{k=0}^{n-1} \kappa_{j_k, j_{k+1}} \right) \quad (14.15)$$

Proof First we show that, for fixed $x \in E$, (14.15) defines a probability measure $\kappa(x, \cdot)$. Define the family $(P_J : J \subset I \text{ finite}, 0 \in J)$ by $P_J := \delta_x \otimes \bigotimes_{k=0}^{n-1} \kappa_{j_k, j_{k+1}}$. By Kolmogorov's extension theorem, it is enough to show that this family is consistent. In fact, if for $0 \notin J \subset I$ finite, we define P_J as the projection of $P_{J \cup \{0\}}$ to E^J , then the family $(P_J : J \subset I \text{ finite})$ is projective. Hence, let $0 \in L \subset J \subset I$ with $J \subset I$ finite. We have to show that $P_J \circ (X_L^J)^{-1} = P_L$. We may assume that $L = J \setminus \{j_l\}$ for some $l = 1, \dots, n$. The general case can be inferred inductively.

First consider $l = n$. Let $A_{j_0}, \dots, A_{j_{n-1}} \in \mathcal{B}(E)$ and $A := \times_{j \in L} A_j$. Then

$$\begin{aligned} P_J \circ (X_L^J)^{-1}(A) &= P_J(A \times E) = P_L \otimes \kappa_{j_{n-1}, j_n}(A \times E) \\ &= \int_A P_L(d(\omega_0, \dots, \omega_{n-1})) \kappa_{j_{n-1}, j_n}(\omega_{n-1}, E) = P_L(A) \end{aligned}$$

Now let $l \in \{1, \dots, n-1\}$. For all $j \in L$, let $A_j \in \mathcal{B}(E)$ and $A_{j_l} := E$. Define $A := \times_{j \in L} A_j$, and abbreviate $A' = \times_{k=0}^{l-1} A_{j_k}$ and $P' = \delta_x \otimes \bigotimes_{k=0}^{l-2} \kappa_{j_k, j_{k+1}}$. For $i = 0, \dots, n-1$, let

$$f_i(\omega_i) = \left(\bigotimes_{k=i}^{n-1} \kappa_{j_k, j_{k+1}} \right) (\omega_i, A_{j_{i+1}} \times \dots \times A_{j_n})$$

By assumption and using Fubini's theorem, we get

$$\begin{aligned}
f_{l-1}(\omega_{l-1}) &= \int_E \kappa_{j_{l-1}, j_l}(\omega_{l-1}, d\omega_l) \int_{A_{j_{l+1}}} \kappa_{j_l, j_{l+1}}(\omega_l, d\omega_{l+1}) f_{l+1}(\omega_{l+1}) \\
&= \int_{A_{j_{l+1}}} \kappa_{j_{l-1}, j_{l+1}}(\omega_{l-1}, d\omega_{l+1}) f_{l+1}(\omega_{l+1})
\end{aligned}$$

This implies

$$\begin{aligned}
P_J \circ (X_L^J)^{-1}(A) &= \int_{A'} P'(d(\omega_0, \dots, \omega_{l-1})) f_{l-1}(\omega_{l-1}) \\
&= \int_{A'} P'(d(\omega_0, \dots, \omega_{l-1})) \int_{A_{j_{l+1}}} (\kappa_{j_{l-1}, j_{l+1}})(\omega_{l-1}, d\omega_{l+1}) f_{l+1}(\omega_{l+1}) \\
&= P_L(A)
\end{aligned}$$

It remains to show that κ is a stochastic kernel. That is, we have to show that $x \mapsto \kappa(x, A)$ is measurable with respect to $\mathcal{B}(E) - \mathcal{B}(E)^{\otimes I}$. By Remark 8.26, it suffices to check this for rectangular cylinders with a finite base $A \in \mathcal{Z}^R$ since $\mathcal{Z}^R$ is a π system that generates $\mathcal{B}(E)^{\otimes I}$. Hence, let $0 = t_0 < t_1 < \dots < t_n$ and $B_0, \dots, B_n \in \mathcal{B}(E)$ as well as $A = \bigcap_{i=0}^n X_{t_i}^{-1}(B_i)$. However, by Corollary 14.24, the following map is measurable,

$$x \mapsto \mathbf{P}_x[A] = \left(\delta_x \otimes \bigotimes_{i=0}^{n-1} \kappa_{t_i, t_{i+1}} \right) \left(\begin{matrix} n \\ \bigcup_{i=0}^n B_i \end{matrix} \right)$$

Corollary 14.43 (Measures by consistent families of kernels) Under the assumptions of Theorem 14.42, for every probability measure μ on E , there exists a unique probability measure $\mathbf{P}_\mu$ on $(E^I, \mathcal{B}(E)^{\otimes I})$ with the following property: For any choice of finitely many numbers $0 = j_0 < j_1 < j_2 < \dots < j_n$ from I , and letting $J := \{j_0, \dots, j_n\}$, we have $\mathbf{P}_\mu \circ X_J^{-1} = \mu \otimes \bigotimes_{k=0}^{n-1} \kappa_{j_k, j_{k+1}}$.

Proof Take $\mathbf{P}_\mu = \int \mu(dx) \kappa(x, \cdot)$.

As a simple conclusion of Lemma 14.41 and Theorem 14.42, we get the following statement that we formulate separately because it will play a central role later.

Corollary 14.44 (Measures via Markov semigroups) Let $(\kappa_t : t \in I)$ be a Markov semigroup on the Polish space E . Then there exists a unique stochastic kernel κ from $(E, \mathcal{B}(E))$ to $(E^I, \mathcal{B}(E)^{\otimes I})$ with the property: For all $x \in E$ and for any choice of finitely many numbers $0 = t_0 < t_1 < t_2 < \dots < t_n$ from I , and letting $J := \{t_0, \dots, t_n\}$, we have

$$\kappa(x, \cdot) \circ X_J^{-1} = \left(\delta_x \otimes \bigotimes_{k=0}^{n-1} \kappa_{t_{k+1}-t_k} \right) \quad (14.16)$$

For any probability measure μ on E , there exists a unique probability measure $\mathbf{P}_\mu$ on $(E^I, \mathcal{B}(E)^{\otimes I})$ with the property: For any choice of finitely many numbers $0 = t_0 < t_1 < t_2 < \dots < t_n$ from I , and letting $J := \{t_0, \dots, t_n\}$, we have $\mathbf{P}_\mu \circ X_J^{-1} = \mu \otimes \bigotimes_{k=0}^{n-1} \kappa_{t_{k+1}-t_k}$. We denote $\mathbf{P}_x = \mathbf{P}_{\delta_x} = \kappa(x, \cdot)$ for $x \in E$.

Example 14.45 (Independent normally distributed increments) Let $I = [0, \infty)$ and $\Omega_i = \mathbb{R}, i \in [0, \infty)$, equipped with the Borel σ -algebra $\mathcal{B} = \mathcal{B}(\mathbb{R})$. Further, let $\Omega = \mathbb{R}^{[0, \infty)}, \mathcal{A} = \mathcal{B}^{\otimes [0, \infty)}$ and let X_t be the coordinate map for $t \in [0, \infty)$. In the sense of Definition 14.6, $X = (X_t)_{t \geq 0}$ is thus the canonical process on $(\Omega, \mathcal{A})$.

We construct a probability measure $\mathbf{P}$ on $(\Omega, \mathcal{A})$ such that the stochastic process X has independent, stationary, normally distributed increments (recall Definition 9.7). That is, it should hold that

$$(X_{t_i} - X_{t_{i-1}})_{i=1,\dots,n} \text{ is independent for all } 0 =: t_0 < t_1 < \dots < t_n, \quad (14.17)$$

$$\mathbf{P}_{X_t - X_s} = \mathcal{N}_{0,t-s} \quad \text{for all } t > s \quad (14.18)$$

To this end, define stochastic kernels $\kappa_t(x, dy) := \delta_x * \mathcal{N}_{0,t}(dy)$ for $t \in [0, \infty)$ where $\mathcal{N}_{0,0} = \delta_0$. By Lemma 14.27, the Chapman-Kolmogorov equation holds since (compare Exercise 14.2.1(i))

$$\kappa_s \cdot \kappa_t(x, dy) = \delta_x * (\mathcal{N}_{0,s} * \mathcal{N}_{0,t})(dy) = \delta_x * \mathcal{N}_{0,s+t}(dy) = \kappa_{s+t}(x, dy)$$

Let $P_0 = \delta_0$ and let $\mathbf{P}$ be the unique probability measure on Ω corresponding to P_0 and $(\kappa_t : t \geq 0)$ according to Corollary 14.44. By Theorem 14.28, Eqs. (14.17) and (14.18) hold.

With $(X_t)_{t \geq 0}$, we have almost constructed the so-called Brownian motion. In addition to the properties we required here, Brownian motion has continuous paths; that is, the maps $t \mapsto X_t$ are almost surely continuous. Note that at this point it is not even clear that the paths are measurable maps. We will have some work to do to establish continuity of the paths, and we will come back to this in Chapter 21.

The construction in the preceding example does not depend on the details of the normal distribution but only on the validity of the convolution equation

$$\mathcal{N}_{0,s+t} = \mathcal{N}_{0,s} * \mathcal{N}_{0,t}$$

Hence, in (14.18) we can replace the normal distribution by any parameterized family of distributions $(v_t, t \geq 0)$ with the property $v_{t+s} = v_t * v_s$. Examples include the Gamma distribution $v_t = \Gamma_{\theta,t}$ (for fixed parameter $\theta > 0$), the Poisson distribution $v_t = \text{Poi}_t$, the negative binomial distribution $v_t = b_{t,p}^-$ (for fixed $p \in (0, 1]$), the Cauchy distribution $v_t = \text{Cau}_t$ and others (compare Theorem 15.12 and Corollary 15.13). We establish the result in a theorem.

Definition 14.46 (Convolution semigroup) Let $I \subset [0, \infty)$ be a semigroup. A family $v = (v_t : t \in I)$ of probability distributions on $\mathbb{R}^d$ is called a convolution semigroup if $v_{s+t} = v_s * v_t$ holds for all $s, t \in I$.

If $I = [0, \infty)$ and if in addition $v_t \xrightarrow{t \rightarrow 0} \delta_0$, then the convolution semigroup is called continuous (in the sense of weak convergence).

If $d = 1$ and $v_t((-\infty, 0)) = 0$ for all $t \in I$, then v is called a nonnegative convolution semigroup.

For the following theorem, compare Definition 9.7.

Theorem 14.47 For any convolution semigroup $(v_t : t \in I)$ and any $x \in \mathbb{R}^d$, there exists a probability measure $\mathbf{P}_x$ on the product space $(\Omega, \mathcal{A}) = \left((\mathbb{R}^d)^I, \mathcal{B}(\mathbb{R}^d)^{\otimes I} \right)$ such that the canonical process $(X_t)_{t \in I}$ is a stochastic process with $\mathbf{P}_x[X_0 = x] = 1$, with stationary independent increments and with $\mathbf{P}_x \circ (X_t - X_s)^{-1} = v_{t-s}$ for $t > s$. On the other hand, every stochastic process $(X_t)_{t \in I}$ (on an arbitrary probability space $(\Omega, \mathcal{A}, \mathbf{P})$) with stationary independent increments defines a convolution semigroup by $v_t = \mathbf{P} \circ (X_t - X_0)^{-1}$ for all $t \in I$.

Exercise 14.4.1 Assume that $(v_t : t \geq 0)$ is a continuous convolution semigroup. Show that $v_t = \lim_{s \rightarrow t} v_s$ for all $t > 0$.

Exercise 14.4.2 Assume that $(v_t : t \geq 0)$ is a convolution semigroup. Show that $v_{t/n} \xrightarrow{n \rightarrow \infty} \delta_0$.

Exercise 14.4.3 Show that a nonnegative convolution semigroup is continuous.

Exercise 14.4.4 Show that a continuous real convolution semigroup $(v_t : t \geq 0)$ with $v_t((-\infty, 0)) = 0$ for some $t > 0$ is nonnegative.

5.15 15 Characteristic Functions and the Central Limit Theorem

The main goal of this chapter is the central limit theorem (CLT) for sums of independent random variables (Theorem 15.37) and for independent arrays of random variables (Lindeberg-Feller theorem, Theorem 15.43). For the latter, we prove only that one of the two implications (Lindeberg's theorem) that is of interest in the applications.

The ideal tools for the treatment of central limit theorems are so-called characteristic functions; that is, Fourier transforms of probability measures. We start with a more general treatment of classes of test functions that are suitable to characterize weak convergence and then study Fourier transforms in greater detail. The subsequent section proves the CLT for real-valued random variables by means of characteristic functions. In the fifth section, we prove a multidimensional version of the CLT.

5.15.1 15.1 Separating Classes of Functions

Let (E, d) be a metric space with Borel σ -algebra $\mathcal{E} = \mathcal{B}(E)$. Denote by $\mathbb{C} = \{u + iv : u, v \in \mathbb{R}\}$ the field of complex numbers. Let

$$\operatorname{Re}(u + iv) = u \quad \text{and} \quad \operatorname{Im}(u + iv) = v$$

denote the real part and the imaginary part, respectively, of $z = u + iv \in \mathbb{C}$. Let $\bar{z} = u - iv$ be the complex conjugate of z and $|z| = \sqrt{u^2 + v^2}$ its modulus. A prominent role will be played by the complex exponential function $\exp : \mathbb{C} \rightarrow \mathbb{C}$, which can be defined either by Euler's formula $\exp(z) = \exp(u)(\cos(v) + i \sin(v))$ or by the power series $\exp(z) = \sum_{n=0}^{\infty} z^n/n!$. It is well-known that $\exp(z_1 + z_2) = \exp(z_1) \cdot \exp(z_2)$. Note that $\operatorname{Re}(z) = (z + \bar{z})/2$ and $\operatorname{Im}(z) = (z - \bar{z})/2i$ imply

$$\cos(x) = \frac{e^{ix} + e^{-ix}}{2} \quad \text{and} \quad \sin(x) = \frac{e^{ix} - e^{-ix}}{2i} \quad \text{for all } x \in \mathbb{R}$$

A map $f : E \rightarrow \mathbb{C}$ is measurable if and only if $\operatorname{Re}(f)$ and $\operatorname{Im}(f)$ are measurable (see Theorem 1.90 with $\mathbb{C} \cong \mathbb{R}^2$). In particular, any continuous function $E \rightarrow \mathbb{C}$ is measurable. If $\mu \in \mathcal{M}(E)$, then we define

$$\int f d\mu := \int \operatorname{Re}(f) d\mu + i \int \operatorname{Im}(f) d\mu$$

if both integrals exist and are finite. Let $C_b(E; \mathbb{C})$ denote the Banach space of continuous, bounded, complex-valued functions on E equipped with the supremum norm $\|f\|_{\infty} = \sup\{|f(x)| : x \in E\}$. We call $\mathcal{C} \subset C_b(E; \mathbb{C})$ a separating class for $\mathcal{M}_f(E)$ if for any two measures $\mu, \nu \in \mathcal{M}_f(E)$ with $\mu \neq \nu$, there is an $f \in \mathcal{C}$ such that $\int f d\mu \neq \int f d\nu$. The analogue of Theorem 13.34 holds for $\mathcal{C} \subset C_b(E; \mathbb{C})$.

Definition 15.1 Let $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{C}$. A subset $\mathcal{C} \subset C_b(E; \mathbb{K})$ is called an algebra if

- (i) $1 \in \mathcal{C}$,
- (ii) if $f, g \in \mathcal{C}$, then $f \cdot g$ and $f + g$ are in $\mathcal{C}$, and
- (iii) if $f \in \mathcal{C}$ and $\alpha \in \mathbb{K}$, then (αf) is in $\mathcal{C}$.

We say that $\mathcal{C}$ separates points if for any two points $x, y \in E$ with $x \neq y$, there is an $f \in \mathcal{C}$ with $f(x) \neq f(y)$.

Theorem 15.2 (Stone-Weierstraß) Let E be a compact Hausdorff space. Let $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{C}$. Let $\mathcal{C} \subset C_b(E; \mathbb{K})$ be an algebra that separates points. If $\mathbb{K} = \mathbb{C}$, then in addition assume that $\mathcal{C}$ is closed under complex conjugation (that is, if $f \in \mathcal{C}$, then the complex conjugate function $\bar{f}$ is also in $\mathcal{C}$).

Then $\mathcal{C}$ is dense in $C_b(E; \mathbb{K})$ with respect to the supremum norm.

Proof We follow the exposition in Dieudonné [34, Chapter VII.3]. First consider the case $\mathbb{K} = \mathbb{R}$. We proceed in several steps.

Step 1. By Weierstraß's approximation theorem (Example 5.15), there is a sequence $(p_n)_{n \in \mathbb{N}}$ of polynomials that approach the map $[0, 1] \rightarrow [0, 1], t \mapsto \sqrt{t}$ uniformly. If $f \in \mathcal{C}$, then also

$$|f| = \|f\|_\infty \lim_{n \rightarrow \infty} p_n(f^2/\|f\|_\infty^2)$$

is in the closure $\bar{\mathcal{C}}$ of $\mathcal{C}$ in $C_b(E; \mathbb{R})$.

Step 2. Applying Step 1 to the algebra $\bar{\mathcal{C}}$ yields that, for all $f, g \in \bar{\mathcal{C}}$,

$$f \vee g = \frac{1}{2}(f + g + |f - g|) \quad \text{and} \quad f \wedge g = \frac{1}{2}(f + g - |f - g|)$$

are also in $\bar{\mathcal{C}}$.

Step 3. For any $f \in C_b(E; \mathbb{R})$, any $x \in E$ and any $\varepsilon > 0$, there exists a $g_x \in \bar{\mathcal{C}}$ with $g_x(x) = f(x)$ and $g_x(y) \leq f(y) + \varepsilon$ for all $y \in E$. As $\mathcal{C}$ separates points, for any $z \in E \setminus \{x\}$, there exists an $H_z \in \mathcal{C}$ with $H_z(z) \neq H_z(x) = 0$. For such z , define $h_z \in \mathcal{C}$ by

$$h_z(y) = f(x) + \frac{f(z) - f(x)}{H_z(z)} H_z(y) \quad \text{for all } y \in E.$$

In addition, define $h_x := f$. Then $h_z(x) = f(x)$ and $h_z(z) = f(z)$ for all $z \in E$. Since f and h_z are continuous, for any $z \in E$, there exists an open neighborhood $U_z \ni z$ with $h(y) \leq f(y) + \varepsilon$ for all $y \in U_z$. We construct a finite covering $U_{z_1}, \dots, U_{z_n}$ of E consisting of such neighborhoods and define $g_x = \min(h_{z_1}, \dots, h_{z_n})$. By Step 2, we have $g_x \in \bar{\mathcal{C}}$.

Step 4. Let $f \in C_b(E; \mathbb{R})$, $\varepsilon > 0$ and, for any $x \in E$, let g_x be as in Step 3. As f and g_x are continuous, for any $x \in E$, there exists an open neighborhood $V_x \ni x$ with $g_x(y) \geq f(y) - \varepsilon$ for any $y \in V_x$. We construct a finite covering $V_{x_1}, \dots, V_{x_n}$ of E and define $g := \max(g_{x_1}, \dots, g_{x_n})$. Then $g \in \bar{\mathcal{C}}$ by Step 2 and $\|g - f\|_\infty < \varepsilon$ by construction. Letting $\varepsilon \downarrow 0$, we get $\bar{\mathcal{C}} = C_b(E; \mathbb{R})$.

Step 5. Now consider $\mathbb{K} = \mathbb{C}$. If $f \in \mathcal{C}$, then by assumption $\operatorname{Re}(f) = (f + \bar{f})/2$ and $\operatorname{Im}(f) = (f - \bar{f})/2i$ are in $\mathcal{C}$. In particular, $\mathcal{C}_0 := \{\operatorname{Re}(f) : f \in \mathcal{C}\} \subset \mathcal{C}$ is a real algebra that, by assumption, separates points and contains the constant functions. Hence $\mathcal{C}_0$ is dense in $C_b(E; \mathbb{R})$. Since $\mathcal{C} = \mathcal{C}_0 + i\mathcal{C}_0$, $\mathcal{C}$ is dense in $C_b(E; \mathbb{C})$.

Corollary 15.3 Let E be a compact metric space. Let $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{C}$. Let $\mathcal{C} \subset C_b(E; \mathbb{K})$ be a family that separates points; that is, stable under multiplication and that contains 1. If $\mathbb{K} = \mathbb{C}$, then in addition assume that $\mathcal{C}$ is closed under complex conjugation.

Then $\mathcal{C}$ is a separating family for $\mathcal{M}_f(E)$.

Proof Let $\mu_1, \mu_2 \in \mathcal{M}_f(E)$ with $\int g d\mu_1 = \int g d\mu_2$ for all $g \in \mathcal{C}$. Let $\mathcal{C}'$ be the algebra of finite linear combinations of elements of $\mathcal{C}$. By linearity of the integral, $\int g d\mu_1 = \int g d\mu_2$ for all $g \in \mathcal{C}'$.

For any $f \in C_b(E, \mathbb{R})$ and any $\varepsilon > 0$, by the Stone-Weierstraß theorem, there exists a $g \in \mathcal{C}'$ with $\|f - g\|_\infty < \varepsilon$. By the triangle inequality,

$$\begin{aligned} \left| \int f d\mu_1 - \int f d\mu_2 \right| &\leq \left| \int f d\mu_1 - \int g d\mu_1 \right| + \left| \int g d\mu_1 - \int g d\mu_2 \right| \\ &\quad + \left| \int g d\mu_2 - \int f d\mu_2 \right| \\ &\leq \varepsilon (\mu_1(E) + \mu_2(E)). \end{aligned}$$

Letting $\varepsilon \downarrow 0$, we get equality of the integrals and hence $\mu_1 = \mu_2$ (by Theorem 13.11).

The following theorems are simple consequences of Corollary 15.3.

Theorem 15.4 The distribution of a bounded real random variable X is characterized by its moments.

Proof Without loss of generality, we can assume that X takes values in $E := [0, 1]$. For $n \in \mathbb{N}$, define the map $f_n : [0, 1] \rightarrow [0, 1]$ by $f_n : x \mapsto x^n$. Further, let $f_0 \equiv 1$. The family $\mathcal{C} = \{f_n, n \in \mathbb{N}_0\}$ separates points and is closed under multiplication; hence it is a separating class for $\mathcal{M}_f(E)$. Thus $\mathbf{P}_X$ is uniquely determined by its moments $\mathbf{E}[X^n] = \int x^n \mathbf{P}_X(dx), n \in \mathbb{N}$.

Example 15.5 (due to [73]) In the preceding theorem, we cannot simply drop the assumption that X is bounded without making other assumptions (see Corollary 15.32). Even if all moments exist, the distribution of X is, in general, not uniquely determined by its moments. As an example consider $X := \exp(Y)$, where $Y \sim \mathcal{N}_{0,1}$. The distribution of X is called the log-normal distribution. For every $n \in \mathbb{N}$, nY is distributed as the sum of n^2 independent, standard normally distributed random variables $nY \stackrel{\mathcal{D}}{=} Y_1 + \dots + Y_{n^2}$. Hence, for $n \in \mathbb{N}$,

$$\begin{aligned} \mathbf{E}[X^n] &= \mathbf{E}[e^{nY}] = \mathbf{E}[e^{Y_1 + \dots + Y_{n^2}}] = \prod_{i=1}^{n^2} \mathbf{E}[e^{Y_i}] = \mathbf{E}[e^Y]^{n^2} \\ &= \left(\int_{-\infty}^{\infty} (2\pi)^{-1/2} e^y e^{-y^2/2} dy \right)^{n^2} = e^{n^2/2} \end{aligned} \quad (15.1)$$

We construct a whole family of distributions with the same moments as X . By the transformation formula for densities (Theorem 1.101), the distribution of X has the density

$$f(x) = \frac{1}{\sqrt{2\pi}} x^{-1} \exp\left(-\frac{1}{2} \log(x)^2\right) \text{ for } x > 0$$

For $\alpha \in [-1, 1]$, define probability densities f_α on $(0, \infty)$ by

$$f_\alpha(x) = f(x)(1 + \alpha \sin(2\pi \log(x))).$$

In order to show that f_α is a density and has the same moments as f , it is enough to show that, for all $n \in \mathbb{N}_0$,

$$m(n) := \int_0^\infty x^n f(x) \sin(2\pi \log(x)) dx = 0$$

With the substitution $y = \log(x) - n$, we get (note that $\sin(2\pi(y+n)) = \sin(2\pi y)$)

$$\begin{aligned} m(n) &= \int_{-\infty}^{\infty} e^{yn+n^2} (2\pi)^{-1/2} e^{-(y+n)^2/2} \sin(2\pi(y+n)) dy \\ &= (2\pi)^{-1/2} e^{n^2/2} \int_{-\infty}^{\infty} e^{-y^2/2} \sin(2\pi y) dy = 0 \end{aligned}$$

where the last equality holds since the integrand is an odd function.

Theorem 15.6 (Laplace transform) A finite measure μ on $[0, \infty)$ is characterized by its Laplace transform

$$\mathcal{L}_\mu(\lambda) := \int e^{-\lambda x} \mu(dx) \quad \text{for } \lambda \geq 0$$

Proof We face the problem that the space $[0, \infty)$ is not compact by passing to the one-point compactification $E = [0, \infty]$. For $\lambda \geq 0$, define the continuous function $f_\lambda : [0, \infty] \rightarrow [0, 1]$ by $f_\lambda(x) = e^{-\lambda x}$ if $x < \infty$ and $f_\lambda(\infty) = \lim_{x \rightarrow \infty} e^{-\lambda x}$. Then $\mathcal{C} = \{f_\lambda, \lambda \geq 0\}$ separates points, $f_0 = 1 \in \mathcal{C}$ and $f_\mu \cdot f_\lambda = f_{\mu+\lambda} \in \mathcal{C}$. By Corollary 15.3, $\mathcal{C}$ is a separating class for $\mathcal{M}_f([0, \infty])$ and thus also for $\mathcal{M}_f([0, \infty))$.

Definition 15.7 For $\mu \in \mathcal{M}_f(\mathbb{R}^d)$, define the map $\varphi_\mu : \mathbb{R}^d \rightarrow \mathbb{C}$ by

$$\varphi_\mu(t) := \int e^{i\langle t, x \rangle} \mu(dx)$$

φ_μ is called the characteristic function of μ .

Theorem 15.8 (Characteristic function) A finite measure $\mu \in \mathcal{M}_f(\mathbb{R}^d)$ is characterized by its characteristic function.

Proof Let $\mu_1, \mu_2 \in \mathcal{M}_f(\mathbb{R}^d)$ with $\varphi_{\mu_1}(t) = \varphi_{\mu_2}(t)$ for all $t \in \mathbb{R}^d$. By Theorem 13.11(ii), $C_c(\mathbb{R}^d)$ is a separating class for $\mathcal{M}_f(\mathbb{R}^d)$. Hence, it is enough to show that $\int f d\mu_1 = \int f d\mu_2$ for all $f \in C_c(\mathbb{R}^d)$.

Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be continuous with compact support and let $\varepsilon > 0$. Assume that $K > 0$ is large enough such that $f(x) = 0$ for $x \notin (-K/2, K/2)^d$ and such that $\mu_i(\mathbb{R}^d \setminus (-K, K)^d) < \varepsilon$, $i = 1, 2$. Consider the torus $E := \mathbb{R}^d / (2K\mathbb{Z}^d)$ and define $\tilde{f} : E \rightarrow \mathbb{R}$ by

$$\tilde{f}(x + 2K\mathbb{Z}^d) = f(x) \quad \text{for } x \in [-K, K]^d.$$

Since the support of f is contained in $(-K, K)^d$, $\tilde{f}$ is continuous. For $m \in \mathbb{Z}^d$ define

$$g_m : \mathbb{R}^d \rightarrow \mathbb{C}, \quad x \mapsto \exp(i\langle \pi m / K, x \rangle).$$

Let $\mathcal{C}$ be the algebra of finite linear combinations of the g_m . For $g \in \mathcal{C}$, we have $g(x) = g(x + 2Kn)$ for all $x \in \mathbb{R}^d$ and $n \in \mathbb{Z}^d$. Hence, the map

$$\tilde{g} : E \rightarrow \mathbb{C}, \quad \tilde{g}(x + 2K\mathbb{Z}^d) = g(x)$$

is well-defined, continuous and bounded. Furthermore, $\tilde{\mathcal{C}} := \{\tilde{g} : g \in \mathcal{C}\} \subset C_b(E; \mathbb{C})$ is an algebra that separates points and is closed under complex conjugation. As E is compact, by the Stone-Weierstraß theorem, there is a $\tilde{g} \in \tilde{\mathcal{C}}$ such that $\|\tilde{g} - \tilde{f}\|_\infty < \varepsilon$. We infer

$$\|(f - g)\mathbb{1}_{[-K, K]^d}\|_\infty < \varepsilon$$

and

$$\|(f - g)\mathbb{1}_{\mathbb{R}^d \setminus [-K, K]^d}\|_\infty \leq \|g\|_\infty = \|\tilde{g}\|_\infty \leq \|\tilde{f}\|_\infty + \varepsilon = \|f\|_\infty + \varepsilon.$$

By assumption of the theorem, $\int g d\mu_1 = \int g d\mu_2$. Hence, using the triangle inequality, we conclude

$$\begin{aligned} \left| \int f d\mu_1 - \int f d\mu_2 \right| &\leq \int |f - g| d\mu_1 + \int |f - g| d\mu_2 \\ &\leq \varepsilon (2\|f\|_\infty + 2\varepsilon + \mu_1(\mathbb{R}^d) + \mu_2(\mathbb{R}^d)) \end{aligned}$$

As $\varepsilon > 0$ was arbitrary, the integrals coincide.

Corollary 15.9 A finite measure μ on $\mathbb{Z}^d$ is uniquely determined by the values

$$\varphi_\mu(t) = \int e^{i\langle t, x \rangle} \mu(dx), \quad t \in [-\pi, \pi]^d$$

Proof This is obvious since $\varphi_\mu(t + 2\pi k) = \varphi_\mu(t)$ for all $k \in \mathbb{Z}^d$.

While the preceding corollary only yields an abstract uniqueness statement, we will profit also from an explicit inversion formula for Fourier transforms.

Theorem 15.10 (Discrete Fourier inversion formula) Let $\mu \in \mathcal{M}_f(\mathbb{Z}^d)$ with characteristic function φ_μ . Then, for every $x \in \mathbb{Z}^d$,

$$\mu(\{x\}) = (2\pi)^{-d} \int_{[-\pi, \pi]^d} e^{-i\langle t, x \rangle} \varphi_\mu(t) dt$$

Proof By the dominated convergence theorem,

$$\begin{aligned} \int_{[-\pi, \pi]^d} e^{-i\langle t, x \rangle} \varphi_\mu(t) dt &= \int_{[-\pi, \pi]^d} e^{-i\langle t, x \rangle} \left(\lim_{n \rightarrow \infty} \sum_{|y| \leq n} e^{i\langle t, y \rangle} \mu(\{y\}) \right) dt \\ &= \lim_{n \rightarrow \infty} \int_{[-\pi, \pi]^d} e^{-i\langle t, x \rangle} \sum_{|y| \leq n} e^{i\langle t, y \rangle} \mu(\{y\}) dt \\ &= \sum_{y \in \mathbb{Z}^d} \mu(\{y\}) \int_{[-\pi, \pi]^d} e^{i\langle t, y-x \rangle} dt \end{aligned}$$

The claim follows since, for $y \in \mathbb{Z}^d$,

$$\int_{[-\pi, \pi]^d} e^{i\langle t, y-x \rangle} dt = \begin{cases} (2\pi)^d, & \text{if } x = y \\ 0, & \text{else.} \end{cases}$$

Similar inversion formulas hold for measures μ on $\mathbb{R}^d$. Particularly simple is the case where μ possesses an integrable density $f := \frac{d\mu}{d\lambda}$ with respect to d -dimensional Lebesgue measure λ . In this case, we have the Fourier inversion formula,

$$f(x) = (2\pi)^{-d} \int_{\mathbb{R}^d} e^{-i\langle t, x \rangle} \varphi_\mu(t) \lambda(dt) \quad (15.2)$$

Furthermore, by Plancherel's theorem, $f \in \mathcal{L}^2(\lambda)$ if and only if $\varphi_\mu \in \mathcal{L}^2(\lambda)$. In this case, $\|f\|_2 = \|\varphi\|_2$.

Since we will not need these statements in the following, we only refer to the standard literature (e.g., [173, Chapter VI.2] or [54, Theorem XV.3.3 and Eq. (XV.3.8)]).

Exercise 15.1.1 Show that, in the Stone-Weierstraß theorem, compactness of E is essential. Hint: Let $E = \mathbb{R}$ and use the fact that $C_b(\mathbb{R}) = C_b(\mathbb{R}; \mathbb{R})$ is not separable. Construct a countable algebra $\mathcal{C} \subset C_b(\mathbb{R})$ that separates points.

Exercise 15.1.2 Let $d \in \mathbb{N}$ and let μ be a finite measure on $[0, \infty)^d$. Show that μ is characterized by its Laplace transform $\mathcal{L}_\mu(\lambda) = \int e^{-\langle \lambda, x \rangle} \mu(dx)$, $\lambda \in [0, \infty)^d$.

Exercise 15.1.3 Show that, under the assumptions of Theorem 15.10, Plancherel's equation holds:

$$\sum_{x \in \mathbb{Z}^d} \mu(\{x\})^2 = (2\pi)^{-d} \int_{[-\pi, \pi]^d} |\varphi_\mu(t)|^2 dt$$

Exercise 15.1.4 (Mellin transform) Let X be a nonnegative real random variable. For $s \geq 0$, define the Mellin transform of $\mathbf{P}_X$ by

$$m_X(s) = \mathbf{E}[X^s]$$

(with values in $[0, \infty]$).

Assume there is an $\varepsilon_0 > 0$ with $m_X(\varepsilon_0) < \infty$ (respectively $m_X(-\varepsilon_0) < \infty$). Show that, for any $\varepsilon > 0$, the distribution $\mathbf{P}_X$ is characterized by the values $m_X(s)$ (respectively $m_X(-s)$), $s \in [0, \varepsilon]$.

Hint: For continuous $f : [0, \infty) \rightarrow [0, \infty)$, let

$$\phi_f(z) = \int_0^\infty t^{z-1} f(t) dt$$

for those $z \in \mathbb{C}$ for which the integral is well-defined. By a standard result of complex analysis if $\phi_f(s) < \infty$ for an $s > 1$, then ϕ_f is holomorphic in $\{z \in \mathbb{C} : \operatorname{Re}(z) \in (1, s)\}$ (and is thus uniquely determined by the values $\phi_f(r), r \in (1, 1 + \varepsilon)$ for any $\varepsilon > 0$). Furthermore, for all $r \in (1, s)$,

$$f(t) = \frac{1}{2\pi i} \int_{-\infty}^{\infty} t^{-(r+i\rho)} \phi_f(r+i\rho) d\rho$$

- (i) Conclude the statement for X with a continuous density.
- (ii) For $\delta > 0$, let $Y_\delta \sim \mathcal{U}_{[1-\delta, 1]}$ be independent of X . Show that XY_δ has a continuous density.
- (iii) Compute m_{XY_δ} , and show that $m_{XY_\delta} \rightarrow m_X$ for $\delta \downarrow 0$.
- (iv) Show that $XY_\delta \implies X$ for $\delta \downarrow 0$.

Exercise 15.1.5 Let X, Y, Z be independent nonnegative random variables such that $\mathbf{P}[Z > 0] > 0$ and such that the Mellin transform $m_{XZ}(s)$ is finite for some $s > 0$.

Show that if $XZ \stackrel{\mathcal{D}}{=} YZ$ holds, then $X \stackrel{\mathcal{D}}{=} Y$.

Exercise 15.1.6 Let μ be a probability measure on $\mathbb{R}$ with integrable characteristic function φ_μ and hence $\varphi_\mu \in \mathcal{L}^1(\lambda)$, where λ is the Lebesgue measure on $\mathbb{R}$. Show that μ is absolutely continuous with bounded continuous density $f = \frac{d\mu}{d\lambda}$ given by

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-itx} \varphi_\mu(t) dt \quad \text{for all } x \in \mathbb{R}$$

Hint: Show this first for the normal distribution $\mathcal{N}_{0,\varepsilon}, \varepsilon > 0$. Then show that $\mu * \mathcal{N}_{0,\varepsilon}$ is absolutely continuous with density f_ε , which converges pointwise to f (as $\varepsilon \rightarrow 0$).

Exercise 15.1.7 Let (Ω, τ) be a separable topological space that satisfies the $T_{3\frac{1}{2}}$ separation axiom: For any closed set $A \subset \Omega$ and any point $x \in \Omega \setminus A$, there exists a continuous function $f : \Omega \rightarrow [0, 1]$ with $f(x) = 0$ and $f(y) = 1$ for all $y \in A$. (Note in particular that every metric space is a $T_{3\frac{1}{2}}$ -space.)

Show that $\sigma(C_b(\Omega)) = \mathcal{B}(\Omega)$; that is, the Borel σ -algebra is generated by the bounded continuous functions $\Omega \rightarrow \mathbb{R}$.

5.15.2 15.2 Characteristic Functions: Examples

Recall that $\operatorname{Re}(z)$ is the real part of $z \in \mathbb{C}$. We collect some simple properties of characteristic functions.

Lemma 15.11 Let X be a random variable with values in $\mathbb{R}^d$ and characteristic function $\varphi_X(t) = \mathbf{E}[e^{i\langle t, X \rangle}]$. Then:

- (i) $|\varphi_X(t)| \leq 1$ for all $t \in \mathbb{R}^d$ and $\varphi_X(0) = 1$.
- (ii) $\varphi_{aX+b}(t) = \varphi_X(at)e^{i\langle b, t \rangle}$ for all $a \in \mathbb{R}$ and $b \in \mathbb{R}^d$.
- (iii) $\mathbf{P}_X = \mathbf{P}_{-X}$ if and only if φ is real-valued.
- (iv) If X and Y are independent, then $\varphi_{X+Y} = \varphi_X \cdot \varphi_Y$.
- (v) $0 \leq 1 - \operatorname{Re}(\varphi_X(2t)) \leq 4(1 - \operatorname{Re}(\varphi_X(t)))$ for all $t \in \mathbb{R}^d$.

Proof (i) and (ii) are trivial.

(iii) $\overline{\varphi_X(t)} = \varphi_X(-t) = \varphi_{-X}(t)$.

(iv) As $e^{i\langle t, X \rangle}$ and $e^{i\langle t, Y \rangle}$ are independent random variables, we have

$$\varphi_{X+Y}(t) = \mathbf{E}[e^{i\langle t, X \rangle} \cdot e^{i\langle t, Y \rangle}] = \mathbf{E}[e^{i\langle t, X \rangle}] \mathbf{E}[e^{i\langle t, Y \rangle}] = \varphi_X(t) \varphi_Y(t).$$

(v) By the addition theorem for trigonometric functions,

$$1 - \cos(\langle 2t, X \rangle) = 2(1 - (\cos(\langle t, X \rangle))^2) \leq 4(1 - \cos(\langle t, X \rangle)).$$

Now take the expectations of both sides.

In the next theorem, we collect the characteristic functions for some of the most important distributions.

Theorem 15.12 (Characteristic functions of some distributions) For some distributions P with density $x \mapsto f(x)$ on $\mathbb{R}$ or weights $P(\{k\}), k \in \mathbb{N}_0$, the characteristic function $\varphi(t)$ is given explicitly in Table 15.1.

Proof (i) (Normal distribution) By Lemma 15.11, it is enough to consider the case $\mu = 0$ and $\sigma^2 = 1$. By virtue of the differentiation lemma (Theorem 6.28) and using partial integration, we get

$$\frac{d}{dt}\varphi(t) = \int_{-\infty}^{\infty} e^{itx} ixe^{-x^2/2} dx = -t\varphi(t)$$

This linear differential equation with initial value $\varphi(0) = 1$ has the unique solution $\varphi(t) = e^{-t^2/2}$.

(ii) (Uniform distribution) This is immediate.

(iii) (Triangle distribution) Note that $\text{Tri}_a = \mathcal{U}_{[-a/2, a/2]} * \mathcal{U}_{[-a/2, a/2]}$; hence

$$\varphi_{\text{Tri}_a}(t) = \varphi_{\mathcal{U}_{[-a/2, a/2]}}(t)^2 = 4 \frac{\sin(at/2)^2}{a^2 t^2} = 2 \frac{1 - \cos(at)}{a^2 t^2}.$$

Here we used the fact that by the addition theorem for trigonometric functions

$$1 - \cos(x) = \sin(x/2)^2 + \cos(x/2)^2 - \cos(x) = 2 \sin(x/2)^2.$$

Table 15.1 Characteristic functions of some distributions

Distribution	Parameter	on	Density / Weights	Char. fct.
Name Symbol				$\varphi(t)$
Normal $\mathcal{N}_{\mu, \sigma^2}$	$\mu \in \mathbb{R}, \sigma^2 > 0$	$\mathbb{R}$	$\frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$	$e^{i\mu t} \cdot e^{-\sigma^2 t^2/2}$
Uniform $\mathcal{U}_{[0, a]}$	$a > 0$	$[0, a]$	$1/a$	$\frac{e^{iat} - 1}{iat}$
Uniform $\mathcal{U}_{[-a, a]}$	$a > 0$	$[-a, a]$	$1/2a$	$\frac{\sin(at)}{at}$
Triangle Tri_a	$a > 0$	$[-a, a]$	$\frac{1}{a}(1 - x /a)^+$	$2 \frac{1 - \cos(at)}{a^2 t^2}$
N.N.	$a > 0$	$\mathbb{R}$	$\frac{1}{a} \frac{1 - \cos(ax)}{ax^2}$	$(1 - t /a)^+$
Gamma $\Gamma_{\theta, r}$	$\theta > 0, r > 0$	$[0, \infty)$	$\frac{\theta^r}{\Gamma(r)} x^{r-1} e^{-\theta x}$	$(1 - it/\theta)^{-r}$
Exponential $\exp_{\theta}$	$\theta > 0$	$[0, \infty)$	$\theta e^{-\theta x}$	$\frac{\theta}{\theta - it}$
Two-sided exponential $\exp_{\theta}^2$	$\theta > 0$	$\mathbb{R}$	$\frac{\theta}{2} e^{-\theta x }$	$\frac{1}{1 + (t/a)^2}$
Cauchy Cau_a	$a > 0$	$\mathbb{R}$	$\frac{1}{a\pi} \frac{1}{1 + (x/a)^2}$	$e^{-a t }$
Binomial $b_{n, p}$	$n \in \mathbb{N}, p \in [0, 1]$	$\{0, \dots, n\}$	$\binom{n}{k} p^k (1-p)^{n-k}$	$((1-p) + pe^{it})^n$
Negative binomial	$r > 0, p \in (0, 1]$	$\mathbb{N}_0$	$\binom{-r}{k} (-1)^k p^r (1-p)^k$	$\left(\frac{p}{1 - (1-p)e^{it}}\right)^r$
$b_{r, p}^-$ Poisson Poi_{λ}	$\lambda > 0$	$\mathbb{N}_0$	$e^{-\lambda} \frac{\lambda^k}{k!}$	$\exp(\lambda(e^{it} - 1))$

(iv) (N.N.) This can either be computed directly or can be deduced from (iii) by using the Fourier inversion formula (Eq. (15.2)).

(v) (Gamma distribution) Again it suffices to consider the case $\theta = 1$. For $0 \leq b < c \leq \infty$ and $t \in \mathbb{R}$, let $\gamma_{b, c, t}$ be the linear path in $\mathbb{C}$ from $b - ibt$ to $c - ict$, let $\delta_{b, t}$ be the linear path from b to $b - ibt$ and let $\epsilon_{c, t}$ be the linear path from $c - ict$ to c . Substituting $z = (1 - it)x$, we get

$$\varphi(t) = \frac{1}{\Gamma(r)} \int_0^\infty x^{r-1} e^{-x} e^{itx} dx = \frac{(1-it)^{-r}}{\Gamma(r)} \int_{\gamma_{0,\infty,t}} z^{r-1} e^{-z} dz.$$

Hence, it suffices to show that $\int_{\gamma_{0,\infty,t}} z^{r-1} \exp(-z) dz = \Gamma(r)$.

The function $z \mapsto z^{r-1} \exp(-z)$ is holomorphic in the right complex plane. Hence, by the residue theorem for $0 < b < c < \infty$,

$$\begin{aligned} \int_b^c x^{r-1} \exp(-x) dx &= \int_{\gamma_{b,c,t}} z^{r-1} \exp(-z) dz \\ &\quad + \int_{\delta_{b,t}} z^{r-1} \exp(-z) dz + \int_{\epsilon_{c,t}} z^{r-1} \exp(-z) dz \end{aligned}$$

Recall that $\int_0^\infty x^{r-1} \exp(-x) dx =: \Gamma(r)$. Hence, it is enough to show that the integrals along $\delta_{b,t}$ and $\epsilon_{c,t}$ vanish if $b \rightarrow 0$ and $c \rightarrow \infty$.

However, $|z^{r-1} \exp(-z)| \leq (1+t^2)^{(r-1)/2} b^{r-1} \exp(-b)$ for $z \in \delta_{b,t}$. As the path $\delta_{b,t}$ has length $b|t|$, we get the estimate

$$\left| \int_{\delta_{b,t}} z^{r-1} e^{-z} dz \right| \leq b^r e^{-b} (1+t^2)^{r/2} \rightarrow 0 \quad \text{for } b \rightarrow 0$$

Similarly,

$$\left| \int_{\epsilon_{c,t}} z^{r-1} e^{-z} dz \right| \leq c^r e^{-c} (1+t^2)^{r/2} \rightarrow 0 \quad \text{for } c \rightarrow \infty$$

(vi) (Exponential distribution) This follows from (v) since $\exp_\theta = \Gamma_{\theta,1}$.

(vii) (Two-sided exponential distribution) If X and Y are independent $\exp_\theta$ distributed random variables, then it is easy to check that $X - Y \sim \exp_\theta^2$. Hence

$$\varphi_{\exp_\theta^2}(t) = \varphi_{\exp_\theta}(t) \varphi_{\exp_\theta}(-t) = \frac{1}{1-it/\theta} \frac{1}{1+it/\theta} = \frac{1}{1+(t/\theta)^2}$$

(viii) (Cauchy distribution) This can either be computed directly using residue calculus or can be inferred from the statement for the two-sided exponential distribution by the Fourier inversion formula (Eq. (15.2)).

(ix) (Binomial distribution) By the binomial theorem,

$$\varphi(t) = \sum_{k=0}^n \binom{n}{k} (1-p)^{n-k} (pe^{it})^k = (1-p+pe^{it})^n.$$

(x) (Negative binomial distribution) By the generalized binomial theorem (Lemma 3.5), for all $x \in \mathbb{C}$ with $|x| < 1$,

$$(1-x)^{-r} = \sum_{k=0}^{\infty} \binom{-r}{k} (-x)^k.$$

Using this formula with $x = (1-p)e^{it}$ gives the claim.

(xi) (Poisson distribution) Clearly,

$$\varphi_{\text{Poi}_\lambda}(t) = \sum_{n=0}^{\infty} e^{-\lambda} \frac{(\lambda e^{it})^n}{n!} = e^{\lambda(e^{it}-1)}.$$

Corollary 15.13 The following convolution formulas hold.

- (i) $\mathcal{N}_{\mu_1, \sigma_1^2} * \mathcal{N}_{\mu_2, \sigma_2^2} = \mathcal{N}_{\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2}$ for $\mu_1, \mu_2 \in \mathbb{R}$ and $\sigma_1^2, \sigma_2^2 > 0$.
- (ii) $\Gamma_{\theta, r} * \Gamma_{\theta, s} = \Gamma_{\theta, r+s}$ for $\theta, r, s > 0$.
- (iii) $\text{Cau}_a * \text{Cau}_b = \text{Cau}_{a+b}$ for $a, b > 0$.
- (iv) $b_{m,p} * b_{n,p} = b_{m+n,p}$ for $m, n \in \mathbb{N}$ and $p \in [0, 1]$.
- (v) $b_{r,p}^- * b_{s,p}^- = b_{r+s,p}^-$ for $r, s > 0$ and $p \in (0, 1]$.
- (vi) $\text{Poi}_\lambda * \text{Poi}_\mu = \text{Poi}_{\lambda+\mu}$ for $\lambda, \mu \geq 0$.

Proof This follows by Theorem 15.12 and by $\varphi_{\mu * \nu} = \varphi_\mu \varphi_\nu$ (Lemma 15.11).

The following theorem gives two simple procedures for calculating the characteristic functions of compound distributions.

Theorem 15.14

- (i) Let $\mu_1, \mu_2, \dots \in \mathcal{M}_f(\mathbb{R}^d)$ and let $p_1, p_2, \dots$ be nonnegative numbers with $\sum_{n=1}^{\infty} p_n \mu_n(\mathbb{R}^d) < \infty$. Then the measure $\mu := \sum_{n=1}^{\infty} p_n \mu_n \in \mathcal{M}_f(\mathbb{R}^d)$ has characteristic function

$$\varphi_\mu = \sum_{n=1}^{\infty} p_n \varphi_{\mu_n} \quad (15.3)$$

- (ii) Let $N, X_1, X_2, \dots$ be independent random variables. Assume $X_1, X_2, \dots$ are identically distributed on $\mathbb{R}^d$ with characteristic function φ_X . Assume N takes values in $\mathbb{N}_0$ and has the probability generating function f_N . Then $Y := \sum_{n=1}^N X_n$ has the characteristic function $\varphi_Y(t) = f_N(\varphi_X(t))$.

- (iii) In particular, if we let $N \sim \text{Poi}_\lambda$ in (ii), then $\varphi_Y(t) = \exp(\lambda(\varphi_X(t) - 1))$.

Proof (i) Define $\nu_n = \sum_{k=1}^n p_k \mu_k$. By the linearity of the integral, $\varphi_{\nu_n} = \sum_{k=1}^n p_k \varphi_{\mu_k}$. By assumption, $\mu = \text{w-lim}_{n \rightarrow \infty} \nu_n$; hence also $\varphi_\mu(t) = \lim_{n \rightarrow \infty} \varphi_{\nu_n}(t)$.

- (ii) Clearly,

$$\begin{aligned} \varphi_Y(t) &= \sum_{n=0}^{\infty} \mathbf{P}[N = n] \mathbf{E}[e^{i\langle t, X_1 + \dots + X_n \rangle}] \\ &= \sum_{n=0}^{\infty} \mathbf{P}[N = n] \varphi_X(t)^n = f_N(\varphi_X(t)) \end{aligned}$$

- (iii) In this special case, $f_N(z) = e^{\lambda(z-1)}$ for $z \in \mathbb{C}$ with $|z| \leq 1$.

Example 15.15 Let $n \in \mathbb{N}$, and assume that the points $0 = a_0 < a_1 < \dots < a_n$ and $1 = y_0 > y_1 > \dots > y_n = 0$ are given. Let $\varphi : \mathbb{R} \rightarrow [0, \infty)$ have the properties that

- $\varphi(a_k) = y_k$ for all $k = 0, \dots, n$ and φ is linearly interpolated between the points a_k ,




Fig. 15.1 The characteristic function φ from Example 15.15 with $n = 4$

- $\varphi(x) = 0$ for $|x| > a_n$, and
- φ is even (that is, $\varphi(x) = \varphi(-x)$).

Assume in addition that the y_k are chosen such that φ is convex on $[0, \infty)$. This is equivalent to the condition that $m_1 \leq m_2 \leq \dots \leq m_n \leq 0$, where $m_k := \frac{y_k - y_{k-1}}{a_k - a_{k-1}}$ is the slope on the k th interval. We want to show that φ is the characteristic function of a probability measure $\mu \in \mathcal{M}_1(\mathbb{R})$.

Define $p_k = a_k(m_{k+1} - m_k)$ for $k = 1, \dots, n$.

Let $\mu_k \in \mathcal{M}_1(\mathbb{R})$ be the distribution on $\mathbb{R}$ with density $\frac{1}{\pi} \frac{1 - \cos(a_k \pi)}{a_k x^2}$. By Theorem 15.12, μ_k has the characteristic function $\varphi_{\mu_k}(t) = \left(1 - \frac{|t|}{a_k}\right)^+$. The characteristic function φ_μ of $\mu := \sum_{k=1}^n p_k \mu_k$ is then

$$\varphi_\mu(t) = \sum_{k=1}^n p_k (1 - |t|/a_k)^+.$$

This is a continuous, symmetric, real function with $\varphi_\mu(0) = 1$. It is linear on each of the intervals $[a_{k-1}, a_k]$. See Fig. 15.1 for an example with $n = 4$. By partial summation, for all $k = 1, \dots, n$ (since $m_{n+1} = 0$),

$$\begin{aligned} \varphi_\mu(a_l) &= \sum_{k=1}^n a_k (m_{k+1} - m_k) \left(1 - \frac{a_l}{a_k}\right)^+ = \sum_{k=l}^n (a_k - a_l) (m_{k+1} - m_k) \\ &= [(a_n - a_l) m_{n+1} - (a_l - a_l) m_l] - \sum_{k=l+1}^n (a_k - a_{k-1}) m_k \\ &= - \sum_{k=l+1}^n (y_k - y_{k-1}) = y_l = \varphi(a_l) \end{aligned}$$

Hence $\varphi_\mu = \varphi$.

Example 15.16 Define the function $\varphi : \mathbb{R} \rightarrow [0, 1]$ for $t \in [-\pi, \pi)$ by $\varphi(t) = 1 - 2|t|/\pi$, and assume φ is periodic (with period 2π). By the discrete Fourier inversion formula (Theorem 15.10), φ is the characteristic function of the probability measure $\mu \in \mathcal{M}_1(\mathbb{Z})$ with $\mu(\{x\}) = (2\pi)^{-1} \int_{-\pi}^{\pi} \cos(tx) \varphi(t) dt$. In fact, in order that μ be a measure (not only a signed measure), we still have to show that all of the masses $\mu(\{x\})$ are nonnegative. Clearly, $\mu(\{0\}) = 0$. For $x \in \mathbb{Z} \setminus \{0\}$, use partial integration to compute the integral,

$$\begin{aligned} \int_{-\pi}^{\pi} \cos(tx) \varphi(t) dt &= 2 \int_0^{\pi} \cos(tx) (1 - 2t/\pi) dt \\ &= \frac{4}{x} \left(1 - \frac{2}{\pi}\right) \sin(\pi x) - \frac{4}{x} \sin(0) + \frac{4}{\pi x} \int_0^{\pi} \sin(tx) dt \\ &= \frac{4}{\pi x^2} (1 - \cos(\pi x)) \end{aligned}$$

Summing up, we have

$$\mu(\{x\}) = \begin{cases} \frac{4}{\pi^2 x^2}, & \text{if } x \text{ is odd} \\ 0, & \text{else} \end{cases}$$

Since $\mu(\mathbb{Z}) = \varphi(0) = 1$, μ is indeed a probability measure.

Example 15.17 Define the function $\psi : \mathbb{R} \rightarrow [0, 1]$ for $t \in [-\pi/2, \pi/2)$ by $\psi(t) = 1 - 2|t|/\pi$. Assume ψ is periodic with period π . If φ is the characteristic function of the measure μ from

the previous example, then clearly $\psi(t) = |\varphi(t)|$. On the other hand, $\psi(t) = \frac{1}{2} + \frac{1}{2}\varphi(2t)$. By Theorem 15.14 and Lemma 15.11(ii), we infer that ψ is the characteristic function of the measure ν with $\nu(A) = \frac{1}{2}\delta_0(A) + \frac{1}{2}\mu(A/2)$ for $A \subset \mathbb{R}$. Hence,

$$v(\{x\}) = \begin{cases} \frac{1}{2}, & \text{if } x = 0 \\ \frac{8}{\pi^2 x^2}, & \text{if } \frac{x}{2} \in \mathbb{Z} \text{ is odd} \\ 0, & \text{else} \end{cases}$$

Example 15.18 Let $\varphi(t) = (1 - 2|t|/\pi)^+$ be the characteristic function of the distribution "N.N." from Theorem 15.12 (with $a = \pi/2$) and let ψ be the characteristic function from the preceding example. Note that $\varphi(t) = \psi(t)$ for $|t| \leq \pi/2$ and $\varphi(t) = 0$ for $|t| > \pi/2$; hence $\varphi^2 = \varphi \cdot \psi$. Now let X, Y, Z be independent real random variables with characteristic functions $\varphi_X = \varphi_Y = \varphi$ and $\varphi_Z = \psi$. Then $\varphi_X \varphi_Y = \varphi_X \varphi_Z$; hence $X + Y \stackrel{\mathcal{D}}{=} X + Z$. However, the distributions of Y and Z do not coincide.

Exercise 15.2.1 Let φ be the characteristic function of the d -dimensional random variable X . Assume that $\varphi(t) = 1$ for some $t \neq 0$. Show that $\mathbf{P}[X \in H_t] = 1$, where

$$\begin{aligned} H_t &= \{x \in \mathbb{R}^d : \langle x, t \rangle \in 2\pi\mathbb{Z}\} \\ &= \{y + z \cdot (2\pi t / \|t\|_2^2) : z \in \mathbb{Z}, y \in \mathbb{R}^d \text{ with } \langle y, t \rangle = 0\}. \end{aligned}$$

Infer that $\varphi(t + s) = \varphi(s)$ for all $s \in \mathbb{R}^d$.

Exercise 15.2.2 Show that there are real random variables X, X' and Y, Y' with the properties (i) $X \stackrel{\mathcal{D}}{=} X'$ and $Y \stackrel{\mathcal{D}}{=} Y'$, (ii) X' and Y' are independent, (iii) $X + Y \stackrel{\mathcal{D}}{=} X' + Y'$, and (iv) X and Y are not independent.

Exercise 15.2.3 Let X be a real random variable with characteristic function φ . X is called lattice distributed if there are $a, d \in \mathbb{R}$ such that $\mathbf{P}[X \in a + d\mathbb{Z}] = 1$. Show that X is lattice distributed if and only if there exists a $u \neq 0$ such that $|\varphi(u)| = 1$.

Exercise 15.2.4 Let X be a real random variable with characteristic function φ . Assume that there is a sequence $(t_n)_{n \in \mathbb{N}}$ of real numbers such that $|t_n| \downarrow 0$ and $|\varphi(t_n)| = 1$ for any n . Show that there exists a $b \in \mathbb{R}$ such that $X = b$ almost surely. If in addition, $\varphi(t_n) = 1$ for all n , then $X = 0$ almost surely.

5.15.3 15.3 Lévy's Continuity Theorem

The main statement of this section is Lévy's continuity theorem (Theorem 15.23). Roughly speaking, it says that a sequence of characteristic functions converges pointwise to a continuous function if and only if the limiting function is a characteristic function and the corresponding probability measures converge weakly. We prepare for the proof of this theorem by assembling some analytic tools.

Lemma 15.19 Let $\mu \in \mathcal{M}_1(\mathbb{R}^d)$ with characteristic function φ . Then

$$|\varphi(t) - \varphi(s)|^2 \leq 2(1 - \operatorname{Re}(\varphi(t - s))) \quad \text{for all } s, t \in \mathbb{R}^d.$$

Proof By the Cauchy-Schwarz inequality,

$$\begin{aligned}
|\varphi(t) - \varphi(s)|^2 &= \left| \int_{\mathbb{R}^d} e^{i\langle t, x \rangle} - e^{i\langle s, x \rangle} \mu(dx) \right|^2 \\
&= \left| \int_{\mathbb{R}^d} (e^{i\langle t-s, x \rangle} - 1) e^{i\langle s, x \rangle} \mu(dx) \right|^2 \\
&\leq \int_{\mathbb{R}^d} |e^{i\langle t-s, x \rangle} - 1|^2 \mu(dx) \cdot \int_{\mathbb{R}^d} |e^{i\langle s, x \rangle}|^2 \mu(dx) \\
&= \int_{\mathbb{R}^d} (e^{i\langle t-s, x \rangle} - 1) (e^{-i\langle t-s, x \rangle} - 1) \mu(dx) \\
&= 2(1 - \operatorname{Re}(\varphi(t-s)))
\end{aligned}$$

Definition 15.20 Let (E, d) be a metric space. A family $(f_i, i \in I)$ of maps $E \rightarrow \mathbb{R}$ is called uniformly equicontinuous if, for every $\varepsilon > 0$, there exists a $\delta > 0$ such that $|f_i(t) - f_i(s)| < \varepsilon$ for all $i \in I$ and all $s, t \in E$ with $d(s, t) < \delta$.

Theorem 15.21 If $\mathcal{F} \subset \mathcal{M}_1(\mathbb{R}^d)$ is a tight family, then $\{\varphi_\mu : \mu \in \mathcal{F}\}$ is uniformly equicontinuous. In particular, every characteristic function is uniformly continuous.

Proof We have to show that, for every $\varepsilon > 0$, there exists a $\delta > 0$ such that, for all $t \in \mathbb{R}^d$, all $s \in \mathbb{R}^d$ with $|t - s| < \delta$ and all $\mu \in \mathcal{F}$, we have

$$|\varphi_\mu(t) - \varphi_\mu(s)| < \varepsilon.$$

As $\mathcal{F}$ is tight, there exists an $N \in \mathbb{N}$ with $\mu([-N, N]^d) > 1 - \varepsilon^2/6$ for all $\mu \in \mathcal{F}$. Furthermore, there exists a $\delta > 0$ such that, for $x \in [-N, N]^d$ and $u \in \mathbb{R}^d$ with $|u| < \delta$, we have $|1 - e^{i\langle u, x \rangle}| < \varepsilon^2/6$. Hence we get for all $\mu \in \mathcal{F}$

$$\begin{aligned}
1 - \operatorname{Re}(\varphi_\mu(u)) &\leq \int_{\mathbb{R}^d} |1 - e^{i\langle u, x \rangle}| \mu(dx) \\
&\leq \frac{\varepsilon^2}{3} + \int_{[-N, N]^d} |1 - e^{i\langle u, x \rangle}| \mu(dx) \\
&\leq \frac{\varepsilon^2}{3} + \frac{\varepsilon^2}{6} = \frac{\varepsilon^2}{2}.
\end{aligned}$$

Thus, for $|t - s| < \delta$ by Lemma 15.19, $|\varphi_\mu(t) - \varphi_\mu(s)| \leq \varepsilon$.

Lemma 15.22 Let (E, d) be a metric space and let $f, f_1, f_2, \dots$ be maps $E \rightarrow \mathbb{R}$ with $f_n \xrightarrow{n \rightarrow \infty} f$ pointwise. If $(f_n)_{n \in \mathbb{N}}$ is uniformly equicontinuous, then f is uniformly continuous and $(f_n)_{n \in \mathbb{N}}$ converges to f uniformly on compact sets; that is, for every compact set $K \subset E$, we have

$$\sup_{s \in K} |f_n(s) - f(s)| \xrightarrow{n \rightarrow \infty} 0$$

Proof Fix $\varepsilon > 0$, and choose $\delta > 0$ such that $|f_n(t) - f_n(s)| < \varepsilon$ for all $n \in \mathbb{N}$ and all $s, t \in E$ with $d(s, t) < \delta$. For these s, t , we thus have

$$|f(s) - f(t)| = \lim_{n \rightarrow \infty} |f_n(s) - f_n(t)| \leq \varepsilon$$

Hence, f is uniformly continuous.

Now let $K \subset E$ be compact. As compact sets are totally bounded, there exists an $N \in \mathbb{N}$ and points $t_1, \dots, t_N \in K$ with $K \subset \bigcup_{i=1}^N B_\delta(t_i)$. Choose $n_0 \in \mathbb{N}$ large enough that $|f_n(t_i) - f(t_i)| \leq \varepsilon$ for all $i = 1, \dots, N$ and $n \geq n_0$.

Now let $s \in K$ and $n \geq n_0$. Choose a t_i with $d(s, t_i) < \delta$. Then

$$|f_n(s) - f(s)| \leq |f_n(s) - f_n(t_i)| + |f_n(t_i) - f(t_i)| + |f(t_i) - f(s)| \leq 3\varepsilon.$$

As $\varepsilon > 0$ was arbitrary, we infer that $f_n \xrightarrow{n \rightarrow \infty} f$ uniformly on K .

A map $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is called partially continuous at $x = (x_1, \dots, x_d)$ if, for any $i = 1, \dots, d$, the map $y_i \mapsto f(x_1, \dots, x_{i-1}, y_i, x_{i+1}, \dots, x_d)$ is continuous at $y_i = x_i$.

Theorem 15.23 (Lévy's continuity theorem) Let $P, P_1, P_2, \dots \in \mathcal{M}_1(\mathbb{R}^d)$ with characteristic functions $\varphi, \varphi_1, \varphi_2, \dots$.

(i) If $P = w - \lim_{n \rightarrow \infty} P_n$, then $\varphi_n \xrightarrow{n \rightarrow \infty} \varphi$ uniformly on compact sets.

(ii) If $\varphi_n \xrightarrow{n \rightarrow \infty} f$ pointwise for some $f : \mathbb{R}^d \rightarrow \mathbb{C}$ that is partially continuous at 0, then there exists a probability measure Q such that $\varphi_Q = f$ and $Q = w - \lim_{n \rightarrow \infty} P_n$.

Proof (i) By the definition of weak convergence, we have $\varphi_n \xrightarrow{n \rightarrow \infty} \varphi$ pointwise. As the family $(P_n)_{n \in \mathbb{N}}$ is tight, by Theorem 15.21, $(\varphi_n)_{n \in \mathbb{N}}$ is uniformly equicontinuous. By Lemma 15.22, this implies uniform convergence on compact sets.

(ii) By Theorem 13.34, it is enough to show that the sequence $(P_n)_{n \in \mathbb{N}}$ is tight. For this purpose, it suffices to show that, for every $k = 1, \dots, n$, the sequence $(P_n^k)_{n \in \mathbb{N}}$ of k th marginal distributions is tight. (Here $P_n^k := P_n \circ \pi_k^{-1}$, where $\pi_k : \mathbb{R}^d \rightarrow \mathbb{R}$ is the projection on the k th coordinate.) Let e_k be the k th unit vector in $\mathbb{R}^d$. Then $\varphi_{P_n^k}(t) = \varphi_n(t e_k)$ is the characteristic function of P_n^k . By assumption, $\varphi_{P_n^k} \xrightarrow{n \rightarrow \infty} f_k$ pointwise for some function f_k that is continuous at 0. We have thus reduced the problem to the one-dimensional situation and will henceforth assume $d = 1$.

As $\varphi_n(0) = 1$ for all $n \in \mathbb{N}$, we have $f(0) = 1$. Define the map $h : \mathbb{R} \rightarrow [0, \infty)$ by $h(x) = 1 - \sin(x)/x$ for $x \neq 0$ and $h(0) = 0$. Clearly, h is continuously differentiable on $\mathbb{R}$. It is easy to see that $\alpha := \inf\{h(x) : |x| \geq 1\} = 1 - \sin(1) > 0$. Now, for $K > 0$, compute (using Markov's inequality and Fubini's theorem)

$$\begin{aligned} P_n([-K, K]^c) &\leq \alpha^{-1} \int_{[-K, K]^c} h(x/K) P_n(dx) \\ &\leq \alpha^{-1} \int_{\mathbb{R}} h(x/K) P_n(dx) \\ &= \alpha^{-1} \int_{\mathbb{R}} \left(\int_0^1 (1 - \cos(tx/K)) dt \right) P_n(dx) \\ &= \alpha^{-1} \int_0^1 \left(\int_{\mathbb{R}} (1 - \cos(tx/K)) P_n(dx) \right) dt \\ &= \alpha^{-1} \int_0^1 (1 - \operatorname{Re}(\varphi_n(t/K))) dt \end{aligned}$$

Using dominated convergence, we conclude that

$$\begin{aligned} \limsup_{n \rightarrow \infty} P_n([-K, K]^c) &\leq \alpha^{-1} \limsup_{n \rightarrow \infty} \int_0^1 (1 - \operatorname{Re}(\varphi_n(t/K))) dt \\ &= \alpha^{-1} \int_0^1 \left(\lim_{n \rightarrow \infty} (1 - \operatorname{Re}(\varphi_n(t/K))) \right) dt \\ &= \alpha^{-1} \int_0^1 (1 - \operatorname{Re}(f(t/K))) dt \end{aligned}$$

As f is continuous and $f(0) = 1$, the last integral converges to 0 for $K \rightarrow \infty$. Hence $(P_n)_{n \in \mathbb{N}}$ is tight.

Applying Lévy's continuity theorem to Example 15.15, we get a theorem of Pólya.

Theorem 15.24 (Pólya) Let $f : \mathbb{R} \rightarrow [0, 1]$ be continuous and even with $f(0) = 1$. Assume that f is convex on $[0, \infty)$. Then f is the characteristic function of a probability measure.

Proof Define f_n by $f_n(k/n) := f(k/n)$ for $k = 0, \dots, n^2$, and assume f_n is linearly interpolated between these points. Furthermore, let f_n be constant to the right of n and for $x < 0$, define

$f_n(x) = f_n(-x)$. This is an approximation of f on $[0, \infty)$ by convex and piecewise linear functions. By Example 15.15, every f_n is a characteristic function of a probability measure μ_n . Clearly, $f_n \xrightarrow{n \rightarrow \infty} f$ pointwise; hence f is the characteristic function of a probability measure $\mu = w - \lim_{n \rightarrow \infty} \mu_n$ on $\mathbb{R}$.

Corollary 15.25 For every $\alpha \in (0, 1]$ and $r > 0$, $\varphi_{\alpha,r}(t) = e^{-|rt|^\alpha}$ is the characteristic function of a symmetric probability measure $\mu_{\alpha,r}$ on $\mathbb{R}$.

Remark 15.26 In fact, $\varphi_{\alpha,r}$ is a characteristic function for every $\alpha \in (0, 2]$ ($\alpha = 2$ corresponds to the normal distribution), see Section 16.2. The distributions $\mu_{\alpha,r}$ are the so-called α -stable distributions (see Definition 16.20): If $X_1, X_2, \dots, X_n$ are independent and $\mu_{\alpha,a}$ -distributed, then $\varphi_{X_1+\dots+X_n}(t) = \varphi_X(t)^n = \varphi_X(n^{1/\alpha}t)$; hence $X_1 + \dots + X_n \stackrel{\mathcal{D}}{=} n^{1/\alpha}X_1$.

The Stone-Weierstraß theorem implies that a characteristic function determines a probability distribution uniquely. Pólya's theorem gives a sufficient condition for a symmetric real function to be a characteristic function. Clearly, that condition is not necessary, as, for example, the normal distribution does not fulfill it. For general education we present Bochner's theorem that formulates a necessary and sufficient condition for a function $\varphi : \mathbb{R}^d \rightarrow \mathbb{C}$ to be the characteristic function of a probability measure.

Definition 15.27 A function $f : \mathbb{R}^d \rightarrow \mathbb{C}$ is called positive semidefinite if, for all $n \in \mathbb{N}$, all $t_1, \dots, t_n \in \mathbb{R}^d$ and all $y_1, \dots, y_n \in \mathbb{C}$, we have

$$\sum_{k,l=1}^n y_k \bar{y}_l f(t_k - t_l) \geq 0,$$

in other words, if the matrix $(f(t_k - t_l))_{k,l=1,\dots,n}$ is positive semidefinite.

Lemma 15.28 If $\mu \in \mathcal{M}_f(\mathbb{R}^d)$ has characteristic function φ , then φ is positive semidefinite.

Proof We have

$$\begin{aligned} \sum_{k,l=1}^n y_k \bar{y}_l \varphi(t_k - t_l) &= \sum_{k,l=1}^n y_k \bar{y}_l \int e^{ix(t_k - t_l)} \mu(dx) \\ &= \int \sum_{k,l=1}^n y_k e^{ixt_k} \overline{y_l e^{ixt_l}} \mu(dx) \\ &= \int \left| \sum_{k=1}^n y_k e^{ixt_k} \right|^2 \mu(dx) \geq 0 \end{aligned}$$

□

In the case $d = 1$, the following theorem goes back to Bochner (1932) [19].

Theorem 15.29 (Bochner) A continuous function $\varphi : \mathbb{R}^d \rightarrow \mathbb{C}$ is the characteristic function of a probability distribution on $\mathbb{R}^d$ if and only if φ is positive semidefinite and $\varphi(0) = 1$.

The statement still holds if $\mathbb{R}^d$ is replaced by a locally compact Abelian group.

Proof For the case $d = 1$ see [19, Section 20, Theorem 23] or [54, Chapter XIX.2, p. 622]. For the general case, see, e.g., [71, p. 293, Theorem 33.3]. □

Exercise 15.3.1 (Compare [50] and [4]) Show that there exist two exchangeable sequences $X = (X_n)_{n \in \mathbb{N}}$ and $Y = (Y_n)_{n \in \mathbb{N}}$ of real random variables with $\mathbf{P}_X \neq \mathbf{P}_Y$ but such that

$$\sum_{k=1}^n X_k \stackrel{\mathcal{D}}{=} \sum_{k=1}^n Y_k \quad \text{for all } n \in \mathbb{N}. \quad (15.4)$$

Hint:

(i) Define the characteristic functions (see Theorem 15.12) $\varphi_1(t) = \frac{1}{1+t^2}$ and $\varphi_2(t) = (1-t/2)^+$. Use Pólya's theorem to show that

$$\psi_1(t) := \begin{cases} \varphi_1(t), & \text{if } |t| \leq 1, \\ \varphi_2(t), & \text{if } |t| > 1, \end{cases}$$

and

$$\psi_2(t) := \begin{cases} \varphi_2(t), & \text{if } |t| \leq 1, \\ \varphi_1(t), & \text{if } |t| > 1, \end{cases}$$

are characteristic functions of probability distributions on $\mathbb{R}$.

(ii) Define independent random variables $X_{n,i}, Y_{n,i}, n \in \mathbb{N}, i = 1, 2$, and $\Theta_n, n \in \mathbb{N}$ such that $X_{n,i}$ has characteristic function φ_i , $Y_{n,i}$ has characteristic function ψ_i and $\mathbf{P}[\Theta_n = 1] = \mathbf{P}[\Theta_n = -1] = \frac{1}{2}$. Define $X_n = X_{n,\Theta_n}$ and $Y_n = Y_{n,\Theta_n}$. Show that (15.4) holds.

(iii) Determine $\mathbf{E}[e^{it_1 X_1 + it_2 X_2}]$ and $\mathbf{E}[e^{it_1 Y_1 + it_2 Y_2}]$ for $t_1 = \frac{1}{2}$ and $t_2 = 2$. Conclude that $(X_1, X_2) \neq (Y_1, Y_2)$ and thus $\mathbf{P}_X \neq \mathbf{P}_Y$.

Exercise 15.3.2 Show that for any $\delta > 0$ and $\varepsilon > 0$, there is a $C < \infty$ such that for any $\mu \in \mathcal{M}_1(\mathbb{R})$ with characteristic function φ , we have

$$\mu([- \delta, \delta]^c) \leq C \int_0^\varepsilon (1 - \operatorname{Re}(\varphi(t))) dt$$

For $\varepsilon \delta \leq 3$ one can choose $C = 12/\delta^2 \varepsilon^3$.

Hint: Proceed as in the proof of Lévy's continuity theorem.

Exercise 15.3.3 Let $(\mu_n)_{n \in \mathbb{N}}$ be a sequence of probability measures on $\mathbb{R}$ and denote by $(\varphi_n)_{n \in \mathbb{N}}$ the corresponding characteristic functions. Assume that $\varphi_n(t) \xrightarrow{n \rightarrow \infty} 1$ for t in a neighborhood of 0. Use Exercise 15.3.2 to show that $\mu_n \xrightarrow{n \rightarrow \infty} \delta_0$.

5.15.4 15.4 Characteristic Functions and Moments

We want to study the connection between the moments of a real random variable X and the derivatives of its characteristic function φ_X . We start with a simple lemma.

Lemma 15.30 For $t \in \mathbb{R}$ and $n \in \mathbb{N}$, we have

$$\left| e^{it} - 1 - \frac{it}{1!} - \dots - \frac{(it)^{n-1}}{(n-1)!} \right| \leq \frac{|t|^n}{n!}.$$

Proof As the n th derivative of e^{it} has modulus 1, this follows by Taylor's formula. $\square$

Theorem 15.31 (Moments and differentiability) Let X be a real random variable with characteristic function φ .

(i) If $\mathbf{E}[|X|^n] < \infty$, then φ is n -times continuously differentiable with derivatives

$$\varphi^{(k)}(t) = \mathbf{E}[(iX)^k e^{itX}] \text{ for } k = 0, \dots, n.$$

(ii) In particular, if $\mathbf{E}[X^2] < \infty$, then

$$\varphi(t) = 1 + it\mathbf{E}[X] - \frac{1}{2}t^2\mathbf{E}[X^2] + \varepsilon(t)t^2$$

with $\varepsilon(t) \rightarrow 0$ for $t \rightarrow 0$.

(iii) Let $h \in \mathbb{R}$. If $\lim_{n \rightarrow \infty} \frac{|h|^n \mathbf{E}[|X|^n]}{n!} = 0$, then, for every $t \in \mathbb{R}$,

$$\varphi(t+h) = \sum_{k=0}^{\infty} \frac{(ih)^k}{k!} \mathbf{E} [e^{itX} X^k].$$

In particular, this holds if $\mathbf{E} [e^{|hX|}] < \infty$.

Proof (i) For $t \in \mathbb{R}$, $h \in \mathbb{R} \setminus \{0\}$ and $k \in \{1, \dots, n\}$, define

$$Y_k(t, h, x) = k! h^{-k} e^{itx} \left(e^{ihx} - \sum_{l=0}^{k-1} \frac{(ihx)^l}{l!} \right).$$

Then

$$\mathbf{E} [Y_k(t, h, X)] = k! h^{-k} \left(\varphi(t+h) - \varphi(t) - \sum_{l=1}^{k-1} \mathbf{E} [e^{itX} (iX)^l] \frac{h^l}{l!} \right).$$

If the limit $\varphi_k(t) := \lim_{h \rightarrow 0} \mathbf{E} [Y_k(t, h, X)]$ exists, then φ is k -times differentiable at t with $\varphi^{(k)}(t) = \varphi_k(t)$.

However (by Lemma 15.30 with $n = k+1$), $Y_k(t, h, x) \xrightarrow{h \rightarrow 0} (ix)^k e^{itx}$ for all $x \in \mathbb{R}$ and (by Lemma 15.30 with $n = k$) $|Y_k(t, h, x)| \leq |x|^k$. As $\mathbf{E} [|X|^k] < \infty$ by assumption, the dominated convergence theorem implies

$$\mathbf{E} [Y_k(t, h, X)] \xrightarrow{h \rightarrow 0} \mathbf{E} [(iX)^k e^{itX}] = \varphi^{(k)}(t).$$

Applying the continuity lemma (Theorem 6.27) yields that $\varphi^{(k)}$ is continuous.

(ii) This is a direct consequence of (i).

(iii) By assumption,

$$\begin{aligned} \left| \varphi(t+h) - \sum_{k=0}^{n-1} \frac{(ih)^k}{k!} \mathbf{E} [e^{itX} X^k] \right| &= \frac{h^n}{n!} |\mathbf{E} [Y_n(t, h, X)]| \\ &\leq \frac{h^n \mathbf{E} [|X|^n]}{n!} \xrightarrow{n \rightarrow \infty} 0. \end{aligned}$$

Corollary 15.32 (Method of moments) Let X be a real random variable with

$$\alpha := \limsup_{n \rightarrow \infty} \frac{1}{n} \mathbf{E} [|X|^n]^{1/n} < \infty$$

Then the characteristic function φ of X is analytic and the distribution of X is uniquely determined by the moments $\mathbf{E} [X^n]$, $n \in \mathbb{N}$. In particular, this holds if $\mathbf{E} [e^{t|X|}] < \infty$ for some $t > 0$.

Proof By Stirling's formula,

$$\lim_{n \rightarrow \infty} \frac{1}{n!} n^n e^{-n} \sqrt{2\pi n} = 1$$

Thus, for $|h| < 1/(3\alpha)$,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \mathbf{E} [|X|^n] \cdot |h|^n / n! &= \limsup_{n \rightarrow \infty} \sqrt{2\pi n} \left(\mathbf{E} [|X|^n]^{1/n} \cdot |h| \cdot e/n \right)^n \\ &\leq \limsup_{n \rightarrow \infty} \sqrt{2\pi n} (e/3)^n = 0 \end{aligned}$$

Hence the characteristic function can be expanded about any point $t \in \mathbb{R}$ in a power series with radius of convergence at least $1/(3\alpha)$. In particular, it is analytic and is hence determined by the coefficients of its power series about $t = 0$; that is, by the moments of X .

Example 15.33

(i) Let $X \sim \mathcal{N}_{\mu, \sigma^2}$. Then, for every $t \in \mathbb{R}$,

$$\begin{aligned}\mathbf{E}[e^{tX}] &= (2\pi\sigma^2)^{-1/2} \int_{-\infty}^{\infty} e^{tx} e^{-(x-\mu)^2/2\sigma^2} dx \\ &= e^{\mu t + t^2\sigma^2/2} (2\pi\sigma^2)^{-1/2} \int_{-\infty}^{\infty} e^{-(x-\mu-t\sigma^2)^2/2\sigma^2} dx \\ &= e^{\mu t + t^2\sigma^2/2} < \infty\end{aligned}$$

Hence the distribution of X is characterized by its moments. The characteristic function $\varphi(t) = e^{i\mu t} e^{-\sigma^2 t^2/2}$ that we get by the above calculation with t replaced by it is indeed analytic.

(ii) Let X be exponentially distributed with parameter $\theta > 0$. Then, for $t \in (0, \theta)$,

$$\mathbf{E}[e^{tX}] = \theta \int_0^{\infty} e^{tx} e^{-\theta x} dx = \frac{\theta}{\theta - t} < \infty$$

Hence the distribution of X is characterized by its moments. The above calculation with t replaced by it yields $\varphi(t) = \theta/(\theta - it)$, and this function is indeed analytic. The fact that in the complex plane φ has a singularity at $t = -i\theta$ implies that the power series of φ about 0 has radius of convergence θ . In particular, this implies that not all exponential moments are finite. This is reflected by the above calculation that shows that, for $t \geq \theta$, the exponential moments are infinite.

(iii) Let X be log-normally distributed (see Example 15.5). Then $\mathbf{E}[X^n] = e^{n^2/2}$. In particular, here $\alpha = \infty$. In fact, in Example 15.5, we saw that here the moments do not determine the distribution of X .

(iv) If X takes values in $\mathbb{N}_0$ and if $\beta := \limsup_{n \rightarrow \infty} \mathbf{E}[X^n]^{1/n} < 1$, then by Hadamard's criterion $\psi_X(z) := \sum_{k=1}^{\infty} \mathbf{P}[X = k] z^k < \infty$ for $|z| < 1/\beta$. In particular, the probability generating function X is characterized by its derivatives $\psi_X^{(n)}(1)$, $n \in \mathbb{N}$, and thus by the moments of X . Compare Theorem 3.2(iii).

Theorem 15.34 Let X be a real random variable and let φ be its characteristic function. Let $n \in \mathbb{N}$, and assume that φ is $2n$ -times differentiable at 0 with derivative $\varphi^{(2n)}(0)$. Then $\mathbf{E}[X^{2n}] = (-1)^n \varphi^{(2n)}(0) < \infty$.

Proof We carry out the proof by induction on $n \in \mathbb{N}_0$. For $n = 0$, the claim is trivially true. Now, let $n \in \mathbb{N}$, and assume φ is $2n$ -times (not necessarily continuously) differentiable at 0. Define $u(t) = \operatorname{Re}(\varphi(t))$. Then u is also $2n$ -times differentiable at 0 and $u^{(2k-1)}(0) = 0$ for $k = 1, \dots, n$ since u is even. Since $\varphi^{(2n)}(0)$ exists, $\varphi^{(2n-1)}$ is continuous at 0 and $\varphi^{(2n-1)}(t)$ exists for all $t \in (-\varepsilon, \varepsilon)$ for some $\varepsilon > 0$. Furthermore, $\varphi^{(k)}$ exists in $(-\varepsilon, \varepsilon)$ and is continuous on $(-\varepsilon, \varepsilon)$ for any $k = 0, \dots, 2n - 2$. By Taylor's formula, for every $t \in (-\varepsilon, \varepsilon)$,

$$\left| u(t) - \sum_{k=0}^{n-1} u^{(2k)}(0) \frac{t^{2k}}{(2k)!} \right| \leq \frac{|t|^{2n-1}}{(2n-1)!} \sup_{\theta \in (0,1]} |u^{(2n-1)}(\theta t)|. \quad (15.5)$$

Define a continuous function $f_n : \mathbb{R} \rightarrow [0, \infty)$ by $f_n(0) = 1$ and

$$f_n(x) = (-1)^n (2n)! x^{-2n} \left[\cos(x) - \sum_{k=0}^{n-1} (-1)^k \frac{x^{2k}}{(2k)!} \right] \quad \text{for } x \neq 0.$$

By the induction hypothesis, $\mathbf{E}[X^{2k}] = u^{(2k)}(0)$ for all $k = 1, \dots, n-1$. Using (15.5), we infer

$$\begin{aligned}\mathbf{E}[f_n(tX)X^{2n}] &\leq \frac{2n}{|t|} \sup_{\theta \in (0,1]} |u^{(2n-1)}(\theta t)| \leq g_n(t) \\ &:= 2n \sup_{\theta \in (0,1]} \frac{|u^{(2n-1)}(\theta t)|}{\theta|t|}\end{aligned}$$

Now Fatou's lemma implies

$$\begin{aligned}\mathbf{E}[X^{2n}] &= \mathbf{E}[f_n(0)X^{2n}] \leq \liminf_{t \rightarrow 0} \mathbf{E}[f_n(tX)X^{2n}] \\ &\leq \liminf_{t \rightarrow 0} g_n(t) = 2n |u^{(2n)}(0)| < \infty.\end{aligned}$$

By Theorem 15.31, this implies $\mathbf{E}[X^{2n}] = (-1)^n u^{(2n)}(0) = (-1)^n \varphi^{(2n)}(0)$.

Remark 15.35 For odd moments, the statement of the theorem may fail (see, e.g., Exercise 15.4.4 for the first moment). Indeed, φ is differentiable at 0 with derivative im for some $m \in \mathbb{R}$ if and only if $x\mathbf{P}[|X| > x] \xrightarrow{x \rightarrow \infty} 0$ and $\mathbf{E}[X \mathbb{1}_{\{|X| \leq x\}}] \xrightarrow{x \rightarrow \infty} m$. (See [54, Chapter XVII.2a, p.565].)

Exercise 15.4.1 Let X and Y be nonnegative random variables with

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \mathbf{E}[|X|^n]^{1/n} < \infty, \quad \limsup_{n \rightarrow \infty} \frac{1}{n} \mathbf{E}[|Y|^n]^{1/n} < \infty$$

and

$$\mathbf{E}[X^m Y^n] = \mathbf{E}[X^m] \mathbf{E}[Y^n] \quad \text{for all } m, n \in \mathbb{N}_0$$

Show that X and Y are independent.

Hint: Consider the random variable Y with respect to the probability measure $X^m \mathbf{P}[\cdot] / \mathbf{E}[X^m]$, and use Corollary 15.32 to show that

$$\mathbf{E}[X^m \mathbb{1}_A(Y)] / \mathbf{E}[X^m] = \mathbf{P}[Y \in A] \quad \text{for all } A \in \mathcal{B}(R) \text{ and } m \in \mathbb{N}_0$$

Now apply Corollary 15.32 to the random variable X with respect to the probability measure $\mathbf{P}[\cdot | Y \in A]$.

Exercise 15.4.2 Let $r, s > 0$ and let $Z \sim \Gamma_{1,r+s}$ and $B \sim \beta_{r,s}$ be independent (see Example 1.107). Use Exercise 15.4.1 to show that the random variables $X := BZ$ and $Y := (1 - B)Z$ are independent with $X \sim \Gamma_{1,r}$ and $Y \sim \Gamma_{1,s}$.

Exercise 15.4.3 Show that, for $\alpha > 2$, the function $\phi_\alpha(t) = e^{-|t|^\alpha}$ is not a characteristic function.

Hint: Assume the contrary and show that the corresponding random variable would have variance zero.

Exercise 15.4.4 Let $X_1, X_2, \dots$ be i.i.d. real random variables with characteristic function φ . Show the following.

(i) If φ is differentiable at 0, then $\varphi'(0) = im$ for some $m \in \mathbb{R}$.

(ii) φ is differentiable at 0 with $\varphi'(0) = im$ if and only if $(X_1 + \dots + X_n)/n \xrightarrow{n \rightarrow \infty} m$ in probability.

(iii) Assume that φ is differentiable at 0 and that $X_1 \geq 0$ almost surely. Then $\mathbf{E}[X_1] = -i\varphi'(0) < \infty$.

Hint: Use (ii) and the law of large numbers.

(iv) The distribution of X_1 can be chosen such that φ is differentiable at 0 but $\mathbf{E}[|X_1|] = \infty$.

Exercise 15.4.5 Let $X_1, X_2, \dots$ be real random variables. For $r > 0$ let $M_r(X_n) = \mathbf{E}[|X_n|^r]$ be the r th absolute moment. For $k \in \mathbb{N}$ let $m_k(X_n) = \mathbf{E}[X_n^k]$ be the k th moment if $M_k(X_n) < \infty$.

(i) Assume that X is a real random variable and that $(X_{n_l})_{l \in \mathbb{N}}$ is a subsequence such that

$$\mathbf{P}_{X_{n_l}} \xrightarrow{l \rightarrow \infty} \mathbf{P}_X \quad \text{weakly.}$$

Assume further that there is an $r > 0$ such that $\sup_{n \in \mathbb{N}} M_r(X_n) < \infty$. Show that for any $k \in \mathbb{N} \cap (0, r)$ and $s \in (0, r)$ we have $M_s(X) < \infty$ as well as

$$M_s(X_{n_l}) \xrightarrow{l \rightarrow \infty} M_s(X) \quad \text{and} \quad m_k(X_{n_l}) \xrightarrow{l \rightarrow \infty} m_k(X).$$

(ii) Assume that for any $k \in \mathbb{N}$ the limit

$$m_k := \lim_{n \rightarrow \infty} m_k(X_n)$$

exists and is finite (note that finitely many of the $m_k(X_n)$ may be undefined for any k). Show that there exists a real random variable X with $m_k = m_k(X)$ for all $k \in \mathbb{N}$ and a subsequence $(X_{n_l})_{l \in \mathbb{N}}$ such that

$$\mathbf{P}_{X_{n_l}} \xrightarrow{l \rightarrow \infty} \mathbf{P}_X \quad \text{weakly.}$$

(iii) Show the theorem of Fréchet-Shohat: If in (ii) the distribution of X is determined by its moments $m_k(X), k \in \mathbb{N}$ (see Corollary 15.32), then

$$\mathbf{P}_{X_n} \xrightarrow{n \rightarrow \infty} \mathbf{P}_X \quad \text{weakly.}$$

Exercise 15.4.6 Let $X_1, X_2, \dots$ be i.i.d. real random variables with $\mathbf{E}[X_1] = 0$ and $\mathbf{E}[|X_1|^k] < \infty$ for all $k \in \mathbb{N}$.

(i) Show that there exist finite numbers $(d_k)_{k \in \mathbb{N}}$ (depending on the distribution $\mathbf{P}_{X_1}$) such that for any $k, n \in \mathbb{N}$ we have

$$\left| \mathbf{E}[(X_1 + \dots + X_n)^{2k-1}] \right| \leq d_{2k-1} n^{k-1}$$

and

$$\left| \mathbf{E}[(X_1 + \dots + X_n)^{2k}] - \frac{(2k)!}{2^k k!} \mathbf{E}[X_1^2]^k n^k \right| \leq d_{2k} n^{k-1}.$$

Hint: Expand the bracket expression, sort the terms by the different mixed moments and compute by combinatorial means the number of each type of summand. The number of summands of the type $\mathbf{E}[X_{l_1}^2 \dots X_{l_k}^2]$ (for different $l_1, \dots, l_k$) is of particular importance.

(ii) Let $Y \sim \mathcal{N}_{0,1}$. Use Theorem 15.31(i) to show that for any $k \in \mathbb{N}$ we have

$$\mathbf{E}[Y^{2k-1}] = 0 \quad \text{and} \quad \mathbf{E}[Y^{2k}] = \frac{(2k)!}{2^k k!}$$

(iii) Let $S_n^* = (X_1 + \dots + X_n) / \sqrt{n \operatorname{Var}[X_1]}$. Use Exercise 15.4.5 to infer the statement of the central limit theorem (compare Theorem 15.37)

$$\mathbf{P}_{S_n^*} \xrightarrow{n \rightarrow \infty} \mathcal{N}_{0,1} \quad \text{weakly.}$$

5.15.5 15.5 The Central Limit Theorem

In the strong law of large numbers, we saw that, for large n , the order of magnitude of the sum $S_n = X_1 + \dots + X_n$ of i.i.d. integrable random variables is $n \cdot \mathbf{E}[X_1]$. Of course, for any n , the actual value of S_n will sometimes be smaller than $n \cdot \mathbf{E}[X_1]$ and sometimes larger. In the central limit theorem (CLT), we study the size and shape of the typical fluctuations around $n \cdot \mathbf{E}[X_1]$ in the case where the X_i have a finite variance.

We prepare for the proof of the CLT with a lemma.

Lemma 15.36 Let $X_1, X_2, \dots$ be i.i.d. real random variables with $\mathbf{E}[X_1] = \mu$ and $\text{Var}[X_1] = \sigma^2 \in (0, \infty)$. Let

$$S_n^* := \frac{1}{\sqrt{n\sigma^2}} \sum_{k=1}^n (X_k - \mu)$$

be the normalized n th partial sum. Then

$$\lim_{n \rightarrow \infty} \varphi_{S_n^*}(t) = e^{-t^2/2} \quad \text{for all } t \in \mathbb{R}$$

Proof Let $\varphi = \varphi_{X_k - \mu}$. Then, by Theorem 15.31(ii),

$$\varphi(t) = 1 - \frac{\sigma^2}{2}t^2 + \varepsilon(t)t^2$$

where the error term $\varepsilon(t)$ goes to 0 if $t \rightarrow 0$. By Lemma 15.11(iv) and (ii),

$$\varphi_{S_n^*}(t) = \varphi\left(\frac{t}{\sqrt{n\sigma^2}}\right)^n$$

Now $\left(1 - \frac{t^2}{2n}\right)^n \xrightarrow{n \rightarrow \infty} e^{-t^2/2}$ and

$$\begin{aligned} \left| \left(1 - \frac{t^2}{2n}\right)^n - \varphi\left(\frac{t}{\sqrt{n\sigma^2}}\right)^n \right| &\leq n \left| 1 - \frac{t^2}{2n} - \varphi\left(\frac{t}{\sqrt{n\sigma^2}}\right) \right| \\ &\leq n \frac{t^2}{n\sigma^2} \left| \varepsilon\left(\frac{t}{\sqrt{n\sigma^2}}\right) \right| \xrightarrow{n \rightarrow \infty} 0. \end{aligned}$$

(Note that $|u^n - v^n| \leq |u - v| \cdot n \cdot \max(|u|, |v|)^{n-1}$ for all $u, v \in \mathbb{C}$.) $\square$

Theorem 15.37 (Central limit theorem (CLT)) Let $X_1, X_2, \dots$ be i.i.d. real random variables with $\mu := \mathbf{E}[X_1] \in \mathbb{R}$ and $\sigma^2 := \text{Var}[X_1] \in (0, \infty)$. For $n \in \mathbb{N}$, let $S_n^* := \frac{1}{\sqrt{\sigma^2 n}} \sum_{i=1}^n (X_i - \mu)$. Then

$$\mathbf{P}_{S_n^*} \xrightarrow{n \rightarrow \infty} \mathcal{N}_{0,1} \quad \text{weakly.}$$

For $-\infty \leq a < b \leq +\infty$, we have $\lim_{n \rightarrow \infty} \mathbf{P}[S_n^* \in [a, b]] = \frac{1}{\sqrt{2\pi}} \int_a^b e^{-x^2/2} dx$.

Proof By Lemma 15.36 and Lévy's continuity theorem (Theorem 15.23), $(\mathbf{P}_{S_n^*})$ converges to the distribution with characteristic function $\varphi(t) = e^{-t^2/2}$. By Theorem 15.12(i), this is $\mathcal{N}_{0,1}$. The additional claim follows by the Portemanteau theorem (Theorem 13.16) since $\mathcal{N}_{0,1}$ has a density; hence $\mathcal{N}_{0,1}(\partial[a, b]) = 0$. $\square$

Remark 15.38 If we prefer to avoid the continuity theorem, we could argue as follows: For every $K > 0$ and $n \in \mathbb{N}$, we have $\mathbf{P}[|S_n^*| > K] \leq \text{Var}[S_n^*]/K^2 = 1/K^2$; hence the sequence $(\mathbf{P}_{S_n^*})$ is tight. As characteristic functions determine distributions, the claim follows by Theorem 13.34. $\square$

We want to weaken the assumption in Theorem 15.37 that the random variables are identically distributed. In fact, we can even take a different set of summands for every n . The essential assumptions are that the summands are independent, each summand contributes only a little to the sum and the sum is centered and has variance 1.

Definition 15.39 For every $n \in \mathbb{N}$, let $k_n \in \mathbb{N}$ and let $X_{n,1}, \dots, X_{n,k_n}$ be real random variables. We say that $(X_{n,l}) = (X_{n,l}, l = 1, \dots, k_n, n \in \mathbb{N})$ is an array of random variables. Its row sum is denoted by $S_n = X_{n,1} + \dots + X_{n,k_n}$. The array is called

- independent if, for every $n \in \mathbb{N}$, the family $(X_{n,l})_{l=1, \dots, k_n}$ is independent,

- centered if $X_{n,l} \in \mathcal{L}^1(\mathbf{P})$ and $\mathbf{E}[X_{n,l}] = 0$ for all n and l , and
- normed if $X_{n,l} \in \mathcal{L}^2(\mathbf{P})$ and $\sum_{l=1}^{k_n} \text{Var}[X_{n,l}] = 1$ for all $n \in \mathbb{N}$.

A centered array is called a null array if its individual components are asymptotically negligible in the sense that, for all $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \max_{1 \leq l \leq k_n} \mathbf{P}[|X_{n,l}| > \varepsilon] = 0.$$

Definition 15.40 A centered array of random variables $(X_{n,l})$ with $X_{n,l} \in \mathcal{L}^2(\mathbf{P})$ for every $n \in \mathbb{N}$ and $l = 1, \dots, k_n$ is said to satisfy the Lindeberg condition if, for all $\varepsilon > 0$,

$$L_n(\varepsilon) := \frac{1}{\text{Var}[S_n]} \sum_{l=1}^{k_n} \mathbf{E} \left[X_{n,l}^2 \mathbb{1}_{\{X_{n,l}^2 > \varepsilon^2 \text{Var}[S_n]\}} \right] \xrightarrow{n \rightarrow \infty} 0. \quad (15.6)$$

The array fulfills the Lyapunov condition if there exists a $\delta > 0$ such that

$$\lim_{n \rightarrow \infty} \frac{1}{\text{Var}[S_n]^{1+(\delta/2)}} \sum_{l=1}^{k_n} \mathbf{E} \left[|X_{n,l}|^{2+\delta} \right] = 0. \quad (15.7)$$

Lemma 15.41 The Lyapunov condition implies the Lindeberg condition.

Proof For $x \in \mathbb{R}$, we have $x^2 \mathbb{1}_{\{|x| > \varepsilon'\}} \leq (\varepsilon')^{-\delta} |x|^{2+\delta} \mathbb{1}_{\{|x| > \varepsilon'\}} \leq (\varepsilon')^{-\delta} |x|^{2+\delta}$. Letting $\varepsilon' := \varepsilon \sqrt{\text{Var}[S_n]}$, we get

$$L_n(\varepsilon) \leq \varepsilon^{-\delta} \frac{1}{\text{Var}[S_n]^{1+(\delta/2)}} \sum_{l=1}^{k_n} \mathbf{E} \left[|X_{n,l}|^{2+\delta} \right].$$

Example 15.42 Let $(Y_n)_{n \in \mathbb{N}}$ be i.i.d. with $\mathbf{E}[Y_n] = 0$ and $\text{Var}[Y_n] = 1$. Let $k_n = n$ and $X_{n,l} = \frac{Y_l}{\sqrt{n}}$. Then $(X_{n,l})$ is independent, centered and normed. Clearly, $\mathbf{P}[|X_{n,l}| > \varepsilon] = \mathbf{P}[|Y_1| > \sqrt{\varepsilon n}] \xrightarrow{n \rightarrow \infty} 0$; hence $(X_{n,l})$ is a null array. Furthermore, $L_n(\varepsilon) = \mathbf{E} \left[Y_1^2 \mathbb{1}_{\{|Y_1| > \varepsilon \sqrt{n}\}} \right] \xrightarrow{n \rightarrow \infty} 0$; hence $(X_{n,l})$ satisfies the Lindeberg condition.

If $Y_1 \in \mathcal{L}^{2+\delta}(\mathbf{P})$ for some $\delta > 0$, then

$$\sum_{l=1}^n \mathbf{E} \left[|X_{n,l}|^{2+\delta} \right] = n^{-(\delta/2)} \mathbf{E} \left[|Y_1|^{2+\delta} \right] \xrightarrow{n \rightarrow \infty} 0.$$

In this case, $(X_{n,l})$ also satisfies the Lyapunov condition.

The following theorem is due to Lindeberg (1922, see [108]) for the implication (i) $\implies$ (ii) and is attributed to Feller (1935 and 1937, see [51, 52]) for the converse implication (ii) $\implies$ (i). As most applications only need (i) $\implies$ (ii), we only prove that implication. For a proof of (ii) $\implies$ (i) see, e.g., [154, Theorem III.4.3].

Theorem 15.43 (Central limit theorem of Lindeberg-Feller) Let $(X_{n,l})$ be an independent centered and normed array of real random variables. For every $n \in \mathbb{N}$, let $S_n = X_{n,1} + \dots + X_{n,k_n}$. Then the following are equivalent.

- The Lindeberg condition holds.
- $(X_{n,l})$ is a null array and $\mathbf{P}_{S_n} \xrightarrow{n \rightarrow \infty} \mathcal{N}_{0,1}$.

We prepare for the proof of Lindeberg's theorem with a couple of lemmas.

Lemma 15.44 If (i) of Theorem 15.43 holds, then $(X_{n,l})$ is a null array.

Proof For $\varepsilon > 0$, by Chebyshev's inequality,

$$\sum_{l=1}^{k_n} \mathbf{P}[|X_{n,l}| > \varepsilon] \leq \varepsilon^{-2} \sum_{l=1}^{k_n} \mathbf{E} \left[X_{n,l}^2 \mathbb{1}_{\{|X_{n,l}| > \varepsilon\}} \right] = L_n(\varepsilon) \xrightarrow{n \rightarrow \infty} 0.$$

In the following, $\varphi_{n,l}$ and φ_n will always denote the characteristic functions of $X_{n,l}$ and S_n .
 Lemma 15.45 For every $n \in \mathbb{N}$ and $t \in \mathbb{R}$, we have

$$\sum_{l=1}^{k_n} |1 - \varphi_{n,l}(t)| \leq \frac{t^2}{2}$$

Proof For every $x \in \mathbb{R}$, we have $|e^{itx} - 1 - itx| \leq \frac{t^2 x^2}{2}$. Since $\mathbf{E}[X_{n,l}] = 0$,

$$\begin{aligned} \sum_{l=1}^{k_n} |\varphi_{n,l}(t) - 1| &= \sum_{l=1}^{k_n} |\mathbf{E}[e^{itX_{n,l}} - 1]| \\ &\leq \sum_{l=1}^{k_n} \mathbf{E}[|e^{itX_{n,l}} - itX_{n,l} - 1|] + |\mathbf{E}[itX_{n,l}]| \\ &\leq \sum_{l=1}^{k_n} \frac{t^2}{2} \mathbf{E}[X_{n,l}^2] = \frac{t^2}{2}. \end{aligned}$$

Lemma 15.46 If (i) of Theorem 15.43 holds, then

$$\lim_{n \rightarrow \infty} \left| \log \varphi_n(t) - \sum_{l=1}^{k_n} \mathbf{E}[e^{itX_{n,l}} - 1] \right| = 0$$

Proof Let $m_n := \max_{l=1, \dots, k_n} |\varphi_{n,l}(t) - 1|$. Note that, for all $\varepsilon > 0$,

$$|e^{itx} - 1| \leq \begin{cases} 2x^2/\varepsilon^2, & \text{if } |x| > \varepsilon, \\ \varepsilon t, & \text{if } |x| \leq \varepsilon. \end{cases}$$

This implies

$$\begin{aligned} |\varphi_{n,l}(t) - 1| &\leq \mathbf{E}[|e^{itX_{n,l}} - 1| \mathbb{1}_{\{|X_{n,l}| \leq \varepsilon\}}] + \mathbf{E}[|e^{itX_{n,l}} - 1| \mathbb{1}_{\{|X_{n,l}| > \varepsilon\}}] \\ &\leq \varepsilon t + 2\varepsilon^{-2} \mathbf{E}[X_{n,l}^2 \mathbb{1}_{\{|X_{n,l}| > \varepsilon\}}]. \end{aligned}$$

Hence, for all $\varepsilon > 0$,

$$\limsup_{n \rightarrow \infty} m_n \leq \limsup_{n \rightarrow \infty} (\varepsilon t + 2\varepsilon^{-2} L_n(\varepsilon)) = \varepsilon t,$$

and thus $\lim_{n \rightarrow \infty} m_n = 0$. Now $|\log(1+x) - x| \leq x^2$ for all $x \in \mathbb{C}$ with $|x| \leq \frac{1}{2}$. If n is sufficiently large that $m_n < \frac{1}{2}$, then

$$\begin{aligned}
\left| \log \varphi_n(t) - \sum_{l=1}^{k_n} \mathbf{E} [e^{itX_{n,l}} - 1] \right| &= \left| \sum_{l=1}^{k_n} \log(\varphi_{n,l}(t)) - \mathbf{E} [e^{itX_{n,l}} - 1] \right| \\
&\leq \sum_{l=1}^{k_n} (\varphi_{n,l}(t) - 1)^2 \\
&\leq m_n \sum_{l=1}^{k_n} |\varphi_{n,l}(t) - 1| \\
&\leq \frac{1}{2} m_n t^2 \quad (\text{by Lemma 15.45}) \\
&\longrightarrow 0 \quad \text{for } n \rightarrow \infty
\end{aligned}$$

The fundamental trick of the proof, which is worth remembering, consists in the introduction of the function

$$f_t(x) := \begin{cases} \frac{1+x^2}{x^2} (e^{itx} - 1 - \frac{itx}{1+x^2}), & \text{if } x \neq 0 \\ -\frac{t^2}{2}, & \text{if } x = 0 \end{cases} \quad (15.8)$$

and the measures $\mu_n, v_n \in \mathcal{M}_f(\mathbb{R}), n \in \mathbb{N}$,

$$v_n(dx) := \sum_{l=1}^{k_n} x^2 \mathbf{P}_{X_{n,l}}(dx) \quad \text{and} \quad \mu_n(dx) := \sum_{l=1}^{k_n} \frac{x^2}{1+x^2} \mathbf{P}_{X_{n,l}}(dx)$$

Lemma 15.47 For every $t \in \mathbb{R}$, we have $f_t \in C_b(\mathbb{R})$.

Proof For all $|x| \geq 1$, we have $\frac{1+x^2}{x^2} \leq 2$; hence

$$|f_t(x)| \leq 2 \left(|e^{itx}| + 1 + \frac{tx}{1+x^2} \right) \leq 4 + 2|t|$$

We have to show that f_t is continuous at 0. By Taylor's formula (Lemma 15.30), we get

$$e^{itx} = 1 + itx - \frac{t^2 x^2}{2} + R(tx)$$

where the error term is bounded by $|R(tx)| \leq \frac{1}{6}|tx|^3$. Hence, for fixed t ,

$$\lim_{0 \neq x \rightarrow 0} f_t(x) = \lim_{0 \neq x \rightarrow 0} \frac{1}{x^2} \left(itx \left(1 - \frac{1}{1+x^2} \right) - \frac{t^2 x^2}{2} + R(tx) \right) = -\frac{t^2}{2}.$$

Lemma 15.48 If (i) of Theorem 15.43 holds, then $v_n \xrightarrow{n \rightarrow \infty} \delta_0$ weakly.

Proof For every $n \in \mathbb{N}$, we have $v_n \in \mathcal{M}_1(\mathbb{R})$ since

$$v_n(\mathbb{R}) = \sum_{l=1}^{k_n} \int x^2 \mathbf{P}_{X_{n,l}}(dx) = \sum_{l=1}^{k_n} \text{Var}[X_{n,l}] = 1$$

However, for $\varepsilon > 0$, we have $v_n((-\varepsilon, \varepsilon)^c) = L_n(\varepsilon) \xrightarrow{n \rightarrow \infty} 0$; hence $v_n \xrightarrow{n \rightarrow \infty} \delta_0$.

Lemma 15.49 If (i) of Theorem 15.43 holds, then

$$\int f_t(x) \mu_n(dx) + it \int \frac{1}{x} \mu_n(dx) \xrightarrow{n \rightarrow \infty} -\frac{t^2}{2}.$$

Proof Since $(x \mapsto f_t(x)/(1+x^2)) \in C_b(\mathbb{R})$, by Lemma 15.48,

$$\int f_t(x) \mu_n(dx) = \int f_t(x) \frac{1}{1+x^2} v_n(dx) \xrightarrow{n \rightarrow \infty} f_t(0) = -\frac{t^2}{2}.$$

Now $(x \mapsto x/(1+x^2)) \in C_b(\mathbb{R})$ and $\mathbf{E}[X_{n,l}] = 0$ for all n and l ; hence

$$\begin{aligned} \int \frac{1}{x} \mu_n(dx) &= \sum_{l=1}^{k_n} \mathbf{E} \left[\frac{X_{n,l}}{1+X_{n,l}^2} \right] = \sum_{l=1}^{k_n} \mathbf{E} \left[\frac{X_{n,l}}{1+X_{n,l}^2} - X_{n,l} \right] \\ &= - \sum_{l=1}^{k_n} \mathbf{E} \left[X_{n,l}^2 \cdot \frac{X_{n,l}}{1+X_{n,l}^2} \right] \\ &= - \int \frac{x}{1+x^2} v_n(dx) \xrightarrow{n \rightarrow \infty} 0 \end{aligned}$$

Proof of Theorem 15.43 "(i) $\implies$ (ii)" We have to show that $\lim_{n \rightarrow \infty} \log \varphi_n(t) = -\frac{t^2}{2}$ for every $t \in \mathbb{R}$. By Lemma 15.46, this is equivalent to

$$\lim_{n \rightarrow \infty} \sum_{l=1}^{k_n} (\varphi_{n,l}(t) - 1) = -\frac{t^2}{2}.$$

Now $f_t(x) \frac{x^2}{1+x^2} = e^{itx} - 1 - \frac{itx}{1+x^2}$. Hence

$$\begin{aligned} \sum_{l=1}^{k_n} (\varphi_{n,l}(t) - 1) &= \sum_{l=1}^{k_n} \int \left(f_t(x) \frac{x^2}{1+x^2} + \frac{itx}{1+x^2} \right) \mathbf{P}_{X_{n,l}}(dx) \\ &= \int f_t d\mu_n + it \int \frac{1}{x} \mu_n(dx) \\ &\xrightarrow{n \rightarrow \infty} -\frac{t^2}{2} \quad (\text{by Lemma 15.49}). \end{aligned}$$

As an application of the Lindeberg-Feller theorem, we give the so-called threeseries theorem, which is due to Kolmogorov.

Theorem 15.50 (Kolmogorov's three-series theorem) Let $X_1, X_2, \dots$ be independent real random variables. Let $K > 0$ and $Y_n := X_n \mathbb{1}_{\{|X_n| \leq K\}}$ for all $n \in \mathbb{N}$.

The series $\sum_{n=1}^{\infty} X_n$ converges almost surely if and only if each of the following three conditions holds:

- (i) $\sum_{n=1}^{\infty} \mathbf{P}[|X_n| > K] < \infty$.
- (ii) $\sum_{n=1}^{\infty} \mathbf{E}[Y_n]$ converges.
- (iii) $\sum_{n=1}^{\infty} \text{Var}[Y_n] < \infty$.

Proof ". Assume that (i), (ii) and (iii) hold. By Exercise 7.1.1, since (iii) holds, the series $\sum_{n=1}^{\infty} (Y_n - \mathbf{E}[Y_n])$ converges a.s. As (ii) holds, $\sum_{n=1}^{\infty} Y_n$ converges almost surely. By the Borel-Cantelli lemma, there exists an $N = N(\omega)$ such that $|X_n| < K$; hence $X_n = Y_n$ for all $n \geq N$. Hence $\sum_{n=1}^{\infty} X_n = \sum_{n=1}^{N-1} X_n + \sum_{n=N}^{\infty} Y_n$ converges a.s.

" $\implies$ " Assume that $\sum_{n=1}^{\infty} X_n$ converges a.s. Clearly, this implies (i) (otherwise, by the Borel-Cantelli lemma, $|X_n| > K$ infinitely often, contradicting the assumption).

We assume that (iii) does not hold and produce a contradiction. To this end, let $\sigma_n^2 = \sum_{k=1}^n \text{Var}[Y_k]$ and define an array $(X_{n,l}; l = 1, \dots, n, n \in \mathbb{N})$ by $X_{n,l} = (Y_l - \mathbf{E}[Y_l])/\sigma_n$. This array is centered and normed. Since $\sigma_n^2 \xrightarrow{n \rightarrow \infty} \infty$, for every $\varepsilon > 0$ and for sufficiently large $n \in \mathbb{N}$, we have $2K < \varepsilon \sigma_n$; thus $|X_{n,l}| \leq \varepsilon$ for all $l = 1, \dots, n$. This implies $L_n(\varepsilon) \xrightarrow{n \rightarrow \infty} 0$, where $L_n(\varepsilon) = \sum_{l=1}^n \mathbf{E} \left[X_{n,l}^2 \mathbb{1}_{\{|X_{n,l}| \geq \varepsilon\}} \right]$ is the quantity of the Lindeberg condition (see (15.6)). By the Lindeberg-Feller theorem, we then get $S_n := X_{n,1} + \dots + X_{n,n} \xrightarrow{n \rightarrow \infty} \mathcal{N}_{0,1}$. As shown in the first part of this proof, almost sure convergence of $\sum_{n=1}^{\infty} X_n$ and (i) imply that

$$\sum_{n=1}^{\infty} Y_n \quad \text{converges almost surely.} \quad (15.9)$$

In particular, $T_n := (Y_1 + \dots + Y_n) / \sigma_n \xrightarrow{n \rightarrow \infty} 0$. Thus, by Slutsky's theorem, we also have $(S_n - T_n) \xrightarrow{n \rightarrow \infty} \mathcal{N}_{0,1}$. On the other hand, for all $n \in \mathbb{N}$, the difference $S_n - T_n$ is deterministic, contradicting the assumption that (iii) does not hold.

Now that we have established (iii), by Exercise 7.1.1, we see that $\sum_{n=1}^{\infty} (Y_n - \mathbf{E}[Y_n])$ converges almost surely. Together with (15.9), we conclude (ii).

As a supplement, we cite a statement about the speed of convergence in the central limit theorem (see, e.g., [154, Chapter III, Section 11] for a proof). With different bounds (instead of 0.8), the statement was found independently by Berry [10] and Esseen [46].

Theorem 15.51 (Berry-Esseen) Let $X_1, X_2, \dots$ be independent and identically distributed with $\mathbf{E}[X_1] = 0$, $\mathbf{E}[X_1^2] = \sigma^2 \in (0, \infty)$ and $\gamma := \mathbf{E}[|X_1|^3] < \infty$. Let $S_n^* := \frac{1}{\sqrt{n\sigma^2}}(X_1 + \dots + X_n)$ and let $\Phi : x \mapsto \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-t^2/2} dt$ be the distribution function of the standard normal distribution. Then, for all $n \in \mathbb{N}$,

$$\sup_{x \in \mathbb{R}} |\mathbf{P}[S_n^* \leq x] - \Phi(x)| \leq \frac{0.8\gamma}{\sigma^3 \sqrt{n}}$$

Example 15.52 Let $\alpha \in (0, 1)$. Consider the distribution μ_α on $\mathbb{R}$ with density

$$f_\alpha(x) = \frac{1}{2\alpha} |x|^{-1-1/\alpha} \mathbb{1}_{\{|x| \geq 1\}}$$

Let $X_1, X_2, \dots$, be i.i.d. random variables with distribution μ_α . Then $\mathbf{E}[X_1] = 0$ and $\sigma^2 := \text{Var}[X_1] = 1/(1 - 2\alpha) < \infty$ if $\alpha < 1/2$. Let F_n denote the distribution function of S_n^* and F_Φ the distribution function of the standard normal distribution.

The closer F_n and F_Φ are, the closer lie the points $(F_\Phi^{-1}(t), F_n^{-1}(t))$ on the diagonal $\{(x, x) : x \in \mathbb{R}\}$. A graphical representation of the points $(F_\Phi^{-1}(t), F_n^{-1}(t))$, $t \in \mathbb{R}$ is called Q-Q-plot or quantile-quantile-plot.

As α approaches $1/2$, the distribution μ_α has less and less moments. Hence we expect the convergence in the central limit theorem to be slower. For fixed n , we expect the deviation of F_n from F_Φ to be larger for larger α . The graphs in Fig. 15.2 illustrate this.


Exercise 15.5.1 The argument of Remark 15.38 is more direct than the argument with Lévy's continuity theorem but is less robust: Give a sequence $X_1, X_2, \dots$ of independent real random variables with $\mathbf{E}[|X_n|] = \infty$ for all $n \in \mathbb{N}$ but such that

$$\frac{X_1 + \dots + X_n}{\sqrt{n}} \xrightarrow{n \rightarrow \infty} \mathcal{N}_{0,1}.$$

Exercise 15.5.2 Let $Y_1, Y_2, \dots$ be i.i.d. with $\mathbf{E}[Y_i] = 0$ and $\mathbf{E}[Y_i^2] = 1$. Let $Z_1, Z_2, \dots$ be independent random variables (and independent of $Y_1, Y_2, \dots$) with

$$\mathbf{P}[Z_i = i] = \mathbf{P}[Z_i = -i] = \frac{1}{2} (1 - \mathbf{P}[Z_i = 0]) = \frac{1}{2} \frac{1}{i^2}$$

For $i, n \in \mathbb{N}$, define $X_i := Y_i + Z_i$ and $S_n = X_1 + \dots + X_n$.


 Fig. 15.2 Q-Q-plots for S_{100}

from Example 15.52 with $\alpha = 0.4$ (left) and $\alpha = 0.48$ (right). The abscissa shows the quantiles of the standard normal distribution. For convenience, also the diagonal is drawn

Show that $n^{-1/2}S_n \xrightarrow{n \rightarrow \infty} \mathcal{N}_{0,1}$ but that $(X_i)_{i \in \mathbb{N}}$ does not satisfy the Lindeberg condition.

Hint: Do not try a direct computation!

Exercise 15.5.3 Let $X_1, X_2, \dots$ be i.i.d. random variables with density

$$f(x) = \frac{1}{|x|^3} \mathbb{1}_{\mathbb{R} \setminus [-1,1]}(x)$$

Then $\mathbf{E}[X_1^2] = \infty$ but there are numbers $A_1, A_2, \dots$, such that

$$\frac{X_1 + \dots + X_n}{A_n} \xrightarrow{n \rightarrow \infty} \mathcal{N}_{0,1}.$$

Determine one such sequence $(A_n)_{n \in \mathbb{N}}$ explicitly.

5.15.6 15.6 Multidimensional Central Limit Theorem

We come to a multidimensional version of the CLT.

Definition 15.53 Let C be a (strictly) positive definite symmetric real $d \times d$ matrix and let $\mu \in \mathbb{R}^d$. A random vector $X = (X_1, \dots, X_d)^T$ is called d -dimensional normally distributed with expectation μ and covariance matrix C if X has the density

$$f_{\mu,C}(x) = \frac{1}{\sqrt{(2\pi)^d \det(C)}} \exp\left(-\frac{1}{2} \langle x - \mu, C^{-1}(x - \mu) \rangle\right) \quad (15.10)$$

for $x \in \mathbb{R}^d$. In this case, we write $X \sim \mathcal{N}_{\mu,C}$.

Theorem 15.54 Let $\mu \in \mathbb{R}^d$ and let C be a real positive definite symmetric $d \times d$ matrix. If $X \sim \mathcal{N}_{\mu,C}$, then the following statements hold.

- (i) $\mathbf{E}[X_i] = \mu_i$ for all $i = 1, \dots, d$.
- (ii) $\text{Cov}[X_i, X_j] = C_{i,j}$ for all $i, j = 1, \dots, d$.
- (iii) $\langle \lambda, X \rangle \sim \mathcal{N}_{\langle \lambda, \mu \rangle, \langle \lambda, C \lambda \rangle}$ for every $\lambda \in \mathbb{R}^d$.
- (iv) $\varphi(t) := \mathbf{E}[e^{i\langle t, X \rangle}] = e^{i\langle t, \mu \rangle} e^{-\frac{1}{2} \langle t, C t \rangle}$ for every $t \in \mathbb{R}^d$.

Moreover, $X \sim \mathcal{N}_{\mu,C} \iff \text{(iii)} \iff \text{(iv)}$.

Proof (i) and (ii) follow by simple computations. The same is true for (iii) and (iv). The implication (iii) $\implies$ (iv) is straightforward. The family

$$\{f_t : x \mapsto e^{i\langle t, x \rangle}, t \in \mathbb{R}^d\}$$

is a separating class for $\mathcal{M}_1(\mathbb{R}^d)$ by the Stone-Weierstraß theorem. Hence φ determines the distribution of X uniquely.

Remark 15.55 For one-dimensional normal distributions, it is natural to define the degenerate normal distribution by $\mathcal{N}_{\mu,0} := \delta_\mu$. For the multidimensional situation, there are various possibilities for degeneracy depending on the size of the kernel of C . If C is only positive semidefinite (and symmetric, of course), we define $\mathcal{N}_{\mu,C}$ as that distribution on $\mathbb{R}^n$ with characteristic function $\varphi(t) = e^{i\langle t, \mu \rangle} e^{-\frac{1}{2}\langle t, Ct \rangle}$.

Theorem 15.56 (Cramér-Wold device) Let $X_n = (X_{n,1}, \dots, X_{n,d})^T \in \mathbb{R}^d, n \in \mathbb{N}$, be random vectors. Then, the following are equivalent:

- (i) There is a random vector X such that $X_n \xrightarrow{n \rightarrow \infty} X$.
- (ii) For any $\lambda \in \mathbb{R}^d$, there is a random variable X^λ such that $\langle \lambda, X_n \rangle \xrightarrow{n \rightarrow \infty} X^\lambda$.

If (i) and (ii) hold, then $X^\lambda \stackrel{\mathcal{D}}{=} \langle \lambda, X \rangle$ for all $\lambda \in \mathbb{R}^d$.

Proof Assume (i). Let $\lambda \in \mathbb{R}^d$ and $s \in \mathbb{R}$. The map $\mathbb{R}^d \rightarrow \mathbb{C}, x \mapsto e^{is\langle \lambda, x \rangle}$ is continuous and bounded; hence we have $\mathbf{E}[e^{is\langle \lambda, X_n \rangle}] \xrightarrow{n \rightarrow \infty} \mathbf{E}[e^{is\langle \lambda, X \rangle}]$. Thus (ii) holds with $X^\lambda := \langle \lambda, X \rangle$.

Now assume (ii). Then $(\mathbf{P}_{X_{n,l}})_{n \in \mathbb{N}}$ is tight for every $l = 1, \dots, d$. Hence $(\mathbf{P}_{X_n})_{n \in \mathbb{N}}$ is tight and thus relatively sequentially compact (Prohorov's theorem). For any weak limit point Q for $(\mathbf{P}_{X_n})_{n \in \mathbb{N}}$ and for any $\lambda \in \mathbb{R}^d$, we have

$$\int Q(dx) e^{i\langle \lambda, x \rangle} = \mathbf{E}[e^{iX^\lambda}]$$

Hence the limit point Q is unique and thus $(\mathbf{P}_{X_n})_{n \in \mathbb{N}}$ converges weakly to Q . That is, (i) holds.

If (ii) holds, then the distributions of the limiting random variables X^λ are uniquely determined and by what we have shown already, $X^\lambda = \langle \lambda, X \rangle$ is one possible choice. Thus $X^\lambda \stackrel{\mathcal{D}}{=} \langle \lambda, X \rangle$.

Theorem 15.57 (Central limit theorem in $\mathbb{R}^d$) Let $(X_n)_{n \in \mathbb{N}}$ be i.i.d. random vectors with $\mathbf{E}[X_{n,i}] = 0$ and $\mathbf{E}[X_{n,i}X_{n,j}] = C_{ij}, i, j = 1, \dots, d$. Let $S_n^* := \frac{X_1 + \dots + X_n}{\sqrt{n}}$. Then

$$\mathbf{P}_{S_n^*} \xrightarrow{n \rightarrow \infty} \mathcal{N}_{0,C} \quad \text{weakly.}$$

Proof Let $\lambda \in \mathbb{R}^d$. Define $X_n^\lambda = \langle \lambda, X_n \rangle, S_n^\lambda = \langle \lambda, S_n^* \rangle$ and $S_\infty \sim \mathcal{N}_{0,C}$. Then $\mathbf{E}[X_n^\lambda] = 0$ and $\text{Var}[X_n^\lambda] = \langle \lambda, C\lambda \rangle$. By the one-dimensional central limit theorem, we have $\mathbf{P}_{S_n^\lambda} \xrightarrow{n \rightarrow \infty} \mathcal{N}_{0, \langle \lambda, C\lambda \rangle} = \mathbf{P}_{\langle \lambda, S_\infty \rangle}$. By Theorem 15.56, this yields the claim.

Exercise 15.6.1 Let $\mu \in \mathbb{R}^d$, let C be a symmetric positive semidefinite real $d \times d$ matrix and let $X \sim \mathcal{N}_{\mu,C}$ (in the sense of Remark 15.55). Show that $AX \sim \mathcal{N}_{A\mu, ACAT}$ for every $m \in \mathbb{N}$ and every real $m \times d$ matrix A .

Exercise 15.6.2 (Cholesky factorization) Let C be a positive definite symmetric real $d \times d$ matrix. Then there exists a real $d \times d$ matrix $A = (a_{kl})$ with $A \cdot A^T = C$. The matrix A can be chosen to be lower triangular. Let $W := (W_1, \dots, W_d)^T$, where $W_1, \dots, W_d$ are independent and $\mathcal{N}_{0,1}$ -distributed. Define $X := AW + \mu$. Show that $X \sim \mathcal{N}_{\mu,C}$.

5.16 16 Infinitely Divisible Distributions

For every n , the normal distribution $\mathcal{N}_{\mu, \sigma^2}$ is the n th convolution power of a probability measure (namely, of $\mathcal{N}_{\mu/n, \sigma^2/n}$). This property is called infinite divisibility and is shared by other probability distributions such as the Poisson distribution and the Gamma distribution. In the first section, we study which probability measures on $\mathbb{R}$ are infinitely divisible and give an exhaustive description of this class of distributions by means of the Lévy-Khinchin formula.

Unlike the Poisson distribution, the normal distribution is the limit of rescaled sums of i.i.d. random variables (central limit theorem). In the second section, we investigate briefly which subclass of the infinitely divisible measures on $\mathbb{R}$ shares this property.

5.16.1 16.1 Lévy-khinchin Formula

For the sake of brevity, in this section, we use the shorthand "CFP" for "characteristic function of a probability measure on $\mathbb{R}$ ".

Definition 16.1 A measure $\mu \in \mathcal{M}_1(\mathbb{R})$ is called infinitely divisible if, for every $n \in \mathbb{N}$, there is a $\mu_n \in \mathcal{M}_1(\mathbb{R})$ such that $\mu_n^{*n} = \mu$. Analogously, a CFP φ is called infinitely divisible if, for every $n \in \mathbb{N}$, there is a CFP φ_n such that $\varphi = \varphi_n^n$. A real random variable X is called infinitely divisible if, for every $n \in \mathbb{N}$, there exist i.i.d. random variables $X_{n,1}, \dots, X_{n,n}$ such that $X \stackrel{D}{=} X_{n,1} + \dots + X_{n,n}$.

Manifestly, all three notions of infinite divisibility are equivalent, and we will use them synonymously. Note that the uniqueness of μ_n and φ_n , respectively, is by no means evident. Indeed, n -fold divisibility alone does not imply uniqueness of the n th convolution root $\mu^{*1/n} := \mu_n$ or of φ_n , respectively. As an example for even n , choose a real-valued CFP φ for which $|\varphi| \neq \varphi$ is also a CFP (see Examples 15.16 and 15.17). Then $\varphi^n = |\varphi|^n$ is n -fold divisible; however, the factors are not unique.

By virtue of Lévy's continuity theorem, one can show that (see Exercise 16.1.1) $\varphi(t) \neq 0$ for all $t \in \mathbb{R}$ if φ is infinitely divisible. The probabilistic meaning of this fact is that as a continuous function $\log(\varphi(t))$ is uniquely defined and thus there exists only one continuous function $\varphi^{1/n} = \exp(\log(\varphi)/n)$. The n th convolution roots are thus unique if the distribution is infinitely divisible.

Example 16.2

- (i) δ_x is infinitely divisible with $\delta_{x/n}^{*n} = \delta_x$ for every $n \in \mathbb{N}$.
- (ii) The normal distribution is infinitely divisible with $\mathcal{N}_{m,\sigma^2} = \mathcal{N}_{m/n,\sigma^2/n}^{*n}$.
- (iii) The Cauchy distribution Cau_a with density $x \mapsto (a\pi)^{-1} (1 + (x/a)^2)^{-1}$ is infinitely divisible with $\text{Cau}_a = \text{Cau}_{a/n}^{*n}$. Indeed, Cau_a has CFP $\varphi_a(t) = e^{-a|t|}$; hence $\varphi_{a/n}^n = \varphi_a$.
- (iv) Every symmetric stable distribution with index $\alpha \in (0, 2]$ and scale parameter $\gamma > 0$ (that is, the distribution with CFP $\varphi_{\alpha,\gamma}(t) = e^{-|\gamma t|^\alpha}$) is infinitely divisible. Indeed, $\varphi_{\alpha,\gamma/n^{1/\alpha}}^n = \varphi_{\alpha,\gamma}$. (To be precise, we have shown only for $\alpha \in (0, 1]$ (in Corollary 15.25) and for $\alpha = 2$ (normal distribution) that $\varphi_{\alpha,\gamma}$ is in fact a CFP. In Section 16.2, we will show that this is true for all $\alpha \in (0, 2]$. For $\alpha > 2$, $\varphi_{\alpha,\gamma}$ is not a CFP, see Exercise 15.4.3.)
- (v) The Gamma distribution $\Gamma_{\theta,r}$ with CFP $\varphi_{\theta,r}(t) = \exp(r\psi_\theta(t))$, where $\psi_\theta(t) = \log(1 - it/\theta)$, is infinitely divisible with $\Gamma_{\theta,r} = \Gamma_{\theta/r,n}^{*n}$.
- (vi) The Poisson distribution is infinitely divisible with $\text{Poi}_\lambda = \text{Poi}_{\lambda/n}^{*n}$.
- (vii) The negative binomial distribution $b_{r,p}^-(\{k\}) = \binom{-r}{k} (-1)^k p^r (1-p)^k$, $k \in \mathbb{N}_0$, with parameters $r > 0$ and $p \in (0, 1)$, is infinitely divisible with $b_{r,p}^- = \left(b_{r/n,p}^-\right)^{*n}$. Indeed, $\varphi_{r,p}(t) = e^{r\psi_p(t)}$, where

$$\psi_p(t) = \log(p) - \log(1 - (1-p)e^{it})$$

(viii) Let X and Y be independent with $X \sim \mathcal{N}_{0,\sigma^2}$ and $Y \sim \Gamma_{\theta,r}$, where $\sigma^2, \theta, r > 0$. It can be shown that the random variable $Z := X/\sqrt{Y}$ is infinitely divisible (see [65] or [131]). In particular, Student's t -distribution with $k \in \mathbb{N}$ degrees of freedom is infinitely divisible (this is the case where $\sigma^2 = 1$ and $\theta = r = k/2$).

(ix) The binomial distribution $b_{n,p}$ with parameters $n \in \mathbb{N}$ and $p \in (0, 1)$ is not infinitely divisible (why?).

(x) Somewhat more generally, there is no nontrivial infinitely divisible distribution that is concentrated on a bounded interval.

A main goal of this section is to show that every infinitely divisible distribution can be composed of three generic ones:

- the Dirac measures δ_x with $x \in \mathbb{R}$,
- the normal distributions $\mathcal{N}_{\mu, \sigma^2}$ with $\mu \in \mathbb{R}$ and $\sigma^2 > 0$, and
- (limits of) convolutions of Poisson distributions.

As convolutions of Poisson distributions play a special role, we will consider them separately.

If $v \in \mathcal{M}_1(\mathbb{R})$ with CFP φ_v and if $\lambda > 0$, then one can easily check that $\varphi(t) = \exp(\lambda(\varphi_v(t) - 1))$ is the CFP of $\mu_\lambda = \sum_{k=0}^{\infty} e^{-\lambda} \frac{\lambda^k}{k!} \nu^{*k}$. Hence, formally we can write $\mu_\lambda = e^{*\lambda(v-\delta_0)}$. Indeed, μ_λ is infinitely divisible with $\mu_\lambda = \mu_{\lambda/n}^{*n}$. We want to combine the two parameters λ and v into one parameter λv . For $v \in \mathcal{M}_f(\mathbb{R})$, we can define $\nu^{*n} = \nu(\mathbb{R})^n (\nu/\nu(\mathbb{R}))^{*n}$ (and $\nu^{*n} = 0$ if $\nu = 0$). In both cases, let $\nu^{*0} := \delta_0$. Hence we make the following definition.

Definition 16.3 The compound Poisson distribution with intensity measure $\nu \in \mathcal{M}_f(\mathbb{R})$ is the following probability measure on $\mathbb{R}$:

$$\text{CPoi}_\nu := e^{*(v-\nu(\mathbb{R})\delta_0)} := e^{-v(\mathbb{R})} \sum_{n=0}^{\infty} \frac{v^{*n}}{n!}.$$

The CFP of CPoi_ν is given by

$$\varphi_\nu(t) = \exp \left(\int (e^{itx} - 1) v(dx) \right) \quad (16.1)$$

In particular, $\text{CPoi}_{\mu+\nu} = \text{CPoi}_\mu * \text{CPoi}_\nu$; hence CPoi_ν is infinitely divisible.

Example 16.4 For every measurable set $A \subset \mathbb{R} \setminus \{0\}$ and every $r > 0$,

$$r^{-1} \text{CPoi}_{rv}(A) = e^{-rv(\mathbb{R})} v(A) + e^{-rv(\mathbb{R})} \sum_{k=2}^{\infty} \frac{r^{k-1} v^{*k}(A)}{k!} \xrightarrow{r \downarrow 0} v(A).$$

We use this in order to show that $b_{r,p}^- = \text{CPoi}_{rv}$ for some $\nu \in \mathcal{M}_f(\mathbb{N})$. To this end, for $k \in \mathbb{N}$, we compute

$$r^{-1} b_{r,p}^-(\{k\}) = \frac{r(r+1) \cdot 3 \cdot (r+k-1)}{rk!} p^r (1-p)^k \xrightarrow{r \downarrow 0} \frac{(1-p)^k}{k}.$$

If we had $b_{r,p}^- = \text{CPoi}_{rv}$ for some $v \in \mathcal{M}_f(\mathbb{N})$, then we would have $v(\{k\}) = (1-p)^k/k$. We compute the CFP of CPoi_{rv} for this v ,

$$\varphi_{rv}(t) = \exp \left(r \sum_{k=1}^{\infty} \frac{((1-p)e^{it})^k - (1-p)^k}{k} \right) = p^r (1 - (1-p)e^{it})^{-r}.$$

However, this is the CFP of $b_{r,p}^-$; hence indeed $b_{r,p}^- = \text{CPoi}_{rv}$.

Not every infinitely divisible distribution is of the type CPoi_ν , however we have the following theorem.

Theorem 16.5 A probability measure μ on $\mathbb{R}$ is infinitely divisible if and only if there is a sequence $(v_n)_{n \in \mathbb{N}}$ in $\mathcal{M}_f(\mathbb{R} \setminus \{0\})$ such that $\text{CPoi}_{v_n} \xrightarrow{n \rightarrow \infty} \mu$.

Since every CPoi_{v_n} is infinitely divisible, on the one hand we have to show that this property is preserved under weak limits. On the other hand, we show that, for infinitely divisible μ , the sequence $v_n = n\mu^{*1/n}$ does the trick. We prepare for the proof of Theorem 16.5 with a further theorem.

Theorem 16.6 Let $(\varphi_n)_{n \in \mathbb{N}}$ be a sequence of CFPs. Then the following are equivalent.

- (i) For every $t \in \mathbb{R}$, the limit $\varphi(t) = \lim_{n \rightarrow \infty} \varphi_n^n(t)$ exists and φ is continuous at 0.
 - (ii) For every $t \in \mathbb{R}$, the limit $\psi(t) = \lim_{n \rightarrow \infty} n(\varphi_n(t) - 1)$ exists and ψ is continuous at 0.
- If (i) and (ii) hold, then $\varphi = e^\psi$ is a CFP.

Proof The proof is based on a Taylor expansion of the logarithm,

$$|\log(z) - (z - 1)| \leq \frac{1}{2}|z - 1|^2 \quad \text{for } z \in \mathbb{C} \text{ with } |z - 1| < \frac{1}{2}$$

In particular, for $(z_n)_{n \in \mathbb{N}}$ in $\mathbb{C}$,

$$\limsup_{n \rightarrow \infty} n|z_n - 1| < \infty \iff \limsup_{n \rightarrow \infty} |n \log(z_n)| < \infty, \quad (16.2)$$

and $\lim_{n \rightarrow \infty} n(z_n - 1) = \lim_{n \rightarrow \infty} n \log(z_n)$ if one of the limits exists.

Applying this to $z_n = \varphi_n(t)$, we see that (ii) implies (i). On the other hand, (i) implies (ii) if $\liminf_{n \rightarrow \infty} n \log(|\varphi_n(t)|) > -\infty$ and hence if $\varphi(t) \neq 0$ for all $t \in \mathbb{R}$.

Since φ is continuous at 0 and since $\varphi(0) = 1$, there is an $\varepsilon > 0$ with $|\varphi(t)| > \frac{1}{2}$ for all $t \in [-\varepsilon, \varepsilon]$. Since φ and φ_n are CFPs, $|\varphi|^2$ and $|\varphi_n|^2$ are also CFPs. Thus, since $|\varphi_n(t)|^{2n}$ converges to $|\varphi(t)|^2$ pointwise, Lévy's continuity theorem implies uniform convergence on compact sets. Now apply (16.2) with $z_n = |\varphi_n(t)|^2$. Thus $(n(1 - |\varphi_n(t)|^2))_{n \in \mathbb{N}}$ is bounded for $t \in [-\varepsilon, \varepsilon]$. Hence, by Lemma 15.11(v), $n(1 - |\varphi_n(2t)|^2) \leq 4n(1 - |\varphi_n(t)|^2)$ also is bounded; thus

$$|\varphi(2t)|^2 \geq \liminf_{n \rightarrow \infty} \exp(4n(|\varphi_n(t)|^2 - 1)) = (|\varphi(t)|^2)^4.$$

Inductively, we get $|\varphi(t)| \geq 2^{-(4^k)}$ for $|t| \leq 2^k \varepsilon$. Hence there is a $\gamma > 0$ such that

$$|\varphi(t)| > \frac{1}{2} e^{-\gamma t^2} \quad \text{for all } t \in \mathbb{R} \quad (16.3)$$

If (i) and (ii) hold, then

$$\log \varphi(t) = \lim_{n \rightarrow \infty} n \log(\varphi_n(t)) = \lim_{n \rightarrow \infty} n(\varphi_n(t) - 1) = \psi(t)$$

By Lévy's continuity theorem, as a continuous limit of CFPs, φ is a CFP.

Corollary 16.7 If the conditions of Theorem 16.6 hold, then φ^r is a CFP for every $r > 0$. In particular, $\varphi = (\varphi^{1/n})^n$ is infinitely divisible.

Proof If φ_n is the CFP of $\mu_n \in \mathcal{M}_1(\mathbb{R})$, then $e^{rn(\varphi_n - 1)}$ is the CFP of $\text{CPoi}_{rn\mu_n}$. Being a limit of CFPs that is continuous at 0, by Lévy's continuity theorem, $\varphi^r = e^{r\psi} = \lim_{n \rightarrow \infty} e^{rn(\varphi_n - 1)}$ is a CFP. Letting $r = \frac{1}{n}$, we get that $\varphi = (\varphi^{1/n})^n$ is infinitely divisible.

Corollary 16.8 Let $\varphi : \mathbb{R} \rightarrow \mathbb{C}$ be continuous at 0. φ is an infinitely divisible CFP if and only if there is a sequence $(\varphi_n)_{n \in \mathbb{N}}$ of CFPs such that $\varphi_n^n(t) \rightarrow \varphi(t)$ for all $t \in \mathbb{R}$.

Proof One implication has been shown already in Corollary 16.7. Hence, let φ be an infinitely divisible CFP. Then $\varphi_n = \varphi^{1/n}$ serves the purpose.

Corollary 16.9 If $(\mu_n)_{n \in \mathbb{N}}$ is a (weakly) convergent sequence of infinitely divisible probability measures on $\mathbb{R}$, then $\mu = \lim_{n \rightarrow \infty} \mu_n$ is infinitely divisible.

Proof Apply Theorem 16.6, where φ_n is the CFP of $\mu_n^{*1/n}$.

Corollary 16.10 If $\mu \in \mathcal{M}_1(\mathbb{R})$ is infinitely divisible, then there exists a continuous convolution semigroup $(\mu_t)_{t \geq 0}$ with $\mu_1 = \mu$ and a stochastic process $(X_t)_{t \geq 0}$ with independent, stationary increments $X_t - X_s \sim \mu_{t-s}$ for $t > s$.

Proof Let φ be the CFP of μ . The existence of the convolution semigroup follows by Corollaries 16.8 and 16.7 if we define μ_r by φ^r . Since $\varphi^r(t) \neq 0$ for all $t \in \mathbb{R}$, we have $\varphi^r \rightarrow 1$ for $r \rightarrow 0$ and thus the semigroup is continuous. Finally, Theorem 14.47 implies the existence of the process X .

Corollary 16.11 If φ is an infinitely divisible CFP, then there exists a $\gamma > 0$ with $|\varphi(t)| \geq \frac{1}{2}e^{-\gamma t^2}$ for all $t \in \mathbb{R}$. In particular, for $\alpha > 2, t \mapsto e^{-|t|^\alpha}$ is not a CFP.

Proof This is a direct consequence of (16.3).

Proof of Theorem 16.5 As every CPoi_{v_n} is infinitely divisible, by Corollary 16.9, the weak limit is also infinitely divisible.

Now let μ be infinitely divisible with CFP φ . Fix probability measures μ_n with CFP φ_n as in Corollary 16.8. By Theorem 16.6, $e^{n(\varphi_n-1)} \xrightarrow{n \rightarrow \infty} \varphi$; hence we have $\text{CPoi}_{n\mu_n} \xrightarrow{n \rightarrow \infty} v$.

Without proof, we quote the following strengthening of Corollary 16.8 that relies on a finer analysis using the arguments from the proof of Theorem 16.6.

Theorem 16.12 Let $(\varphi_{n,l}; l = 1, \dots, k_n, n \in \mathbb{N})$ be an array of CFPs with the property

$$\sup_{L>0} \limsup_{n \rightarrow \infty} \sup_{t \in [-L, L]} \sup_{l=1, \dots, k_n} |\varphi_{n,l}(t) - 1| = 0 \quad (16.4)$$

Assume that, for every $t \in \mathbb{R}$, the limit $\varphi(t) := \lim_{n \rightarrow \infty} \prod_{l=1}^{k_n} \varphi_{n,l}(t)$ exists and that φ is continuous at 0. Then φ is an infinitely divisible CFP.

Proof See, e.g., [54, Chapter Xv.7].

In the special case where for every n , the individual $\varphi_{n,l}$ are equal and where $k_n \xrightarrow{n \rightarrow \infty} \infty$, Eq. (16.4) holds automatically if the product converges to a continuous function. Thus, the theorem is in fact an improvement of Corollary 16.8.

The benefit of this theorem will become clear through the following observation. Let $(X_{n,l}; l = 1, \dots, k_n, n \in \mathbb{N})$ be an array of real random variables with CFPs $\varphi_{n,l}$. This array is a null array if and only if (16.4) holds. In fact, if $\mathbf{P}[|X_{n,l}| > \varepsilon] < \delta$, then we have $|\varphi_{n,l}(t) - 1| \leq 2\varepsilon + \delta$ for all $t \in [-1/\varepsilon, 1/\varepsilon]$. Hence (16.4) holds if the array $(X_{n,l})$ is a null array. On the other hand, (16.4) implies $\varphi_{n,l_n} \xrightarrow{n \rightarrow \infty} 1$ for every sequence (l_n) with $l_n \leq k_n$. Hence $X_{n,l_n} \xrightarrow{n \rightarrow \infty} 0$ in probability.

From these considerations and from Theorem 16.12, we conclude the following theorem.

Theorem 16.13 Let $(X_{n,l}; l = 1, \dots, k_n, n \in \mathbb{N})$ be an independent null array of real random variables. If there exists a random variable S with

$$X_{n,1} + \dots + X_{n,k_n} \xrightarrow{n \rightarrow \infty} S,$$

then S is infinitely divisible.

As a direct application of Theorem 16.5, we give a complete description of the class of infinitely divisible probability measures on $[0, \infty)$ in terms of their Laplace transforms. The following theorem is of independent interest. Here, however, it is primarily used to provide familiarity with the techniques that will be needed for the more challenging classification of the infinitely divisible probability measures on $\mathbb{R}$.

Theorem 16.14 (Lévy-Khinchin formula on $[0, \infty)$) Let $\mu \in \mathcal{M}_1([0, \infty))$ and let $u : [0, \infty) \rightarrow [0, \infty), t \mapsto -\log \int e^{-tx} \mu(dx)$ be the log-Laplace transform μ . μ is infinitely divisible if and only if there exists an $\alpha \geq 0$ and $\alpha\sigma$ -finite measure $v \in \mathcal{M}((0, \infty))$ with

$$\int (1 \wedge x) v(dx) < \infty \quad (16.5)$$

and such that

$$u(t) = \alpha t + \int (1 - e^{-tx}) v(dx) \quad \text{for } t \geq 0 \quad (16.6)$$

In this case, the pair (α, v) is unique. v is called the canonical measure or Lévy measure of μ , and α is called the deterministic part.

Proof " $\implies$ " First assume μ is infinitely divisible. The case $\mu = \delta_0$ is trivial. Now let $\mu \neq \delta_0$; hence $u(1) > 0$.

By Theorem 16.5, there exist $v_1, v_2, \dots \in \mathcal{M}_f(\mathbb{R} \setminus \{0\})$ with $\text{CPoi}_{v_n} \xrightarrow{n \rightarrow \infty} \mu$. Evidently, we can assume $v_n((-\infty, 0)) = 0$. If we define $u_n(t) := \int (1 - e^{-tx}) v_n(dx)$, then (as in (16.1)) $u_n(t) \xrightarrow{n \rightarrow \infty} u(t)$ for all $t \geq 0$. In particular, $u_n(1) > 0$ for sufficiently large n . Define $\tilde{v}_n \in \mathcal{M}_1([0, \infty))$ by $\tilde{v}_n(dx) := \frac{1-e^{-x}}{u_n(1)} v_n(dx)$. Hence, for all $t \geq 0$,

$$\int e^{-tx} \tilde{v}_n(dx) = \frac{u_n(t+1) - u_n(t)}{u_n(1)} \xrightarrow{n \rightarrow \infty} \frac{u(t+1) - u(t)}{u(1)}.$$

Therefore, the weak limit $\tilde{v} := w - \lim \tilde{v}_n$ in $\mathcal{M}_1([0, \infty))$ exists and is uniquely determined by u . Let $\alpha := \tilde{v}(\{0\})u(1)$ and define $v \in \mathcal{M}((0, \infty))$ by

$$v(dx) = u(1) (1 - e^{-x})^{-1} \mathbb{1}_{(0, \infty)}(x) \tilde{v}(dx)$$

Since $1 \wedge x \leq 2(1 - e^{-x})$ for all $x \geq 0$, clearly

$$\int (1 \wedge x) v(dx) \leq 2 \int (1 - e^{-x}) v(dx) \leq 2u(1) < \infty$$

For all $t \geq 0$, the function (compare (15.8))

$$f_t : [0, \infty) \rightarrow [0, \infty), \quad x \mapsto \begin{cases} \frac{1-e^{-tx}}{1-e^{-x}}, & \text{if } x > 0 \\ t, & \text{if } x = 0 \end{cases}$$

is continuous and bounded (by $t \wedge 1$). Hence we have

$$\begin{aligned} u(t) &= \lim_{n \rightarrow \infty} u_n(t) = \lim_{n \rightarrow \infty} u_n(1) \int f_t d\tilde{v}_n \\ &= u(1) \int f_t d\tilde{v} = \alpha t + \int (1 - e^{-tx}) v(dx) \end{aligned}$$

" $\Leftarrow$ " Now assume that α and v are given. Define the intervals $I_0 = [1, \infty)$ and $I_k = [1/(k+1), 1/k)$ for $k \in \mathbb{N}$. Let $X_0, X_1, \dots$ be independent random variables with $\mathbf{P}_{X_k} = \text{CPoi}_{(v|_{I_k})}$ for $k = 0, 1, \dots$, and let $X := \alpha + \sum_{k=0}^{\infty} X_k$. For every $k \in \mathbb{N}$, we have $\mathbf{E}[X_k] = \int_{I_k} xv(dx)$; hence $\sum_{k=1}^{\infty} \mathbf{E}[X_k] = \int_{(0,1)} xv(dx) < \infty$. Thus $X < \infty$ almost surely and $\alpha + \sum_{k=0}^n X_k \xrightarrow{n \rightarrow \infty} X$. Therefore,

$$-\log \mathbf{E}[e^{-tX}] = \alpha t - \sum_{k=0}^{\infty} \log \mathbf{E}[e^{-tX_k}] = \alpha t + \int (1 - e^{-tx}) v(dx)$$

Example 16.15 For an infinitely divisible distribution μ on $[0, \infty)$, we can compute the Lévy measure ν by the vague limit

$$v = v - \lim_{n \rightarrow \infty} n\mu^{*1/n} \Big|_{(0, \infty)} \quad (16.7)$$

Often α is also easy to obtain (e.g., via the representation from Exercise 16.1.3). For example, for the Gamma distribution, we get $\alpha = 0$ and

$$n\Gamma_{\theta, 1/n}(A) = \frac{\theta^{1/n}}{\Gamma(1/n)/n} \int_A x^{(1/n)-1} e^{-\theta x} dx \xrightarrow{n \rightarrow \infty} \int_A x^{-1} e^{-\theta x} dx$$

hence $\nu(dx) = x^{-1} e^{-\theta x} dx$.

For infinitely divisible distributions on $\mathbb{R}$, we would like to obtain a description similar to that in

the preceding theorem. However, an infinitely divisible real random variable X is not simply the difference of two infinitely divisible nonnegative random variables, as the normal distribution shows. In addition, we have more freedom if, as in the last proof, we want to express X as a sum of independent random variables X_k .

Hence we define the real random variable X as the sum of independent random variables,

$$X = b + X^N + X_0 + \sum_{k=1}^{\infty} (X_k - \alpha_k), \quad (16.8)$$

where $b \in \mathbb{R}$, $X^N = \mathcal{N}_{0,\sigma^2}$ for some $\sigma^2 \geq 0$ and $\mathbf{P}_{X_k} = \text{CPoi}_{v_k}$ with intensity measure ν_k that is concentrated on $I_k := (-1/k, -1/(k+1)] \cup [1/(k+1), 1/k)$ (with the convention $1/0 = \infty$), $k \in \mathbb{N}_0$. Furthermore, $\alpha_k = \mathbf{E}[X_k] = \int x \nu_k(dx)$ for $k \geq 1$. In order for the series to converge almost surely, it is sufficient (and also necessary, as a simple application of Kolmogorov's three-series theorem shows) that

$$\sum_{k=1}^{\infty} \text{Var}[X_k] < \infty \quad (16.9)$$

(In contrast to the situation in Theorem 16.14, here it is not necessary to have $\sum_{k=1}^{\infty} \mathbf{E}[|X_k - \alpha_k|] < \infty$. This allows for greater freedom in the choice of v than in the case of nonnegative random variables.) Now $\text{Var}[X_k] = \int x^2 \nu_k(dx)$. Hence, if we let $v = \sum_{k=0}^{\infty} \nu_k$, then (16.9) is equivalent to $\int_{(-1,1)} x^2 v(dx) < \infty$. As v_0 is always finite, this in turn is equivalent to $\int (x^2 \wedge 1) \nu(dx) < \infty$.

Definition 16.16 A σ -finite measure ν on $\mathbb{R}$ is called a canonical measure if $\nu(\{0\}) = 0$ and

$$\int (x^2 \wedge 1) \nu(dx) < \infty \quad (16.10)$$

If $\sigma^2 \geq 0$ and $b \in \mathbb{R}$, then (σ^2, b, ν) is called a canonical triple.

To every canonical triple, by (16.8) there corresponds an infinitely divisible random variable. Define

$$\psi_0(t) = \log \mathbf{E}[e^{itX_0}] = \int_{I_0} (e^{itx} - 1) \nu(dx)$$

For $k \in \mathbb{N}$, let

$$\psi_k(t) = \log \mathbf{E}[e^{it(X_k - \alpha_k)}] = \int_{I_k} (e^{itx} - 1 - itx) \nu(dx)$$

Hence

$$\psi(t) := \log \mathbf{E}[e^{itX}] = -\frac{\sigma^2}{2}t^2 + ibt + \sum_{k=0}^{\infty} \psi_k(t)$$

satisfies the Lévy-Khinchin formula

$$\psi(t) = -\frac{\sigma^2}{2}t^2 + ibt + \int (e^{itx} - 1 - itx \mathbb{1}_{\{|x|<1\}}) \nu(dx) \quad (16.11)$$

Theorem 16.17 (Lévy-Khinchin formula) Let $\mu \in \mathcal{M}_1(\mathbb{R})$ and

$$\psi(t) := \log \int e^{itx} \mu(dx)$$

μ is infinitely divisible if and only if there exists a canonical triple (σ^2, b, ν) such that (16.11) holds. By (16.11), this triple is uniquely determined.

Again, ν is called the Lévy measure of μ , σ^2 is called the Gaussian coefficient and b is called the centering constant.

Proof We have shown already that via (16.11) every canonical triple (σ^2, b, ν) corresponds to an infinitely divisible distribution μ . It remains to show:

- (i) A canonical triple is uniquely determined by (16.11).
 - (ii) For every infinitely divisible distribution, there exists a canonical triple such that (16.11) holds.
- (i) Uniqueness. Define $g_t(x) = e^{itx} - 1 - itx \mathbb{1}_{\{|x| < 1\}}$. For every $x \neq 0$, we have

$$2 \geq \left| \frac{g_t(x)}{t^2 (1 \wedge x^2)} \right| \xrightarrow{t \rightarrow \infty} 0.$$

Since (16.10) holds, by the dominated convergence theorem,

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{\psi(t)}{t^2} &= -\frac{\sigma^2}{2} + \lim_{t \rightarrow \infty} \frac{ib}{t} + \lim_{t \rightarrow \infty} \int_{-\infty}^{\infty} \left(\frac{g_t(x)}{t^2 (1 \wedge x^2)} \right) (1 \wedge x^2) \nu(dx) \\ &= -\frac{\sigma^2}{2} \end{aligned} \quad (16.12)$$

This implies the uniqueness of σ^2 . Thus we can and will assume $\sigma^2 = 0$ in the following. Define

$$\bar{\psi}(t) = \psi(t) - \frac{1}{2} \int_{t-1}^{t+1} \psi(s) ds \quad (16.13)$$

Then

$$\bar{\psi}(t) = \int_{\mathbb{R}} e^{itx} \left(1 - \frac{1}{2} \int_{-1}^1 e^{isx} ds \right) \nu(dx) = \int e^{itx} h(x) \nu(dx) \quad (16.14)$$

where $h(x) = 1 - \frac{\sin(x)}{x}$ for $x \neq 0$ and $h(0) = 0$. Define $\hat{h}(x) = h(x)/(1 \wedge x^2)$ for $x \neq 0$ and $\hat{h}(0) = 1/6$. Clearly, h and $\hat{h}$ are bounded and continuous and

$$0 < 1 - \sin(1) \leq \hat{h}(x) \leq \frac{3}{2} \quad \text{for all } x \in \mathbb{R}$$

$\bar{\psi}$ is the characteristic function of $\tilde{\nu} \in \mathcal{M}_f(\mathbb{R})$, where $\tilde{\nu}(dx) = h(x)\nu(dx)$. Hence $\tilde{\nu}$ is uniquely determined by ψ . Since $\nu(dx) = (\mathbb{1}_{\{x \neq 0\}}/h(x)) \tilde{\nu}(dx)$, ν is also uniquely determined by ψ . Now the number b is the difference of the remaining terms.

- (ii) Existence of a canonical triple. Let μ be infinitely divisible and let

$$\psi(t) = \log \int e^{itx} \mu(dx)$$

Clearly, $\text{Im}(\psi)$ is odd and $\text{Re}(\psi(t)) \leq 0$ for all $t \in \mathbb{R}$. Hence $\bar{\psi}(0) \geq 0$ (with $\bar{\psi}$ from (16.13)) and $\bar{\psi}(0) = 0$ if $\text{Re} \psi(t) = 0$ for all $t \in [-1, 1]$. By Exercise 15.2.4, this is the case if and only if $\mu = \delta_b$ for some $b \in \mathbb{R}$. In this case, $(0, b, 0)$ is the corresponding canonical triple.

Now assume $\bar{\psi}(0) > 0$. By Theorem 16.5, there exists a sequence $(\nu_n)_{n \in \mathbb{N}}$ in $\mathcal{M}_f(\mathbb{R})$ with $\text{CPoi}_{\nu_n} \xrightarrow{n \rightarrow \infty} \mu$ and $\nu_n(\{0\}) = 0$ for any $n \in \mathbb{N}$. Define

$$b_n = \int x \mathbb{1}_{\{|x| < 1\}} \nu_n(dx)$$

Then, by (16.1) and with g_t from (i),

$$\psi_n(t) := \log \int e^{itx} \text{CPoi}_{v_n}(dx) = \int (e^{itx} - 1) v_n(dx) = \int g_t dv_n + ib_n t$$

As in (16.14), we have

$$\bar{\psi}_n(t) := \psi_n(t) - \frac{1}{2} \int_{t-1}^{t+1} \psi_n(s) ds = \int e^{itx} h(x) v_n(dx)$$

As $\psi_n \xrightarrow{n \rightarrow \infty} \psi$ converges uniformly on compact sets (Theorem 15.23(i)), and since ψ is continuous and thus locally bounded, we have $\bar{\psi}_n \xrightarrow{n \rightarrow \infty} \bar{\psi}$ pointwise. Therefore,

$$\int e^{itx} h(x) v_n(dx) \xrightarrow{n \rightarrow \infty} \bar{\psi}(t) \quad (16.15)$$

In particular, $\bar{\psi}_n(0) > 0$ for large n . If we let $\tilde{v}_n(dx) = (h(x)/\bar{\psi}_n(0)) v_n(dx) \in \mathcal{M}_1(\mathbb{R})$, then $\int e^{itx} \tilde{v}_n(dx) \xrightarrow{n \rightarrow \infty} \bar{\psi}(t)/\bar{\psi}(0)$ and the right-hand side is continuous. Hence, by Lévy's continuity theorem, there is a $\tilde{v} \in \mathcal{M}_1(\mathbb{R})$ with $\tilde{v}_n \xrightarrow{n \rightarrow \infty} \tilde{v}$ and

$$\bar{\psi}(t) = \bar{\psi}(0) \int e^{itx} \tilde{v}(dx)$$

Let $\sigma^2 := -6\bar{\psi}(0)\tilde{v}(\{0\})$ and define a canonical measure ν by

$$v(dx) = \frac{\bar{\psi}(0)}{h(x)} \mathbb{1}_{\{x \neq 0\}} \tilde{v}(dx)$$

The map (compare (15.8))

$$f_t : \mathbb{R} \rightarrow \mathbb{C}, \quad x \mapsto \begin{cases} \frac{g_t(x)}{h(x)}, & \text{if } x \neq 0 \\ -3t^2, & \text{if } x = 0 \end{cases}$$

is bounded and continuous. By construction, we have

$$\int g_t dv_n = \bar{\psi}_n(0) \int f_t d\tilde{v}_n \xrightarrow{n \rightarrow \infty} \bar{\psi}(0) \int f_t d\tilde{v} = -\frac{\sigma^2}{2} t^2 + \int g_t dv$$

Hence also the limit

$$itb := \lim_{n \rightarrow \infty} itb_n = \lim_{n \rightarrow \infty} \left(\psi_n(t) - \int g_t dv_n \right) = \psi(t) + \frac{\sigma^2}{2} t^2 - \int g_t dv$$

exists, and we have

$$\psi(t) = -\frac{\sigma^2}{2} t^2 + ibt + \int g_t dv$$

Remark 16.18 There are many versions of the Lévy-Khinchin formula

$$\psi(t) = -\frac{\sigma^2}{2} t^2 + ibt + \int (e^{itx} - 1 - itf(x)) v(dx)$$

that differ in the function $itf(x)$ that is subtracted for the centering in the integral. We chose $f(x) = x \mathbb{1}_{\{|x| < 1\}}$ since this fits best to the construction with the random variables X_k . However, for a given canonical measure ν , any function $\tilde{f}$ for which $\int |f - \tilde{f}| d\nu < \infty$ holds is possible; that is, every $\tilde{f}$ for which

$|f(x) - \tilde{f}(x)|/(1 \wedge x^2)$ is bounded. One common function is, e.g., $\tilde{f}(x) = \sin(x)$. The Lévy measure and the Gaussian coefficient σ^2 do not change but the b differs:

$$\tilde{b} - b = \int (f - \tilde{f}) d\nu$$

If ν is a measure that is concentrated on $(0, \infty)$ and such that $\int (1 \wedge x) \nu(dx) < \infty$ holds, then this f is integrable with respect to ν and can thus be replaced by $\tilde{f} = 0$. Hence we recover Theorem 16.14 as a special case. However, condition (16.10) is weaker than $\int (1 \wedge x) \nu(dx) < \infty$ and thus describes a larger class of measures than is considered in Theorem 16.14. This implies that to a canonical triple $(b, 0, \nu)$ with $\nu((-\infty, 0)) = 0$ and $\int (1 \wedge x) \nu(dx) = \infty$, there corresponds an infinitely divisible probability distribution μ that is not concentrated on $[0, \infty)$, no matter how b is chosen.

For a given infinitely divisible distribution μ , we can compute the canonical measure ν as the vague limit

$$\nu = v_{n \rightarrow \infty} n \mu^{*1/n} \Big|_{(0, \infty)} \quad (16.16)$$

Example 16.19 For the Cauchy distribution Cau_a with $\psi(t) = -a|t|$, by symmetry, we get $b = 0$ and, by (16.12), $\sigma^2 = -2 \lim_{t \rightarrow \infty} \psi(t)/t^2 = 0$. Finally, if $A \subset \mathbb{R}$ with $(-\varepsilon, \varepsilon) \cap A = \emptyset$ for some $\varepsilon > 0$, then

$$n \text{Cau}_{1/n}(A) = \frac{1}{\pi} \int_A \frac{n^2}{1 + (nx)^2} dx \xrightarrow{n \rightarrow \infty} \frac{1}{\pi} \int_A \frac{1}{x^2} dx$$

Hence Cau_1 has the canonical triple $(0, 0, (\pi x^2)^{-1} dx)$.

Exercise 16.1.1 Use a variance argument to show that an infinitely divisible distribution that is concentrated on a bounded interval is a Dirac measure.

Exercise 16.1.2 Let φ be infinitely divisible, and for every $n \in \mathbb{N}$, let φ_n be a CFP with $\varphi_n^n = \varphi$. Use Lévy's continuity theorem to show that $\varphi_n \xrightarrow{n \rightarrow \infty} 1$ uniformly on compact sets $\varphi_n \xrightarrow{n \rightarrow \infty} 1$. Conclude that $\varphi(t) \neq 0$ for all $t \in \mathbb{R}$.

Exercise 16.1.3 Under the conditions of Theorem 16.14, show that

$$\alpha = \sup\{x \geq 0 : \mu([0, x)) = 0\}$$

5.16.2 16.2 Stable Distributions

A distribution μ on the real numbers is called stable if for any $n \in \mathbb{N}$, the n -fold convolution μ^{*n} equals μ up to an affine linear transformation. Hence stability can be interpreted as self-similarity. We first show that the class of stable distributions is rather simple and can easily be parameterized. Then we quote results which say that stable distributions are exactly those distributions that occur as limits of sums of i.i.d. random variables.

Symmetric Stable Distributions

For $\alpha \in (0, 2)$, let

$$\theta_\alpha := \int_{\mathbb{R}} (1 - \cos(x)) |x|^{-\alpha-1} dx = \begin{cases} -2\Gamma(-\alpha) \cos(\alpha\pi/2), & \text{if } \alpha \neq 1 \\ \pi, & \text{if } \alpha = 1 \end{cases}$$

(Note that the integral diverges for $\alpha \in \mathbb{R} \setminus (0, 2)$.) Then $\nu_\alpha(dx) = \theta_\alpha^{-1} |x|^{-\alpha-1} dx$ is a canonical measure since

$$\int (1 \wedge x^2) v_\alpha(dx) = 2\theta_\alpha^{-1} (\alpha^{-1} + (2 - \alpha)^{-1}) < \infty$$

Let ψ_α be the logarithm of the characteristic function that corresponds to the infinitely divisible measure μ_α with canonical triple $(0, 0, v_\alpha)$. By the Lévy-Khinchin formula, we have

$$\begin{aligned} \psi_\alpha(t) &= \int_{-\infty}^{\infty} (e^{itx} - 1 - itx \mathbb{1}_{\{|x| < 1\}}) \theta_\alpha^{-1} |x|^{-\alpha-1} dx \\ &= -\theta_\alpha^{-1} \int_{-\infty}^{\infty} (1 - \cos(tx)) |x|^{-\alpha-1} dx \\ &= -|t|^\alpha \end{aligned}$$

Hence $\varphi_\alpha(t) := e^{-|t|^\alpha}$ is the characteristic function of the infinitely divisible measure μ_α , which is called the symmetric stable distribution with index α . The name is due to the fact that, for i.i.d. random variables $X_1, X_2, \dots$ that are μ_α -distributed, we have

$$X_1 + \dots + X_n \stackrel{\mathcal{D}}{=} n^{1/\alpha} X_n \quad \text{for all } n \in \mathbb{N} \quad (16.17)$$

General Stable Distributions

Motivated by equation (16.17), we present a somewhat more general notion of stability of a distribution.

Definition 16.20 (Stable distribution) Let $\mu \in \mathcal{M}_1(\mathbb{R})$ be a probability distribution on the real numbers that is not concentrated in one point. Assume that $X_1, X_2, \dots$ are i.i.d. random variables with distribution μ . The distribution μ is said to be stable in the broad sense if there exist nonnegative numbers $a_1, a_2, \dots$ and real numbers $d_1, d_2, \dots$ such that

$$X_1 + \dots + X_n \stackrel{\mathcal{D}}{=} a_n X_1 + d_n \quad \text{for all } n \in \mathbb{N}. \quad (16.18)$$

μ is called stable (in the strict sense), if (16.18) holds with $d_1 = d_2 = \dots = 0$.

μ is called stable in the broad sense with index $\alpha \in (0, 2]$, if (16.18) holds with $a_n = n^{1/\alpha}, n \in \mathbb{N}$. It is called stable (in the strict sense) with index $\alpha \in (0, 2]$, if in addition, we can choose $d_1 = d_2 = \dots = 0$.

Remark 16.21 If μ is stable in the broad sense, then it is infinitely divisible.

Theorem 16.22 Let μ be stable in the broad sense.

- (i) There is an $\alpha \in (0, 2]$ such that μ is stable in the broad sense with index α .
- (ii) If $\alpha = 2$, then μ is a normal distribution.
- (iii) If $\alpha \in (0, 2)$, then the Lévy measure ν of μ has the density

$$\frac{\nu(dx)}{dx} = \begin{cases} c^- (-x)^{-\alpha-1}, & \text{if } x < 0 \\ c^+ x^{-\alpha-1}, & \text{if } x > 0 \end{cases} \quad (16.19)$$

for some $c^-, c^+ \geq 0, c^- + c^+ > 0$.

- (iv) If $\alpha \neq 1$, then there exists a $b \in \mathbb{R}$ such that $\mu * \delta_{-b}$ is stable with index α .
- (v) If $\alpha = 1$, then $d_n = (c^+ - c^-) n \log(n), n \in \mathbb{N}$. If $c^- = c^+$, then μ is a Cauchy distribution.

Remark 16.23 If μ is infinitely divisible with Lévy measure ν given by (16.19), then $\psi(t) := \log \int e^{itx} \mu(dx)$ is given by

$$\psi(t) = \begin{cases} |t|^\alpha \Gamma(-\alpha) \left[(c^+ + c^-) \cos\left(\frac{\pi\alpha}{2}\right) + i(c^+ - c^-) \sin\left(\frac{\pi\alpha}{2}\right) \right], & \alpha \neq 1, \\ -|t| (c^+ + c^-) \left[\frac{\pi}{2} + i \operatorname{sign}(t) (c^+ - c^-) \log(|t|) \right], & \alpha = 1. \end{cases} \quad (16.20)$$

Lemma 16.24 Let μ be infinitely divisible with canonical triple (σ^2, b, ν) ; that is, with log-characteristic function $\psi(t) := \log \left(\int e^{itx} \mu(dx) \right)$ given by

$$\psi(t) = -\frac{\sigma^2}{2}t^2 + ibt + \int (e^{itx} - 1 - itx\mathbb{1}_{\{|x|<1\}}) \nu(dx)$$

Further, let $a > 0, d \in \mathbb{R}, n \in \mathbb{N}$ and let $X, X_1, \dots, X_n$ be i.i.d. random variables with distribution μ .

- (i) The canonical triple of $X_1 + \dots + X_n$ is $(n\sigma^2, nb, n\nu)$.
(ii) The canonical triple of $aX + d$ is $(a^2\sigma^2, \tilde{b}, v \circ m_a^{-1})$, where $m_a : \mathbb{R} \rightarrow \mathbb{R}, x \mapsto ax$ is the multiplication by a and

$$\tilde{b} := ab + d + a \int (\mathbb{1}_{\{|x|<1/a\}} - \mathbb{1}_{\{|x|<1\}}) xv(dx) \quad (16.21)$$

Proof (i) The log-characteristic function of $X_1 + \dots + X_n$ is $n\psi$.

- (ii) The log-characteristic function of $aX + d$ is

$$\begin{aligned} \psi_{aX+d}(t) &= \psi(at) + idt \\ &= -\frac{a^2\sigma^2}{2}t^2 + i(ab + d)t + \int (e^{iatx} - 1 - iatx\mathbb{1}_{\{|x|<1\}}) v(dx) \\ &= -\frac{a^2\sigma^2}{2}t^2 + i\tilde{b}t + \int (e^{iatx} - 1 - iatx\mathbb{1}_{\{|x|<1/a\}}) v(dx) \\ &= -\frac{a^2\sigma^2}{2}t^2 + i\tilde{b}t + \int (e^{itx} - 1 - itx\mathbb{1}_{\{|x|<1\}}) v \circ m_a^{-1}(dx) \end{aligned}$$

Lemma 16.25 (Scaling of the canonical triple) Under the assumptions of Theorem 16.22, let (σ^2, b, v) be the canonical triple of μ .

- (i) We have

$$(a_n^2 - n)\sigma^2 = 0 \quad \text{for all } n \in \mathbb{N} \quad (16.22)$$

and (with m_{a_n} as in Lemma 16.24)

$$nv = v \circ m_{a_n}^{-1} \quad \text{for all } n \in \mathbb{N}. \quad (16.23)$$

- (ii) If $v = 0$, then $a_n = n^{1/2}$ for all $n \in \mathbb{N}$ and

$$d_n = b(n - n^{1/2}). \quad (16.24)$$

- (iii) Assume that $\alpha \in (0, 2)$, $a_n = n^{1/\alpha}$, and that v is given by (16.19). Then we have

$$d_n = \left(b + \frac{c^+ - c^-}{\alpha - 1}\right)(n - n^{1/\alpha}) \quad \text{if } \alpha \neq 1 \quad (16.25)$$

and

$$d_n = (c^+ - c^-)n \log(n) \quad \text{if } \alpha = 1 \quad (16.26)$$

Proof (i) Let $(a_n^2\sigma^2, \tilde{b}_n, v \circ m_{a_n}^{-1})$ be the canonical triple of $a_nX + d_n$ as determined in the preceding lemma and let $(n\sigma^2, nb, n\nu)$ be the canonical triple of $X_1 + \dots + X_n$. By (16.18) and due to the uniqueness of the canonical triple (Theorem 16.17), we infer $a_n^2\sigma^2 = n\sigma^2$, $\tilde{b}_n = nb$ and $v \circ m_{a_n}^{-1} = n\nu$.

- (ii) If $v = 0$, then $\sigma^2 > 0$, since by assumption, μ is not concentrated in one point. Hence, by (16.22), we get $a_n = n^{1/2}$. By virtue of Lemma 16.24(ii), we have $nb = \tilde{b}_n = bn^{1/2} + d_n$ and thus (16.24) holds.

- (iii) Using (16.21), we compute $\tilde{b}_n$ more explicitly:

$$\begin{aligned}
nb &= \tilde{b}_n = bn^{1/\alpha} + d_n - n^{1/\alpha} \int \mathbb{1}_{\{n^{-1/\alpha} \leq |x| < 1\}} xv(dx) \\
&= bn^{1/\alpha} + d_n - n^{1/\alpha} (c^+ - c^-) \int_{n^{-1/\alpha}}^1 x^{-\alpha} dx \\
&= bn^{1/\alpha} + d_n - (c^+ - c^-) \begin{cases} (1 - \alpha)^{-1} (n^{1/\alpha} - n), & \text{if } \alpha \neq 1 \\ n \log(n), & \text{if } \alpha = 1 \end{cases}
\end{aligned}$$

Rearranging terms yields (16.25) and (16.26).

Proof of Theorem 16.22 We distinguish the cases $\liminf_{n \rightarrow \infty} a_n n^{-1/2} < \infty$ and " $= \infty$ ".

Case 1. Assume that $\liminf_{n \rightarrow \infty} a_n n^{-1/2} < \infty$. Let $C \in [1, \infty)$ and let $(n_k)_{k \in \mathbb{N}}$ be a subsequence such that $a_{n_k} n_k^{-1/2} \leq C$ for any $k \in \mathbb{N}$. Then for any $x \in \mathbb{R} \setminus \{0\}$, we have

$$C^2 \geq n_k^{-1} (1 \vee a_{n_k}^2) \geq \frac{n_k^{-1} (1 \wedge a_{n_k}^2 x^2)}{1 \wedge x^2} \xrightarrow{k \rightarrow \infty} 0.$$

Using (16.23) and (16.10), the dominated convergence theorem yields

$$\int_{-\infty}^{\infty} (1 \wedge x^2) v(dx) = \int_{-\infty}^{\infty} \frac{n_k^{-1} (1 \wedge a_{n_k}^2 x^2)}{1 \wedge x^2} (1 \wedge x^2) v(dx) \xrightarrow{k \rightarrow \infty} 0$$

That is, we have $v = 0$. By Lemma 16.25(ii), we see that $\mu * \delta_{-b}$ is stable with index 2. This shows (ii).

Case 2. Assume that

$$a_n n^{-1/2} \xrightarrow{n \rightarrow \infty} \infty. \quad (16.27)$$

By (16.22), we have $\sigma^2 = 0$ and hence $v \neq 0$. We define the function

$$F(x) = \begin{cases} v([x, \infty)), & \text{if } x > 0 \\ v((-\infty, x]), & \text{if } x < 0 \end{cases}$$

Since we have $\nu \neq 0$, there is an $x_0 \in \mathbb{R} \setminus \{0\}$ such that $F(x_0) > 0$. By symmetry, we may assume that $x_0 > 0$. Using (16.23), we infer

$$nF(x) = F(x/a_n) \quad \text{for any } x \in \mathbb{R} \setminus \{0\}, n \in \mathbb{N}$$

and thus

$$F\left(\left(\frac{a_{n+1}}{a_n}\right)^k x_0\right) = \left(\frac{n}{n+1}\right)^k F(x_0) \quad \text{for any } k \in \mathbb{Z}$$

We can rephrase this as

$$F(x) = (x/x_0)^{-\alpha_n} F(x_0) \quad \text{for any } x \in \left\{(a_{n+1}/a_n)^k x_0 : k \in \mathbb{Z}\right\},$$

where $\alpha_n := \log((n+1)/n)/\log(a_{n+1}/a_n)$. Since F is monotone decreasing and since $F(x) \xrightarrow{x \rightarrow \infty} 0$, we have $\alpha_n > 0$ for all $n \in \mathbb{N}$, and

$$\left(\frac{m}{m+1}\right) \left(\frac{x}{x_0}\right)^{-\alpha_m} \leq \frac{F(x)}{F(x_0)} \leq \left(\frac{n+1}{n}\right) \left(\frac{x}{x_0}\right)^{-\alpha_n} \quad \text{for } x > 0, m, n \in \mathbb{N}.$$

Letting $x \rightarrow \infty$, we obtain $\alpha_m \geq \alpha_n$. By symmetry, we also get $\alpha_m \leq \alpha_n$. Hence, we define $\alpha := \alpha_1 > 0$ and get $a_n = n^{1/\alpha}$ for all $n \in \mathbb{N}$ (note that (16.18) implies $a_1 = 1$). By the assumption (16.27), we have $\alpha < 2$. This shows (i).

We have $F(1) = x_0^\alpha F(x_0) > 0$ and $F(x) = x^{-\alpha} F(1)$ for all $x > 0$. Similarly, we get $F(x) = (-x)^{-\alpha} F(-1)$ for $x < 0$ (with the same $\alpha \in (0, 2)$ since it is determined by the sequence $(a_n)_{n \in \mathbb{N}}$). Defining $c^+ = \alpha \nu([1, \infty))$ and $c^- := \alpha \nu((-\infty, -1])$, we get (16.19) and thus (iii) and (i).

The statements (iv) and (v) are immediate consequences of Lemma 16.25.

Convergence to Stable Distributions

To complete the picture, we cite theorems from [54, Chapter XVII.5] (see also [62] and [128]) that state that only stable distributions occur as limiting distributions of rescaled sums of i.i.d. random variables $X_1, X_2, \dots$.

In the following, let $X, X_1, X_2, \dots$ be i.i.d. random variables and for $n \in \mathbb{N}$, let $S_n = X_1 + \dots + X_n$.

Definition 16.26 (Domain of attraction) Let $\mu \in \mathcal{M}_1(\mathbb{R})$ be nontrivial. The domain of attraction $\text{Dom}(\mu) \subset \mathcal{M}_1(\mathbb{R})$ is the set of all distributions $\mathbf{P}_X$ with the property that there exist sequences of numbers $(a_n)_{n \in \mathbb{N}}$ and $(d_n)_{n \in \mathbb{N}}$ with

$$\frac{S_n - d_n}{a_n} \xrightarrow{n \rightarrow \infty} \mu$$

If μ is stable (in the broader sense) with index $\alpha \in (0, 2]$, then $\mathbf{P}_X$ is said to be in the domain of normal attraction if we can choose $a_n = n^{1/\alpha}$.

Theorem 16.27 Let $\mu \in \mathcal{M}_1(\mathbb{R})$ be nontrivial. Then $\text{Dom}(\mu) \neq \emptyset$ if and only if μ is stable (in the broader sense). In this case, $\mu \in \text{Dom}(\mu)$.

In the following, an important role is played by the function

$$U(x) := \mathbf{E} [X^2 \mathbb{1}_{\{|X| \leq x\}}] \quad (16.28)$$

A function $H : (0, \infty) \rightarrow (0, \infty)$ is called slowly varying at ∞ if

$$\lim_{x \rightarrow \infty} \frac{H(\gamma x)}{H(x)} = 1 \quad \text{for all } \gamma > 0$$

In the following, we assume that there exists an $\alpha \in (0, 2]$ such that

$$U(x)x^{\alpha-2} \quad \text{is slowly varying at } \infty \quad (16.29)$$

Theorem 16.28

(i) If $\mathbf{P}_X$ is in the domain of attraction of some distribution, then there exists an $\alpha \in (0, 2]$ such that (16.29) holds.

(ii) In the case $\alpha = 2$, we have: If $\mathbf{P}_X$ is not concentrated at one point, then (16.29) implies that $\mathbf{P}_X$ is in the domain of attraction of some distribution.

(iii) In the case $\alpha \in (0, 2)$, we have: $\mathbf{P}_X$ is in the domain of attraction of some distribution if and only if (16.29) holds and the limit

$$p := \lim_{x \rightarrow \infty} \frac{\mathbf{P}[X \geq x]}{\mathbf{P}[|X| \geq x]} \quad \text{exists.} \quad (16.30)$$

Theorem 16.29 Let $\mathbf{P}_X$ be in the domain of attraction of an α -stable distribution (that is, assume that condition (ii) or (iii) of Theorem 16.28 holds), and assume that $(a_n)_{n \in \mathbb{N}}$ is such that

$$C := \lim_{n \rightarrow \infty} \frac{nU(a_n)}{a_n^2} \in (0, \infty)$$

exists. Further, let μ be the stable distribution with index α whose characteristic function is given by (16.20) with $c^+ = Cp$ and $c^- = C(1-p)$.

- (i) In the case $\alpha \in (0, 1)$, let $b_n \equiv 0$.
- (ii) In the case $\alpha = 2$ and $\text{Var}[X] < \infty$, let $\mathbf{E}[X] = 0$.
- (iii) In the case $\alpha \in (1, 2]$, let $d_n = n\mathbf{E}[X]$ for all $n \in \mathbb{N}$.
- (iv) In the case $\alpha = 1$, let $d_n = na_n\mathbf{E}[\sin(X/a_n)]$ for all $n \in \mathbb{N}$.

Then

$$\frac{S_n - d_n}{a_n} \xrightarrow{n \rightarrow \infty} \mu$$

Corollary 16.30 If $\mathbf{P}_X$ is in the domain of attraction of a stable distribution with index α , then $\mathbf{E}[|X|^\beta] < \infty$ for all $\beta \in (0, \alpha)$ and $\mathbf{E}[|X|^\beta] = \infty$ if $\beta > \alpha$ and $\alpha < 2$.

Exercise 16.2.1 Let μ be an α -stable distribution and let φ be its characteristic function.

- (i) Show by a direct computation using only the definition of stability that $|\varphi(t) - 1| \leq C|t|^\alpha$ for t close to 0 (for some $C < \infty$).
- (ii) Use Exercise 15.3.2 to infer that $\mu = \delta_0$ if $\alpha > 2$.
- (iii) Modify the argument in order to show that for $\alpha > 2$, the α -stable distributions in the broad sense are also necessarily trivial.

Exercise 16.2.2 Show that the distribution on $\mathbb{R}$ with density

$$f(x) = \frac{1 - \cos(x)}{\pi x^2}$$

is not infinitely divisible.

Exercise 16.2.3 Let Φ be the distribution function of the standard normal distribution $\mathcal{N}_{0,1}$ and let $F : \mathbb{R} \rightarrow [0, 1]$ be defined by

$$F(x) = \begin{cases} 2(1 - \Phi(x^{-1/2})), & \text{if } x > 0 \\ 0, & \text{else.} \end{cases}$$

Show the following.

- (i) F is the distribution function of a $\frac{1}{2}$ -stable distribution.
- (ii) If $X_1, X_2, \dots$ are i.i.d. with distribution function F , then $\frac{1}{n} \sum_{k=0}^n X_k$ diverges almost surely for $n \rightarrow \infty$.

Hint: Compute the density of F , and show that the Laplace transform is given by $\lambda \mapsto e^{-\sqrt{2\lambda}}$.

Exercise 16.2.4 Which of the following distributions is in the domain of attraction of a stable distribution and for which parameter?

- (i) The distribution on $\mathbb{R}$ with density

$$f(x) = \begin{cases} \varrho \frac{1}{1+\alpha} |x|^\alpha, & \text{if } x < -1 \\ (1 - \varrho) \frac{1}{1+\beta} x^\beta, & \text{if } x > 1 \\ 0, & \text{else.} \end{cases}$$

Here $\alpha, \beta < -1$ and $\varrho \in [0, 1]$.

- (ii) The exponential distribution $\exp_\theta$ for $\theta > 0$.
- (iii) The distribution on $\mathbb{N}$ with weights cn^α if n is even and cn^β if n is odd. Here $\alpha, \beta < -1$ and $c = (2^\alpha \zeta(-\alpha) + (1 - 2^\beta) \zeta(-\beta))^{-1}$ (ζ is the Riemann zeta function) is the normalization constant.

5.17 17 Markov Chains

In spite of their simplicity, Markov processes with countable state space (and discrete time) are interesting mathematical objects with which a variety of real-world phenomena can be modeled. We give an introduction to the basic concepts and then study certain examples in more detail. The connection with discrete potential theory will be investigated later, in Chapter 19. Some readers might prefer to skip the somewhat technical construction of general Markov processes in Section 17.1.

There is a vast literature on Markov chains. For further reading, see, e.g., [21, 27, 64, 66, 91, 116, 123, 124, 143, 152].

5.17.1 17.1 Definitions and Construction

In the following, E is always a Polish space with Borel σ -algebra $\mathcal{B}(E)$, $I \subset \mathbb{R}$ and $(X_t)_{t \in I}$ is an E -valued stochastic process. We assume that $(\mathcal{F}_t)_{t \in I} = \mathbb{F} = \sigma(X)$ is the filtration generated by X .

Definition 17.1 We say that X has the Markov property (MP) if, for every $A \in \mathcal{B}(E)$ and all $s, t \in I$ with $s \leq t$,

$$\mathbf{P}[X_t \in A \mid \mathcal{F}_s] = \mathbf{P}[X_t \in A \mid X_s]$$

Remark 17.2 If E is a countable space, then X has the Markov property if and only if, for all $n \in \mathbb{N}$, all $s_1 < \dots < s_n < t$ and all $i_1, \dots, i_n, i \in E$ with $\mathbf{P}[X_{s_1} = i_1, \dots, X_{s_n} = i_n] > 0$, we have

$$\mathbf{P}[X_t = i \mid X_{s_1} = i_1, \dots, X_{s_n} = i_n] = \mathbf{P}[X_t = i \mid X_{s_n} = i_n] \quad (17.1)$$

In fact, (17.1) clearly implies the Markov property. On the other hand, if X has the Markov property, then (see(8.6)) $\mathbf{P}[X_t = i \mid \mathcal{F}_{s_n}](\omega) = \mathbf{P}[X_t = i \mid X_{s_n} = i_n]$ for almost all $\omega \in \{X_{s_n} = i_n\}$. Hence, for $A := \{X_{s_1} = i_1, \dots, X_{s_n} = i_n\}$ (using the

Markov property in the second equation),

$$\begin{aligned} & \mathbf{P}[X_t = i, X_{s_1} = i_1, \dots, X_{s_n} = i_n] \\ &= \mathbf{E}[\mathbf{E}[\mathbb{1}_{\{X_t = i\}} \mid \mathcal{F}_{s_n}] \mathbb{1}_A] = \mathbf{E}[\mathbf{E}[\mathbb{1}_{\{X_t = i\}} \mid X_{s_n}] \mathbb{1}_A] \\ &= \mathbf{E}[\mathbf{P}[X_t = i \mid X_{s_n} = i_n] \mathbb{1}_A] = \mathbf{P}[X_t = i \mid X_{s_n} = i_n] \mathbf{P}[A] \end{aligned}$$

Dividing both sides by $\mathbf{P}[A]$ yields (17.1).

Definition 17.3 Let $I \subset [0, \infty)$ be closed under addition and assume $0 \in I$. A stochastic process $X = (X_t)_{t \in I}$ is called a time-homogeneous Markov process with distributions $(\mathbf{P}_x)_{x \in E}$ on the space $(\Omega, \mathcal{A})$ if:

- (i) For every $x \in E$, X is a stochastic process on the probability space $(\Omega, \mathcal{A}, \mathbf{P}_x)$ with $\mathbf{P}_x[X_0 = x] = 1$.
- (ii) The map $\kappa : E \times \mathcal{B}(E)^{\otimes I} \rightarrow [0, 1]$, $(x, B) \mapsto \mathbf{P}_x[X \in B]$ is a stochastic kernel.
- (iii) X has the time-homogeneous Markov property (MP): For every $A \in \mathcal{B}(E)$, every $x \in E$ and all $s, t \in I$, we have

$$\mathbf{P}_x[X_{t+s} \in A \mid \mathcal{F}_s] = \kappa_t(X_s, A) \quad \mathbf{P}_x\text{-a.s.}$$

Here, for every $t \in I$, the transition kernel $\kappa_t : E \times \mathcal{B}(E) \rightarrow [0, 1]$ is the stochastic kernel defined for $x \in E$ and $A \in \mathcal{B}(E)$ by

$$\kappa_t(x, A) := \kappa(x, \{y \in E^I : y(t) \in A\}) = \mathbf{P}_x[X_t \in A]$$

The family $(\kappa_t(x, A), t \in I, x \in E, A \in \mathcal{B}(E))$ is also called the family of transition probabilities of X .

We write $\mathbf{E}_x$ for expectation with respect to $\mathbf{P}_x$, $\mathcal{L}_x[X] = \mathbf{P}_x$ and $\mathcal{L}_x[X | \mathcal{F}] = \mathbf{P}_x[X \in \cdot | \mathcal{F}]$ (for a regular conditional distribution of X given $\mathcal{F}$).

If E is countable, then X is called a discrete Markov process.

In the special case $I = \mathbb{N}_0$, X is called a Markov chain. In this case, κ_n is called the family of n -step transition probabilities.

Remark 17.4 We will see that the existence of the transition kernels (κ_t) implies the existence of the kernel κ . Thus, a time-homogeneous Markov process is simply a stochastic process with the Markov property and for which the transition probabilities are time-homogeneous. Although it is sometimes convenient to allow also timeinhomogeneous Markov processes, for a wide range of applications it is sufficient to consider time-homogeneous Markov processes. We will not go into the details but will henceforth assume that all Markov processes are time-homogeneous.

In the following, we will use the somewhat sloppy notation $\mathbf{P}_{X_s}[X \in \cdot] := \kappa(X_s, \cdot)$. That is, we understand X_s as the initial value of a second Markov process with the same distributions $(\mathbf{P}_x)_{x \in E}$.

Example 17.5 Let $Y_1, Y_2, \dots$ be i.i.d. $\mathbb{R}^d$ -valued random variables and let

$$S_n^x = x + \sum_{i=1}^n Y_i \quad \text{for } x \in \mathbb{R}^d \text{ and } n \in \mathbb{N}_0.$$

Define probability measures $\mathbf{P}_x$ on $((\mathbb{R}^d)^{\mathbb{N}_0}, (\mathcal{B}(\mathbb{R}^d))^{\otimes \mathbb{N}_0})$ by $\mathbf{P}_x = \mathbf{P} \circ (S^x)^{-1}$. Then the canonical process $X_n : (\mathbb{R}^d)^{\mathbb{N}_0} \rightarrow \mathbb{R}^d$ is a Markov chain with distributions $(\mathbf{P}_x)_{x \in \mathbb{R}^d}$. The process X is called a random walk on $\mathbb{R}^d$ with initial value x . $\square$

Example 17.6 In the previous example, it is simple to pass to continuous time; that is, $I = [0, \infty)$. To this end, let $(v_t)_{t \geq 0}$ be a convolution semigroup on $\mathbb{R}^d$ and let $\kappa_t(x, dy) = \delta_x * v_t(dy)$. In Theorem 14.47, for every $x \in \mathbb{R}^d$, we constructed a measure $\mathbf{P}_x$ on $((\mathbb{R}^d)^{[0, \infty)}, \mathcal{B}(\mathbb{R}^d)^{\otimes [0, \infty)})$ with

$$\mathbf{P}_x \circ (X_0, X_{t_1}, \dots, X_{t_n})^{-1} = \delta_x \otimes \bigotimes_{i=0}^{n-1} \kappa_{t_{i+1}-t_i}$$

for any choice of finitely many points $0 = t_0 < t_1 < \dots < t_n$. It is easy to check that the map $\kappa : \mathbb{R}^d \times \mathcal{B}(\mathbb{R}^d)^{\otimes [0, \infty)}, (x, A) \mapsto \mathbf{P}_x[A]$ is a stochastic kernel. The timehomogeneous Markov property is immediate from the fact that the increments are independent and stationary. $\square$

Example 17.7 (See Example 9.5 and Theorem 5.36.) Let $\theta > 0$ and $v_t^\theta(\{k\}) = e^{-\theta t} \frac{t^k \theta^k}{k!}$, $k \in \mathbb{N}_0$, the convolution semigroup of the Poisson distribution. The Markov process X on $\mathbb{N}_0$ with this semigroup is called a Poisson process with (jump) rate θ . $\square$

As in Example 17.6, we will construct a Markov process for a more general Markov semigroup of stochastic kernels.

Theorem 17.8 Let $I \subset [0, \infty)$ be closed under addition and let $(\kappa_t)_{t \in I}$ be a Markov semigroup of stochastic kernels from E to E . Then there is a measurable space $(\Omega, \mathcal{A})$ and a Markov process $((X_t)_{t \in I}, (\mathbf{P}_x)_{x \in E})$ on $(\Omega, \mathcal{A})$ with transition probabilities

$$\mathbf{P}_x[X_t \in A] = \kappa_t(x, A) \quad \text{for all } x \in E, A \in \mathcal{B}(E), t \in I. \quad (17.2)$$

Conversely, for every Markov process X , Eq. (17.2) defines a semigroup of stochastic kernels. By (17.2), the finite-dimensional distributions of X are uniquely determined.

Proof " " We construct X as a canonical process. Let $\Omega = E^{[0, \infty)}$ and $\mathcal{A} = \mathcal{B}(E)^{\otimes [0, \infty)}$. Further, let X_t be the projection on the t th coordinate. For $x \in E$,

define (see Corollary 14.43) on $(\Omega, \mathcal{A})$ the probability measure $\mathbf{P}_x$ such that, for finitely many time points $0 = t_0 < t_1 < \dots < t_n$, we have

$$\mathbf{P}_x \circ (X_{t_0}, \dots, X_{t_n})^{-1} = \delta_x \otimes \bigotimes_{i=0}^{n-1} \kappa_{t_{i+1}-t_i}.$$

Then

$$\begin{aligned} \mathbf{P}_x [X_{t_0} \in A_0, \dots, X_{t_n} \in A_n] \\ = \int_{A_{n-1}} \mathbf{P}_x [X_{t_0} \in A_0, \dots, X_{t_{n-2}} \in A_{n-2}, X_{t_{n-1}} \in dx_{n-1}] \\ \times \kappa_{t_n-t_{n-1}}(x_{n-1}, A_n) \end{aligned}$$

hence $\mathbf{P}_x [X_{t_n} \in A_n \mid \mathcal{F}_{t_{n-1}}] = \kappa_{t_n-t_{n-1}}(X_{t_{n-1}}, A_n)$. Thus X is recognized as a Markov process. Furthermore, we have $\mathbf{P}_x [X_t \in A] = (\delta_x \cdot \kappa_t)(A) = \kappa_t(x, A)$.

" $\Leftarrow$ " Now let $(X, (\mathbf{P}_x)_{x \in E})$ be a Markov process. Then a stochastic kernel κ_t is defined by

$$\kappa_t(x, A) := \mathbf{P}_x [X_t \in A] \quad \text{for all } x \in E, A \in \mathcal{B}(E), t \in I.$$

By the Markov property, we have

$$\begin{aligned} \kappa_{t+s}(x, A) &= \mathbf{P}_x [X_{t+s} \in A] = \mathbf{E}_x [\mathbf{P}_{X_s} [X_t \in A]] \\ &= \int \mathbf{P}_x [X_s \in dy] \mathbf{P}_y [X_t \in A] \\ &= \int \kappa_s(x, dy) \kappa_t(y, A) = (\kappa_s \cdot \kappa_t)(x, A) \end{aligned}$$

Hence $(\kappa_t)_{t \in I}$ is a Markov semigroup.

Theorem 17.9 A stochastic process $X = (X_t)_{t \in I}$ is a Markov process if and only if there exists a stochastic kernel $\kappa : E \times \mathcal{B}(E)^{\otimes I} \rightarrow [0, 1]$ such that, for every bounded $\mathcal{B}(E)^{\otimes I} - \mathcal{B}(\mathbb{R})$ -measurable function $f : E^I \rightarrow \mathbb{R}$ and for every $s \geq 0$ and $x \in E$, we have

$$\mathbf{E}_x [f((X_{t+s})_{t \in I}) \mid \mathcal{F}_s] = \mathbf{E}_{X_s} [f(X)] := \int_{E^I} \kappa(X_s, dy) f(y) \quad (17.3)$$

Proof " $\Rightarrow$ " The time-homogeneous Markov property follows by (17.3) with the function $f(y) = \mathbb{1}_A(y(t))$ since $\mathbf{P}_{X_s} [X_t \in A] = \mathbf{P}_x [X_{t+s} \in A \mid \mathcal{F}_s] = \kappa_t(X_s, A)$.

" $\implies$ " By the usual approximation arguments, it is enough to consider functions f that depend only on finitely many coordinates $0 \leq t_1 \leq t_2 \leq \dots \leq t_n$. We perform the proof by induction on n .

For $n = 1$ and f an indicator function, this is the (time-homogeneous) Markov property. For general measurable f , the statement follows by the usual approximation arguments.

Now assume the claim is proved for $n \in \mathbb{N}$. Again it suffices to assume that f is an indicator function of the type $f(x) = \mathbb{1}_{B_1 \times \dots \times B_{n+1}} l(x_{t_1}, \dots, x_{t_{n+1}})$ (with $B_1, \dots, B_{n+1} \in \mathcal{B}(E)$). Using the Markov property (third and fifth equalities in the following equation) and the induction hypothesis (fourth equality), we get

$$\begin{aligned}
& \mathbf{E}_x [f((X_{t+s})_{t \geq 0}) \mid \mathcal{F}_s] \\
&= \mathbf{E}_x [\mathbf{E}_x [f((X_{t+s})_{t \geq 0}) \mid \mathcal{F}_{t_n+s}] \mid \mathcal{F}_s] \\
&= \mathbf{E}_x \left[\mathbf{E}_x \left[\mathbb{1}_{\{X_{t_{n+1}+s} \in B_{n+1}\}} \mid \mathcal{F}_{t_n+s} \right] \mathbb{1}_{B_1}(X_{t_1+s}) \dots \mathbb{1}_{B_n}(X_{t_n+s}) \mid \mathcal{F}_s \right] \\
&= \mathbf{E}_x [\mathbf{P}_{X_{t_n+s}} [X_{t_{n+1}-t_n} \in B_{n+1}] \mathbb{1}_{B_1}(X_{t_1+s}) \dots \mathbb{1}_{B_n}(X_{t_n+s}) \mid \mathcal{F}_s] \\
&= \mathbf{E}_{X_s} [\mathbf{P}_{X_{t_n}} [X_{t_{n+1}-t_n} \in B_{n+1}] \mathbb{1}_{B_1}(X_{t_1}) \dots \mathbb{1}_{B_n}(X_{t_n})] \\
&= \mathbf{E}_{X_s} [\mathbf{P}_{X_0} [X_{t_{n+1}} \in B_{n+1} \mid \mathcal{F}_{t_n}] \mathbb{1}_{B_1}(X_{t_1}) \dots \mathbb{1}_{B_n}(X_{t_n})] \\
&= \mathbf{E}_{X_s} [\mathbf{P}_{X_0} [X_{t_1} \in B_1, \dots, X_{t_{n+1}} \in B_{n+1} \mid \mathcal{F}_{t_n}]] \\
&= \mathbf{E}_{X_s} [f(X)].
\end{aligned}$$

□

Corollary 17.10 A stochastic process $(X_n)_{n \in \mathbb{N}_0}$ is a Markov chain if and only if

$$\mathcal{L}_x [(X_{n+k})_{n \in \mathbb{N}_0} \mid \mathcal{F}_k] = \mathcal{L}_{X_k} [(X_n)_{n \in \mathbb{N}_0}] \quad \text{for every } k \in \mathbb{N}_0. \quad (17.4)$$

Proof If the conditional distributions exist, then, by Theorem 17.9, the equation (17.4) is equivalent to X being a Markov chain. Hence we only have to show that the conditional distributions exist.

Since E is Polish, $E^{\mathbb{N}_0}$ is also Polish and we have $\mathcal{B}(E^{\mathbb{N}_0}) = \mathcal{B}(E)^{\otimes \mathbb{N}_0}$ (see Theorem 14.8). Hence, by Theorem 8.37, there exists a regular conditional distribution of $(X_{n+k})_{n \in \mathbb{N}_0}$ given $\mathcal{F}_k$.

□

Theorem 17.11 Let $I = \mathbb{N}_0$. If $(X_n)_{n \in \mathbb{N}_0}$ is a stochastic process with distributions $(\mathbf{P}_x, x \in E)$, then the Markov property in Definition 17.3(iii) is implied by the existence of a stochastic kernel $\kappa_1 : E \times \mathcal{B}(E) \rightarrow [0, 1]$ with the property that for every $A \in \mathcal{B}(E)$, every $x \in E$ and every $s \in I$, we have

$$\mathbf{P}_x [X_{s+1} \in A \mid \mathcal{F}_s] = \kappa_1(X_s, A). \quad (17.5)$$

In this case, the n -step transition kernels κ_n can be computed inductively by

$$\kappa_n = \kappa_{n-1} \cdot \kappa_1 = \int_E \kappa_{n-1}(\cdot, dx) \kappa_1(x, \cdot)$$

In particular, the family $(\kappa_n)_{n \in \mathbb{N}}$ is a Markov semigroup and the distribution X is uniquely determined by κ_1 .

Proof In Theorem 17.9, let $t_i = i$ for every $i \in \mathbb{N}_0$. For the proof of that theorem, only (17.5) was needed. □

The (time-homogeneous) Markov property of a process means that, for fixed time t , the future (after t) depends on the past (before t) only via the present (that is, via the value X_t). We can generalize this concept by allowing random times τ instead of fixed times t .

Definition 17.12 Let $I \subset [0, \infty)$ be closed under addition. A Markov process $(X_t)_{t \in I}$ with distributions $(\mathbf{P}_x, x \in E)$ has the strong Markov property if, for every a.s. finite stopping time τ , every bounded $\mathcal{B}(E)^{\otimes I} - \mathcal{B}(\mathbb{R})$ measurable function $f : E^I \rightarrow \mathbb{R}$ and every $x \in E$, we have

$$\mathbf{E}_x [f((X_{\tau+t})_{t \in I}) \mid \mathcal{F}_\tau] = \mathbf{E}_{X_\tau} [f(X)] := \int_{E^I} \kappa(X_\tau, dy) f(y) \quad (17.6)$$

Remark 17.13 If I is countable, then the strong Markov property holds if and only if, for every almost surely finite stopping time τ , we have

$$\mathcal{L}_x [(X_{\tau+t})_{t \in \mathbb{N}_0} \mid \mathcal{F}_\tau] = \mathcal{L}_{X_\tau} [(X_t)_{t \in \mathbb{N}_0}] := \kappa(X_\tau, \cdot) \quad (17.7)$$

This follows just as in Corollary 17.10.

Most Markov processes one encounters have the strong Markov property. In particular, for countable time sets, the strong Markov property follows from the Markov property. For continuous time, however, in general, some work has to be done to establish the strong Markov property.

Theorem 17.14 If $I \subset [0, \infty)$ is countable and closed under addition, then every Markov process $(X_n)_{n \in I}$ with distributions $(\mathbf{P}_x)_{x \in E}$ has the strong Markov property.

Proof Let $f : E^I \rightarrow \mathbb{R}$ be measurable and bounded. Then, for every $s \in I$, the random variable $\mathbb{1}_{\{\tau=s\}} \mathbf{E}_x [f((X_{s+t})_{t \in I}) \mid \mathcal{F}_\tau]$ is measurable with respect to $\mathcal{F}_s$. Using the tower property of the conditional expectation and Theorem 17.9 in the third equality, we thus get

$$\begin{aligned} \mathbf{E}_x [f((X_{\tau+t})_{t \in I}) \mid \mathcal{F}_\tau] &= \sum_{s \in I} \mathbb{1}_{\{\tau=s\}} \mathbf{E}_x [f((X_{s+t})_{t \in I}) \mid \mathcal{F}_\tau] \\ &= \sum_{s \in I} \mathbf{E}_x [\mathbb{1}_{\{\tau=s\}} \mathbf{E}_x [f((X_{s+t})_{t \in I}) \mid \mathcal{F}_s] \mid \mathcal{F}_\tau] \\ &= \sum_{s \in I} \mathbf{E}_x [\mathbb{1}_{\{\tau=s\}} \mathbf{E}_{X_s} [f((X_t)_{t \in I})] \mid \mathcal{F}_\tau] \\ &= \mathbf{E}_{X_\tau} [f((X_t)_{t \in I})] \end{aligned}$$

As a simple application of the strong Markov property, we show the reflection principle for random walks.

Theorem 17.15 (Reflection principle) Let $Y_1, Y_2, \dots$ be i.i.d. real random variables with symmetric distribution $\mathcal{L}[Y_1] = \mathcal{L}[-Y_1]$. Define $X_0 = 0$ and $X_n := Y_1 + \dots + Y_n$ for $n \in \mathbb{N}$. Then, for every $n \in \mathbb{N}_0$ and $a > 0$,

$$\mathbf{P} \left[\sup_{m \leq n} X_m \geq a \right] \leq 2\mathbf{P}[X_n \geq a] - \mathbf{P}[X_n = a] \quad (17.8)$$

If we have $\mathbf{P}[Y_1 \in \{-1, 0, 1\}] = 1$, then for $a \in \mathbb{N}$ equality holds in (17.8).

Proof Let $a > 0$ and $n \in \mathbb{N}$. Define the time of first excess of a (truncated at $(n+1)$),

$$\tau := \inf \{m \geq 0 : X_m \geq a\} \wedge (n+1)$$

Then τ is a bounded stopping time and

$$\sup_{m \leq n} X_m \geq a \iff \tau \leq n.$$

Let $f(m, X) = \mathbb{1}_{\{m \leq n\}} (\mathbb{1}_{\{X_{n-m} > a\}} + \frac{1}{2} \mathbb{1}_{\{X_{n-m} = a\}})$. Then

$$f(\tau, (X_{\tau+m})_{m \in \mathbb{N}_0}) = \mathbb{1}_{\{\tau \leq n\}} \left(\mathbb{1}_{\{X_n > a\}} + \frac{1}{2} \mathbb{1}_{\{X_n = a\}} \right)$$

The strong Markov property of X yields

$$\mathbf{E}_0 [f(\tau, (X_{\tau+m})_{m \geq 0}) \mid \mathcal{F}_\tau] = \varphi(\tau, X_\tau)$$

where $\varphi(m, x) = \mathbf{E}_x[f(m, X)]$. (Recall that $\mathbf{E}_x$ denotes the expectation for X if $X_0 = x$.)

Due to the symmetry of Y_i , we have

$$\varphi(m, x) \begin{cases} \geq \frac{1}{2}, & \text{if } m \leq n \text{ and } x \geq a, \\ = \frac{1}{2}, & \text{if } m \leq n \text{ and } x = a, \\ = 0, & m > n. \end{cases}$$

Hence

$$\begin{aligned}\{\tau \leq n\} &= \{\tau \leq n\} \cap \{X_\tau \geq a\} \subset \left\{ \varphi(\tau, X_\tau) \geq \frac{1}{2} \right\} \cap \{\tau \leq n\} \\ &= \{\varphi(\tau, X_\tau) > 0\} \cap \{\tau \leq n\}\end{aligned}$$

Now (17.8) is implied by

$$\begin{aligned}\mathbf{P}[X_n > a] + \frac{1}{2}\mathbf{P}[X_n = a] &= \mathbf{E}[f(\tau, (X_{\tau+m})_{m \geq 0})] \\ &= \mathbf{E}_0[\varphi(\tau, X_\tau) \mathbb{1}_{\{\tau \leq n\}}] \geq \frac{1}{2}\mathbf{P}_0[\tau \leq n]\end{aligned}\quad (17.9)$$

Now assume $\mathbf{P}[Y_1 \in \{-1, 0, 1\}] = 1$ and $a \in \mathbb{N}$. Then $X_\tau = a$ if $\tau \leq n$. Hence

$$\{\varphi(\tau, X_\tau) > 0\} \cap \{\tau \leq n\} = \left\{ \varphi(\tau, X_\tau) = \frac{1}{2} \right\} \cap \{\tau \leq n\}$$

Thus, in the last step of (17.9), equality holds and hence also in (17.8).

Exercise 17.1.1 Let $I \subset \mathbb{R}$ and let $X = (X_t)_{t \in I}$ be a stochastic process. For $t \in I$, define the σ -algebras that code the past before t and the future beginning with t by

$$\mathcal{F}_{\leq t} := \sigma(X_s : s \in I, s \leq t) \quad \text{and} \quad \mathcal{F}_{\geq t} := \sigma(X_s : s \in I, s \geq t).$$

Show that X has the Markov property if and only if, for every $t \in I$, the σ -algebras $\mathcal{F}_{\leq t}$ and $\mathcal{F}_{\geq t}$ are independent given $\sigma(X_t)$ (compare Definition 12.20).

In other words, a process has the (possibly time-inhomogeneous) Markov property if and only if past and future are independent given the present.

5.17.2 17.2 Discrete Markov Chains: Examples

Let E be countable and $I = \mathbb{N}_0$. By Definition 17.3, a Markov process $X = (X_n)_{n \in \mathbb{N}_0}$ on E is a discrete Markov chain (or Markov chain with discrete state space).

If X is a discrete Markov chain, then $(\mathbf{P}_x)_{x \in E}$ is determined by the transition matrix

$$p = (p(x, y))_{x, y \in E} := (\mathbf{P}_x[X_1 = y])_{x, y \in E}.$$

The n -step transition probabilities

$$p^{(n)}(x, y) := \mathbf{P}_x[X_n = y]$$

can be computed as the n -fold matrix product

$$p^{(n)}(x, y) = p^n(x, y)$$

where

$$p^n(x, y) = \sum_{z \in E} p^{n-1}(x, z)p(z, y)$$

and where $p^0 = I$ is the unit matrix.

By induction, we get the Chapman-Kolmogorov equation (see (14.14)) for all $m, n \in \mathbb{N}_0$ and $x, y \in E$,

$$p^{(m+n)}(x, y) = \sum_{z \in E} p^{(m)}(x, z)p^{(n)}(z, y) \quad (17.10)$$

Definition 17.16 A matrix $(p(x, y))_{x, y \in E}$ with nonnegative entries and with

$$\sum_{y \in E} p(x, y) = 1 \quad \text{for all } x \in E$$

is called a stochastic matrix on E .

A stochastic matrix is essentially a stochastic kernel from E to E . In Theorem 17.8 we saw that, for the semigroup of kernels $(p^n)_{n \in \mathbb{N}}$, there exists a unique discrete Markov chain whose transition probabilities are given by p . The arguments we gave there were rather abstract. Here we give a construction for X that could actually be used to implement a computer simulation of X .

Let $(R_n)_{n \in \mathbb{N}_0}$ be an independent family of random variables with values in E^E and with the property

$$\mathbf{P}[R_n(x) = y] = p(x, y) \quad \text{for all } x, y \in E. \quad (17.11)$$

For example, choose $(R_n(x), x \in E, n \in \mathbb{N})$ as an independent family of random variables with values in E and distributions

$$\mathbf{P}[R_n(x) = y] = p(x, y) \quad \text{for all } x, y \in E \text{ and } n \in \mathbb{N}_0.$$

Note, however, that in (17.11) we have required neither independence of the random variables $(R_n(x), x \in E)$ nor that all R_n had the same distribution. Only the onedimensional marginal distributions are determined. In fact, in many applications it is useful to have subtle dependence structures in order to couple Markov chains with different initial chains. We pick up this thread again in Section 18.2.

For $x \in E$, define

$$X_0^x = x \quad \text{and} \quad X_n^x = R_n(X_{n-1}^x) \quad \text{for } n \in \mathbb{N}.$$

Finally, let $\mathbf{P}_x := \mathcal{L}[X^x]$ be the distribution of X^x . Recall that this is a probability measure on the space of sequences $(E^{\mathbb{N}_0}, \mathcal{B}(E)^{\otimes \mathbb{N}_0})$.

Theorem 17.17

- (i) With respect to the distribution $(\mathbf{P}_x)_{x \in E}$, the canonical process X on $(E^{\mathbb{N}_0}, \mathcal{B}(E)^{\otimes \mathbb{N}_0})$ is a Markov chain with transition matrix p .
- (ii) In particular, to any stochastic matrix p , there corresponds a unique discrete Markov chain X with transition probabilities p .

Proof "(ii)" This follows from (i) since Theorem 17.11 yields uniqueness of X .

"(i)" For $n \in \mathbb{N}_0$ and $x, y, z \in E$, by construction,

$$\begin{aligned} \mathbf{P}_x[X_{n+1} = z \mid \mathcal{F}_n, X_n = y] &= \mathbf{P}[X_{n+1}^x = z \mid \sigma(R_m, m \leq n), X_n^x = y] \\ &= \mathbf{P}[R_{n+1}(X_n^x) = z \mid \sigma(R_m, m \leq n), X_n^x = y] \\ &= \mathbf{P}[R_{n+1}(y) = z] \\ &= p(y, z). \end{aligned}$$

Hence, by Theorem 17.11, X is a Markov chain with transition matrix p .

Example 17.18 (Random walk on $\mathbb{Z}$) Let $E = \mathbb{Z}$, and assume

$$p(x, y) = p(0, y - x) \quad \text{for all } x, y \in \mathbb{Z}.$$

In this case, we say that p is translation invariant. A discrete Markov chain X with transition matrix p is a random walk on $\mathbb{Z}$. Indeed, $X_n \stackrel{D}{=} X_0 + Z_1 + \dots + Z_n$, where $(Z_n)_{n \in \mathbb{N}}$ are i.i.d. with $\mathbf{P}[Z_n = x] = p(0, x)$.

The R_n that we introduced in the explicit construction are given by $R_n(x) := x + Z_n$.

Example 17.19 (Computer simulation) Consider the situation where the state space $E = \{1, \dots, k\}$ is finite. The aim is to simulate a Markov chain X with transition matrix p on a computer. Assume that the computer provides a random number generator that generates an i.i.d. sequence $(U_n)_{n \in \mathbb{N}}$ of random variables that are uniformly distributed on $[0, 1]$. (Of course, this is wishful thinking. But modern random number generators produce sequences that for many purposes are close enough to really random sequences.)

Define $r(i, 0) = 0, r(i, j) = p(i, 1) + \dots + p(i, j)$ for $i, j \in E$, and define Y_n by

$$R_n(i) = j \iff U_n \in [r(i, j-1), r(i, j)).$$

Then, by construction, $\mathbf{P}[R_n(i) = j] = r(i, j) - r(i, j-1) = p(i, j)$.

Example 17.20 (Branching process as a Markov chain) We want to understand the Galton-Watson branching process (see Definition 3.9) as a Markov chain on $E = \mathbb{N}_0$.

To this end, let $(q_k)_{k \in \mathbb{N}_0}$ be a probability vector, the offspring distribution of one individual. Define $q_k^{*0} = \mathbb{1}_{\{0\}}(k)$ and

$$q_k^{*n} = \sum_{l=0}^k q_{k-l}^{*(n-1)} q_l \quad \text{for } n \in \mathbb{N}$$

as the n -fold convolutions of q . Hence, for n individuals, q_k^{*n} is the probability to have exactly k offspring. Finally, define the matrix p by $p(x, y) = q_y^{*x}$ for $x, y \in \mathbb{N}_0$.

Now let $(Y_{n,i}, n \in \mathbb{N}_0, i \in \mathbb{N}_0)$ be i.i.d. with $\mathbf{P}[Y_{n,i} = k] = q_k$. For $x \in \mathbb{N}_0$, define the branching process X with x ancestors and offspring distribution q by $X_0 = x$ and $X_n := \sum_{i=1}^{X_{n-1}} Y_{n-1,i}$. In order to show that X is a Markov chain, we compute

$$\begin{aligned} \mathbf{P}[X_n = x_n \mid X_0 = x, X_1 = x_1, \dots, X_{n-1} = x_{n-1}] \\ &= \mathbf{P}[Y_{n-1,1} + \dots + Y_{n-1,x_{n-1}} = x_n] \\ &= \mathbf{P}_{Y_{1,1}}^{*x_{n-1}}(\{x_n\}) = q_{x_n}^{*x_{n-1}} = p(x_{n-1}, x_n) \end{aligned}$$

Hence X is a Markov chain on $\mathbb{N}_0$ with transition matrix p .

Example 17.21 (Wright's evolution model) In population genetics, Wright's evolution model [171] describes the hereditary transmission of a genetic trait with two possible specifications (say A and B); for example, resistance/no resistance to a specific antibiotic. It is assumed that the population has a constant size of $N \in \mathbb{N}$ individuals and the generations change at discrete times and do not overlap. Furthermore, for simplicity, the individuals are assumed to be haploid; that is, cells bear only one copy of each chromosome (like certain protozoans do) and not two copies (as in mammals).

Here we consider the case where none of the traits is favored by selection. Hence, it is assumed that each individual of the new generation chooses independently and uniformly at random one individual of the preceding generation as ancestor and becomes a perfect clone of that. Thus, if the number of individuals of type A in the current generation is $k \in \{0, \dots, N\}$, then in the new generation it will be random and binomially distributed with parameters N and k/N .

The gene frequencies k/N in this model can be described by a Markov chain X on $E = \{0, 1/N, \dots, (N-1)/N, 1\}$ with transition matrix $p(x, y) = b_{N,x}(\{Ny\})$. Note that X is a (bounded) martingale. Hence, by the martingale convergence theorem (Theorem 11.7), X converges $\mathbf{P}_x$ -almost surely to a random variable X_∞ with $\mathbf{E}_x[X_\infty] = \mathbf{E}_x[X_0] = x$. As with

the voter model (see Example 11.16) that is closely related to Wright's model, we can argue that the limit X_∞ can take only the stable values 0 and 1. That is, $\mathbf{P}_x[\lim_{n \rightarrow \infty} X_n = 1] = x = 1 - \mathbf{P}_x[\lim_{n \rightarrow \infty} X_n = 0]$.

Example 17.22 (Discrete Moran model) In contrast to Wright's model, the Moran model also allows overlapping generations. The situation is similar to that of Wright's model; however, now in each time step, only (exactly) one individual gets replaced by a new one, whose type is chosen at random from the whole population.

As the new and the old types of the replaced individual are independent, as a model for the gene frequencies, we obtain a Markov chain X on $E = \{0, \frac{1}{N}, \dots, 1\}$ with transition matrix

$$p(x, y) = \begin{cases} x(1-x), & \text{if } y = x + 1/N \\ x^2 + (1-x)^2, & \text{if } y = x \\ x(1-x), & \text{if } y = x - 1/N \\ 0, & \text{else.} \end{cases}$$

Here also, X is a bounded martingale and we can compute the square variation process,

$$\langle X \rangle_n = \sum_{i=1}^n \mathbf{E}[(X_i - X_{i-1})^2 \mid X_{i-1}] = \frac{2}{N^2} \sum_{i=0}^{n-1} X_i(1 - X_i). \quad (17.12)$$

Exercise 17.2.1 (Discrete martingale problem) Let $E \subset \mathbb{R}$ be countable and let X be a Markov chain on E with transition matrix p and with the property that, for any x , there are at most three choices for the next step; that is, there exists a set $A_x \subset E$ of cardinality 3 with $p(x, y) = 0$ for all $y \in E \setminus A_x$. Let $d(x) := \sum_{y \in E} (y - x)p(x, y)$ for $x \in E$.

(i) Show that $M_n := X_n - \sum_{k=0}^{n-1} d(X_k)$ defines a martingale M with square variation process $\langle M \rangle_n = \sum_{i=0}^{n-1} f(X_i)$ for a unique function $f : E \rightarrow [0, \infty)$.

(ii) Show that the transition matrix p is uniquely determined by f and d .

(iii) For the Moran model (Example 17.22), use the explicit form (17.12) of the square variation process to compute the transition matrix.

5.17.3 17.3 Discrete Markov Processes in Continuous Time

Let E be countable and let $(X_t)_{t \in [0, \infty)}$ be a Markov process on E with transition probabilities $p_t(x, y) = \mathbf{P}_x[X_t = y]$ (for $x, y \in E$). (Some authors call such a process a Markov chain in continuous time.)

Let $x, y \in E$ with $x \neq y$. We say that X jumps with rate $q(x, y)$ from x to y if the following limit exists:

$$q(x, y) := \lim_{t \downarrow 0} \frac{1}{t} \mathbf{P}_x[X_t = y]$$

Henceforth we assume that the limit $q(x, y)$ exists for all $y \neq x$ and that

$$\sum_{y \neq x} q(x, y) < \infty \quad \text{for all } x \in E \quad (17.13)$$

Then we define

$$q(x, x) = - \sum_{y \neq x} q(x, y) \quad (17.14)$$

Finally we assume that (which is equivalent to exchangeability of the limit and the sum over $y \neq x$ in the display preceding (17.13))

$$\lim_{t \downarrow 0} \frac{1}{t} (\mathbf{P}_x [X_t = y] - \mathbb{1}_{\{x=y\}}) = q(x, y) \quad \text{for all } x, y \in E. \quad (17.15)$$

Definition 17.23 If (17.13), (17.14) and (17.15) hold, then q is called the Q -matrix of X . Sometimes q is also called the generator of the semigroup $(p_t)_{t \geq 0}$.

Example 17.24 (Poisson process) The Poisson process with rate $\alpha > 0$ (compare Section 5.5) has the Q -matrix $q(x, y) = \alpha (\mathbb{1}_{\{y=x+1\}} - \mathbb{1}_{\{y=x\}})$.

Theorem 17.25 Let q be an $E \times E$ matrix such that $q(x, y) \geq 0$ for all $x, y \in E$ with $x \neq y$. Assume that (17.13) and (17.14) hold and that

$$\lambda := \sup_{x \in E} |q(x, x)| < \infty. \quad (17.16)$$

Then q is the Q -matrix of a unique Markov process X .

Intuitively, (17.15) suggests that we define $p_t = e^{tq}$ in a suitable sense. Then, formally, $q = \frac{d}{dt} p_t \big|_{t=0}$. The following proof shows that this formal argument can be made rigorous.

Proof Let I be the unit matrix on E . Define

$$p(x, y) = \frac{1}{\lambda} q(x, y) + I(x, y) \quad \text{for } x, y \in E$$

if $\lambda > 0$ and $p = I$ otherwise. Then p is a stochastic matrix and $q = \lambda(p - I)$. Let $((Y_n)_{n \in \mathbb{N}_0}, (\mathbf{P}_x^Y)_{x \in E})$ be a discrete Markov chain with transition matrix p and let $((T_t)_{t \geq 0}, (\mathbf{P}_n^T)_{n \in \mathbb{N}_0})$ be a Poisson process with rate λ . Let $X_t := Y_{T_t}$ and $\mathbf{P}_x = \mathbf{P}_x^Y \otimes \mathbf{P}_0^T$. Then $\mathfrak{X} := ((X_t)_{t \geq 0}, (\mathbf{P}_x)_{x \in E})$ is a Markov process and

$$\begin{aligned} p_t(x, y) &:= \mathbf{P}_x [X_t = y] = \sum_{n=0}^{\infty} \mathbf{P}_0^T [T_t = n] \mathbf{P}_x^Y [Y_n = y] \\ &= e^{-\lambda t} \sum_{n=0}^{\infty} \frac{\lambda^n t^n}{n!} p^n(x, y) \end{aligned}$$

This power series (in t) converges everywhere (note that as a linear operator, p has finite norm $\|p\|_2 \leq 1$) to the matrix exponential function $e^{\lambda t p}(x, y)$. Furthermore,

$$p_t(x, y) = e^{-\lambda t} e^{\lambda t p}(x, y) = e^{\lambda t(p-I)}(x, y) = e^{tq}(x, y).$$

Differentiating the power series termwise yields $\frac{d}{dt} p_t(x, y) \big|_{t=0} = q(x, y)$. Hence $\mathfrak{X}$ is the required Markov process.

Now assume that $(\tilde{p}_t)_{t \geq 0}$ are the transition probabilities of another Markov process $\tilde{\mathfrak{X}}$ with the same generator q ; that is, with

$$\lim_{s \downarrow 0} \frac{1}{s} (\tilde{p}_s(x, y) - I(x, y)) = q(x, y).$$

It is easy to check that

$$\lim_{s \downarrow 0} \frac{1}{s} (p_{t+s}(x, y) - p_t(x, y)) = (q \cdot p_t)(x, y).$$

That is, we have $(d/dt)p_t(x, y) = qp_t(x, y)$. Similarly, we get $(d/dt)\tilde{p}_t = q\tilde{p}_t(x, y)$. Hence also,

$$p_t(x, y) - \tilde{p}_t(x, y) = \int_0^t (q(p_s - \tilde{p}_s))(x, y) ds$$

If we let $r_s = p_s - \tilde{p}_s$, then $\|r_s\|_2 \leq 2$ and $\|q\|_2 \leq 2\lambda$; hence

$$\sup_{s \leq t} \|r_s\|_2 \leq \sup_{s \leq t} \int_0^s \|qr_u\|_2 du \leq \|q\|_2 \sup_{s \leq t} \int_0^s \|r_u\|_2 du \leq 2\lambda t \sup_{s \leq t} \|r_s\|_2$$

For $t < 1/2\lambda$, this implies $r_t = 0$; hence $\tilde{p}_t = p_t$. For general $t > 0$, choose $n \in \mathbb{N}$ such that $t/n < 1/2\lambda$ to obtain $\tilde{p}_t = (\tilde{p}_{t/n})^n = (p_{t/n})^n = p_t$.

Remark 17.26 The condition (17.16) cannot be dropped easily, as the following example shows. Let $E = \mathbb{N}$ and

$$q(x, y) = \begin{cases} x^2, & \text{if } y = x + 1 \\ -x^2, & \text{if } y = x \\ 0, & \text{else.} \end{cases}$$

We construct explicitly a candidate X for a Markov process with Q -matrix q . Let $T_1, T_2, \dots$ be independent, exponentially distributed random variables with $\mathbf{P}_{T_n} = \exp_{n^2}$. Define $S_n = T_1 + \dots + T_{n-1}$ and $X_t = \sup \{n \in \mathbb{N}_0 : S_n \leq t\}$. Then, at any time, X makes at most one step to the right. Furthermore, due to the lack of memory of the exponential distribution (see Exercise 8.1.1),

$$\begin{aligned} \mathbf{P}[X_{t+s} \geq n+1 \mid X_t = n] &= \mathbf{P}[S_{n+1} \leq t+s \mid S_n \leq t, S_{n+1} > t] \\ &= \mathbf{P}[T_n \leq s+t-S_n \mid S_n \leq t, T_n > t-S_n] = \mathbf{P}[T_n \leq s] \\ &= 1 - \exp(-n^2 s) \end{aligned}$$

Therefore,

$$\lim_{s \downarrow 0} s^{-1} \mathbf{P}[X_{t+s} = n+1 \mid X_t = n] = n^2$$

and

$$\lim_{s \downarrow 0} s^{-1} (\mathbf{P}[X_{t+s} = n \mid X_t = n] - 1) = -n^2;$$

hence

$$\lim_{s \downarrow 0} s^{-1} (\mathbf{P}[X_{t+s} = m \mid X_t = n] - I(m, n)) = q(m, n) \quad \text{for all } m, n \in \mathbb{N}.$$

Let

$$\tau^n = \inf \{t \geq 0 : X_t = n\} = S_n \quad \text{for } n \in \mathbb{N}$$

Then $\mathbf{E}_1[\tau^n] = \sum_{k=1}^{n-1} \frac{1}{k^2}$. By monotone convergence, $\mathbf{E}_1[\sup_{n \in \mathbb{N}} \tau^n] < \infty$. That is, in finite time, X exceeds all levels. We say that X explodes. $\square$

Example 17.27 (A variant of Pólya's urn model) Consider a variant of Pólya's urn model with black and red balls (compare Example 12.29). In contrast to the original model, we do not simply add one ball of the same color as the ball that we return. Rather, the number of balls that we add varies from time to time. More precisely, the k th ball of a given color will be returned together with r_k more balls of the same color. The numbers $r_1, r_2, \dots \in \mathbb{N}$ are parameters of the model. In particular, the case $1 = r_1 = r_2 = \dots$ is the classical Pólya's urn model. Let

$$X_n := \begin{cases} 1, & \text{if the } n\text{th ball is black} \\ 0, & \text{else.} \end{cases}$$

For the classical model, we saw (Example 12.29) that the fraction of black balls in the urn converges a.s. to a Beta-distributed random variable Z . Furthermore, given Z , the sequence $X_1, X_2, \dots$ is independent and Ber_Z -distributed. A similar statement holds for the case where $r = r_1 = r_2 = \dots$ for some $r \in \mathbb{N}$. Indeed, here only the parameters of the Beta distribution change. In particular (as the Beta distribution is continuous and, in particular, does not have atoms at 0 or 1), almost surely we draw infinitely many balls of each color. Formally, $\mathbf{P}[B] = 0$ where B is the event where there is one color of which only finitely many balls are drawn.

The situation changes when the numbers r_k grow quickly as $k \rightarrow \infty$. Assume that in the beginning there is one black and one red ball in the urn. Denote by $w_n = 1 + \sum_{k=1}^n r_k$ the total number of balls of a given color after n balls of that color have been drawn already ($n \in \mathbb{N}_0$).

For illustration, first consider the extreme situation where w_n grows very quickly; for example, $w_n = 2^n$ for every $n \in \mathbb{N}$. Denote by

$$S_n = 2(X_1 + \dots + X_n) - n$$

the number of black balls drawn in the first n steps minus the number of red balls drawn in these steps. Then, for every $n \in \mathbb{N}_0$,

$$\mathbf{P}[X_{n+1} = 1 \mid S_n] = \frac{2^{S_n}}{1 + 2^{S_n}} \quad \text{and} \quad \mathbf{P}[X_{n+1} = 0 \mid S_n] = \frac{2^{-S_n}}{1 + 2^{-S_n}}.$$

We conclude that $(Z_n)_{n \in \mathbb{N}_0} := (|S_n|)_{n \in \mathbb{N}_0}$ is a Markov chain on $\mathbb{N}_0$ with transition matrix

$$p(z, z') = \begin{cases} 2^z / (1 + 2^z), & \text{if } z' = z + 1 > 1, \\ 1, & \text{if } z' = z + 1 = 1, \\ 1 / (1 + 2^z), & \text{if } z' = z - 1, \\ 0, & \text{else.} \end{cases}$$

The event B from above can be written as

$$B = \{Z_{n+1} < Z_n \text{ only finitely often}\}.$$

Let $A = \{Z_{n+1} > Z_n \text{ for all } n \in \mathbb{N}_0\}$ denote the event where Z flees directly to ∞ and let $\tau_z = \inf\{n \in \mathbb{N}_0 : Z_n \geq z\}$. Evidently,

$$\mathbf{P}_z[A] = \prod_{z'=z}^{\infty} p(z', z' + 1) \geq 1 - \sum_{z'=z}^{\infty} \frac{1}{1 + 2^{z'}} \geq 1 - 2^{1-z}$$

It is easy to check that $\mathbf{P}_0[\tau_z < \infty] = 1$ for all $z \in \mathbb{N}_0$. Using the strong Markov property, we get that, for all $z \in \mathbb{N}_0$,

$$\mathbf{P}_0[B] \geq \mathbf{P}_0[Z_{n+1} > Z_n \text{ for all } n \geq \tau_z] = \mathbf{P}_z[A] \geq 1 - 2^{1-z},$$

and thus $\mathbf{P}_0[B] = 1$. In prose, almost surely eventually only balls of one color will be drawn.

This example was a bit extreme. In order to find a necessary and sufficient condition on the growth of (w_n) , we need more subtle methods that appeal to the above example of the explosion of a Markov process.

We will show that $\mathbf{P}[B] = 1$ if and only if $\sum_{n=0}^{\infty} \frac{1}{w_n} < \infty$. To this end, consider independent random variables $T_1^s, T_1^r, T_2^s, T_2^r, \dots$ with $\mathbf{P}_{T_n^r} = \mathbf{P}_{T_n^s} = \exp_{w_{n-1}}$. Let $T_\infty^r = \sum_{n=1}^{\infty} T_n^r$ and $T_\infty^s = \sum_{n=1}^{\infty} T_n^s$. Clearly, $\mathbf{E}[T_\infty^r] = \sum_{n=0}^{\infty} 1/w_n < \infty$; hence, in particular, $\mathbf{P}[T_\infty^r < \infty] = 1$. The corresponding statement holds for T_∞^s . Note that T_∞^r and T_∞^s are independent and have densities (since T_1^r and T_1^s have densities); hence we have $\mathbf{P}[T_\infty^r = T_\infty^s] = 0$.

Now let

$$R_t := \sup \{n \in \mathbb{N} : T_1^r + \dots + T_{n-1}^r \leq t\}$$

and

$$S_t := \sup \{n \in \mathbb{N} : T_1^s + \dots + T_{n-1}^s \leq t\}$$

Let $R := \{T_1^r + \dots + T_n^r, n \in \mathbb{N}\}$ and let $S := \{T_1^s + \dots + T_n^s, n \in \mathbb{N}\}$ be the jump times of (R_t) and (S_t) . Define $U := R \cup S = \{u_1, u_2, \dots\}$, where $u_1 < u_2 < \dots$. Let

$$X_n = \begin{cases} 1, & \text{if } u_n \in S \\ 0, & \text{else} \end{cases}$$

Let $L_n = x_1 + \dots + x_n$. Then

$$\begin{aligned} \mathbf{P}[X_{n+1} = 1 \mid X_1 = x_1, \dots, X_n = x_n] \\ &= \mathbf{P}[u_{n+1} \in S \mid (u_k \in S \iff x_k = 1) \text{ for every } k \leq n] \\ &= \mathbf{P}[T_1^s + \dots + T_{L_n+1}^s < T_1^r + \dots + T_{n-L_n+1}^r \mid \\ &\quad T_1^s + \dots + T_{L_n+1}^s > T_1^r + \dots + T_{n-L_n}^r] \\ &= \mathbf{P}[T_{L_n+1}^s < T_{n-L_n+1}^r] = \frac{w_{L_n}}{w_{L_n} + w_{n-L_n}}. \end{aligned}$$

Hence $(X_n)_{n \in \mathbb{N}_0}$ is our generalized urn model with weights $(w_n)_{n \in \mathbb{N}_0}$. Consider now the event B^c where infinitely many balls of each color are drawn. Evidently, $\{X_n = 1 \text{ infinitely often}\} = \{\sup S = \sup U\}$ and $\{X_n = 0 \text{ infinitely often}\} = \{\sup R = \sup U\}$. Since $\sup S = T_\infty^s$ and $\sup R = T_\infty^r$, we thus have $\mathbf{P}[B^c] = \mathbf{P}[T_\infty^r = T_\infty^s] = 0$.

Exercise 17.3.1 Let $r, s, R, S \in \mathbb{N}$. Consider the generalized version of Pólya's urn model $(X_n)_{n \in \mathbb{N}_0}$ with $r_k = r$ and $s_k = s$ for all $k \in \mathbb{N}$. Assume that in the beginning there are R red balls and S black balls in the urn. Show that the fraction of black balls converges almost surely to a random variable Z with a Beta distribution and determine the parameters. Show that $(X_n)_{n \in \mathbb{N}_0}$ is i.i.d. given Z and $X_i \sim \text{Ber}_Z$ for all $i \in \mathbb{N}_0$.

Exercise 17.3.2 Show that, almost surely, infinitely many balls of each color are drawn if

$$\sum_{n=0}^{\infty} \frac{1}{w_n} = \infty$$

5.17.4 17.4 Discrete Markov Chains: Recurrence and Transience

In the following, let $X = (X_n)_{n \in \mathbb{N}_0}$ be a Markov chain on the countable space E with transition matrix p .

Definition 17.28 For any $x \in E$, let $\tau_x := \tau_x^1 := \inf \{n > 0 : X_n = x\}$ and

$$\tau_x^k = \inf \{n > \tau_x^{k-1} : X_n = x\} \quad \text{for } k \in \mathbb{N}, k \geq 2$$

τ_x^k is the k th entrance time of X for x . For $x, y \in E$, let

$$F(x, y) := \mathbf{P}_x[\tau_y^1 < \infty] = \mathbf{P}_x[\text{there is an } n \geq 1 \text{ with } X_n = y]$$

be the probability of ever going from x to y . In particular, $F(x, x)$ is the return probability (after the first jump) from x to x .

Note that $\tau_x^1 > 0$ even if we start the chain at $X_0 = x$.




Fig. 17.1 Markov chain with eight states. The numbers are the transition probabilities for the corresponding arrows. State 2 is absorbing, the states 1, 3, 4 and 5 are transient and the states 6, 7 and 8 are (positive) recurrent

Theorem 17.29 For all $x, y \in E$ and $k \in \mathbb{N}$, we have

$$\mathbf{P}_x [\tau_y^k < \infty] = F(x, y)F(y, y)^{k-1}.$$

Proof We carry out the proof by induction on k . For $k = 1$, the claim is true by definition. Now let $k \geq 2$. Using the strong Markov property of X (see Theorem 17.14), we get

$$\begin{aligned} \mathbf{P}_x [\tau_y^k < \infty] &= \mathbf{E}_x \left[\mathbf{P}_x [\tau_y^k < \infty \mid \mathcal{F}_{\tau_y^{k-1}}] \mathbb{1}_{\{\tau_y^{k-1} < \infty\}} \right] \\ &= \mathbf{E}_x \left[F(y, y) \cdot \mathbb{1}_{\{\tau_y^{k-1} < \infty\}} \right] \\ &= F(y, y) \cdot F(x, y)F(y, y)^{k-2} = F(x, y)F(y, y)^{k-1}. \end{aligned}$$

□

Definition 17.30 A state $x \in E$ is called

- recurrent if $F(x, x) = 1$,
- positive recurrent if $\mathbf{E}_x [\tau_x^1] < \infty$,
- null recurrent if x is recurrent but not positive recurrent,
- transient if $F(x, x) < 1$, and
- absorbing if $p(x, x) = 1$.

The Markov chain X is called (positive/null) recurrent if every state $x \in E$ is (positive/null) recurrent and is called transient if every recurrent state is absorbing.

Remark 17.31 Clearly, we have:

$$\text{"absorbing"} \implies \text{"positive recurrent"} \implies \text{"recurrent"}.$$

Example 17.32

(i) In Fig. 17.1, the state 2 is absorbing. If it does not get trapped in 2, the chain will eventually jump from 5 to 6 and will not return after that. Hence 1, 3, 4 and 5 are transient. The states 6, 7 and 8 are positive recurrent. One can show (see Exercise 17.6.1) that $\mathbf{E}_6 [\tau_6] = \frac{17}{4}$, $\mathbf{E}_7 [\tau_7] = \frac{17}{5}$ and $\mathbf{E}_8 [\tau_8] = \frac{17}{8}$.




Fig. 17.2 Markov chain on $\mathbb{N}_{-0}$ with parameter $r \in (0, 1)$. The chain is positive recurrent if $r \in (0, 1/2)$, null recurrent if $r = 1/2$ and transient if $r \in (1/2, 1)$

(ii) The chain in Fig. 17.2 has a drift to the right if $r > \frac{1}{2}$. Hence, in this case, every state is transient. On the other hand, if $r \in (0, \frac{1}{2})$, then the chain has a drift to the left (except at the point 0) and hence visits every state again and again. Thus the chain is recurrent. With a little thought, one can show (see Exercise 17.6.4) that in this case, the chain is actually positive recurrent and in the remaining case $r = \frac{1}{2}$ it is null recurrent. $\square$

Definition 17.33 Denote by $N(y) = \sum_{n=0}^{\infty} \mathbb{1}_{\{X_n=y\}}$ the total number of visits of X to y and by

$$G(x, y) = \mathbf{E}_x[N(y)] = \sum_{n=0}^{\infty} p^n(x, y)$$

the Green function of X .

Theorem 17.34

(i) For all $x, y \in E$, we have (with the convention $1/0 := \infty, 0/0 := 0$ and $0 \cdot \infty := 0$)

$$G(x, y) = \begin{cases} \frac{F(x, y)}{1 - F(y, y)}, & \text{if } x \neq y \\ \frac{1}{1 - F(y, y)}, & \text{if } x = y \end{cases} = F(x, y)G(y, y) + \mathbb{1}_{\{x=y\}}. \quad (17.17)$$

(ii) A non-absorbing state $x \in E$ is recurrent if and only if $G(x, x) = \infty$.

Proof (ii) follows by (i). Hence, it remains to show (17.17). By Theorem 17.29, we have

$$\begin{aligned} G(x, y) &= \mathbf{E}_x[N(y)] = \sum_{k=1}^{\infty} \mathbf{P}_x[N(y) \geq k] \\ &= \mathbb{1}_{\{x=y\}} + \sum_{k=1}^{\infty} \mathbf{P}_x[\tau_y^k < \infty] = \mathbb{1}_{\{x=y\}} + \sum_{k=1}^{\infty} F(x, y)F(y, y)^{k-1} \\ &= \begin{cases} \frac{F(x, y)}{1 - F(y, y)}, & \text{if } x \neq y \\ \frac{1}{1 - F(y, y)}, & \text{if } x = y \end{cases} \end{aligned}$$

The second equality in (17.17) is an immediate consequence. $\square$

Theorem 17.35 If x is recurrent and $F(x, y) > 0$, then y is also recurrent, and $F(x, y) = F(y, x) = 1$.

Proof Let $x, y \in E, x \neq y$, be such that $F(x, y) > 0$. Then there is a $k \in \mathbb{N}$ and states $x_1, \dots, x_k \in E$ with $x_k = y$ and $x_i \neq x$ for all $i = 1, \dots, k$ and such that

$$\mathbf{P}_x[X_i = x_i \text{ for all } i = 1, \dots, k] > 0.$$

In particular, $p^k(x, y) > 0$. By the Markov property, we have

$$\begin{aligned} 1 - F(x, x) &= \mathbf{P}_x[\tau_x^1 = \infty] \geq \mathbf{P}_x[X_1 = x_1, \dots, X_k = x_k, \tau_x^1 = \infty] \\ &= \mathbf{P}_x[X_1 = x_1, \dots, X_k = x_k] \cdot \mathbf{P}_y[\tau_x^1 = \infty] \\ &= \mathbf{P}_x[X_1 = x_1, \dots, X_k = x_k] (1 - F(y, x)) \end{aligned}$$

If now $F(x, x) = 1$, then also $F(y, x) = 1$. Since $F(y, x) > 0$, there exists an $l \in \mathbb{N}$ with $p^l(y, x) > 0$. Hence, for $n \in \mathbb{N}_0$,

$$p^{l+n+k}(y, y) \geq p^l(y, x)p^n(x, x)p^k(x, y).$$

If x is recurrent, then we conclude that

$$G(y, y) \geq \sum_{n=0}^{\infty} p^{l+n+k}(y, y) \geq p^l(y, x)p^k(x, y)G(x, x) = \infty$$

and hence also that y is recurrent. Changing the roles of x and y in the above argument, we get $F(x, y) = 1$.

Definition 17.36 A discrete Markov chain is called

- irreducible if $F(x, y) > 0$ for all $x, y \in E$, or equivalently $G(x, y) > 0$, and
- weakly irreducible if $F(x, y) + F(y, x) > 0$ for all $x, y \in E$.

Theorem 17.37 An irreducible discrete Markov chain is either recurrent or transient. If $|E| \geq 2$, then there is no absorbing state.

Proof This follows directly from Theorem 17.35.

Theorem 17.38 If E is finite and X is irreducible, then X is recurrent.

Proof Evidently, for all $x \in E$,

$$\sum_{y \in E} G(x, y) = \sum_{n=0}^{\infty} \sum_{y \in E} p^n(x, y) = \sum_{n=0}^{\infty} 1 = \infty.$$

As E is finite, there is a $y \in E$ with $G(x, y) = \infty$. Since $F(y, x) > 0$, there exists a $k \in \mathbb{N}$ with $p^k(y, x) > 0$. Therefore, since $p^{n+k}(x, x) \geq p^n(x, y)p^k(y, x)$, we have

$$G(x, x) \geq \sum_{n=0}^{\infty} p^n(x, y)p^k(y, x) = p^k(y, x)G(x, y) = \infty.$$

Exercise 17.4.1 Let x be positive recurrent and let $F(x, y) > 0$. Show that y is also positive recurrent.

5.17.5 17.5 Application: Recurrence and Transience of Random Walks

In this section, we study recurrence and transience of random walks on the D dimensional integer lattice $\mathbb{Z}^D$, $D = 1, 2, \dots$. A more exhaustive investigation can be found in Spitzer's book [158].

Consider first the simplest situation of symmetric simple random walk X on $\mathbb{Z}^D$. That is, at each step, X jumps to any of its $2D$ neighbors with the same probability $1/2D$. Hence, in terms of the Markov chain notation, we have $E = \mathbb{Z}^D$ and

$$p(x, y) = \begin{cases} \frac{1}{2D}, & \text{if } |x - y| = 1 \\ 0, & \text{else} \end{cases}$$

Is this random walk recurrent or transient?

The central limit theorem suggests that

$$p^n(0, 0) \approx C_D n^{-D/2} \quad \text{as } n \rightarrow \infty$$

for some constant C_D that depends on the dimension D . However, first we have to exclude the case where n is odd since here clearly $p^n(0, 0) = 0$. Thus let $Y_1, Y_2, \dots$ be independent $\mathbb{Z}^D$ -valued random variables with $\mathbf{P}[Y_i = x] = p^2(0, x)$. Then $X_{2n} \stackrel{\mathcal{D}}{=} S_n := Y_1 + \dots + Y_n$ for $n \in \mathbb{N}_0$; hence $G(0, 0) = \sum_{n=0}^{\infty} \mathbf{P}[S_n = 0]$. Clearly, $Y_1 = (Y_1^1, \dots, Y_1^D)$ has covariance matrix $C_{i,j} := \mathbf{E}[Y_1^i \cdot Y_1^j] = \frac{2}{D} \mathbb{1}_{\{i=j\}}$. By the local central limit theorem (see, e.g., [20, pp. 224ff] for a one-dimensional version of that theorem or Exercise 17.5.1 for an analytic derivation), we have

$$n^{D/2} p^{2n}(0, 0) = n^{D/2} \mathbf{P}[S_n = 0] \xrightarrow{n \rightarrow \infty} 2(4\pi/D)^{-D/2} \quad (17.18)$$

Now $\sum_{n=1}^{\infty} n^{-\alpha} < \infty$ if and only if $\alpha > 1$. Hence $G(0, 0) < \infty$ if and only if $D > 2$. We have thus shown the following theorem of Pólya [134].

Theorem 17.39 (Pólya (1921) [134]) Symmetric simple random walk on $\mathbb{Z}^D$ is recurrent if and only if $D \leq 2$.

The procedure we used here to derive Pólya's theorem has the disadvantage that it relies on the local central limit theorem, which we have not proved (and will not). Hence we will consider different methods of proof that yield further insight into the problem.

Consider first the one-dimensional simple random walk that with probability p jumps one step to the right and with probability $1 - p$ jumps one step to the left.

Then

$$G(0, 0) = \sum_{n=0}^{\infty} \binom{2n}{n} (p(1-p))^n = \sum_{n=0}^{\infty} \binom{-1/2}{n} (-4p(1-p))^n$$

Using the generalized binomial theorem (see Lemma 3.5), we get (since we have $(1 - 4p(1-p))^{1/2} = |2p - 1|$)

$$G(0, 0) = \begin{cases} \frac{1}{|2p-1|}, & \text{if } p \neq \frac{1}{2} \\ \infty, & \text{if } p = \frac{1}{2} \end{cases} \quad (17.19)$$

Thus, simple random walk on $\mathbb{Z}$ is recurrent if and only if it is symmetric; that is, if $p = \frac{1}{2}$.

Of course, transience in the case $p \neq \frac{1}{2}$ could also be deduced directly from the strong law of large numbers since $\lim_{n \rightarrow \infty} \frac{1}{n} X_n = \mathbf{E}_0[X_1] = 2p - 1$ almost surely. In fact, this argument is even more robust since it uses only that the single steps of X have an expectation that is not zero.

Consider now the situation where X does not necessarily jump to one of its nearest neighbors but where we still have $\mathbf{E}_0[|X_1|] < \infty$ and $\mathbf{E}_0[X_1] = 0$. The strong law of large numbers does not yield recurrence immediately and we have to do some work:

By the Markov property, for every $N \in \mathbb{N}$ and every $y \neq x$,

$$G_N(x, y) := \sum_{k=0}^N \mathbf{P}_x[X_k = y] = \sum_{k=0}^N \mathbf{P}_x[\tau_y^1 = k] \sum_{l=0}^{N-k} \mathbf{P}_y[X_l = y] \leq G_N(y, y).$$

This implies for all $L \in \mathbb{N}$

$$\begin{aligned} G_N(0, 0) &\geq \frac{1}{2L+1} \sum_{|y| \leq L} G_N(0, y) \\ &= \frac{1}{2L+1} \sum_{k=0}^N \sum_{|y| \leq L} p^k(0, y) \\ &\geq \frac{1}{2L+1} \sum_{k=1}^N \sum_{y: |y/k| \leq L/N} p^k(0, y) \end{aligned}$$

By the weak law of large numbers, we have $\liminf_{k \rightarrow \infty} \sum_{|y| \leq \varepsilon k} p^k(0, y) = 1$ for every $\varepsilon > 0$. Hence, letting $L = \varepsilon N$, we get

$$\liminf_{N \rightarrow \infty} G_N(0, 0) \geq \frac{1}{2\varepsilon} \quad \text{for every } \varepsilon > 0$$

Thus $G(0, 0) = \infty$, which shows that X is recurrent.

We summarize the discussion in a theorem.

Theorem 17.40 A random walk on $\mathbb{Z}$ with $\sum_{x=-\infty}^{\infty} |x|p(0, x) < \infty$ is recurrent if and only if $\sum_{x=-\infty}^{\infty} xp(0, x) = 0$.

Now what about symmetric simple random walk in dimension $D = 2$ or in higher dimensions? In order that the random walk be at the origin after $2n$ steps, it must perform k_i steps in the i th direction and k_i steps in the opposite direction for some numbers $k_1, \dots, k_D \in \mathbb{N}_0$ with $k_1 + \dots + k_D = n$. We thus get

$$p^{2n}(0, 0) = (2D)^{-2n} \sum_{k_1 + \dots + k_D = n} \binom{2n}{k_1, k_1, \dots, k_D, k_D}, \quad (17.20)$$

where $\binom{N}{l_1, \dots, l_r} = \frac{N!}{l_1! \dots l_r!}$ is the multinomial coefficient. In particular, for $D = 2$,

$$\begin{aligned} p^{2n}(0, 0) &= 4^{-2n} \sum_{k=0}^n \frac{(2n)!}{(k!)^2((n-k)!)^2} \\ &= 4^{-2n} \binom{2n}{n} \sum_{k=0}^n \binom{n}{k} \binom{n}{n-k} = \left(2^{-2n} \binom{2n}{n} \right)^2. \end{aligned}$$

Note that in the last step, we used a simple combinatorial identity that follows, e.g., by the convolution formula $(b_{n,p} * b_{n,p})(\{n\}) = b_{2n,p}(\{n\})$. Now, by Stirling's formula,

$$\lim_{n \rightarrow \infty} \sqrt{n} 2^{-2n} \binom{2n}{n} = \frac{1}{\sqrt{\pi}}$$

hence $\lim_{n \rightarrow \infty} n p^{2n}(0, 0) = \frac{1}{\pi}$. In particular, we have $\sum_{n=1}^{\infty} p^{2n}(0, 0) = \infty$. That is, two-dimensional symmetric simple random walk is recurrent.

For $D \geq 3$, the sum over the multinomial coefficients cannot be computed in a satisfactory way. However, it is not too hard to give an estimate that shows that there exists a $c = c_D$ such that $p^{2n}(0, 0) \leq c n^{-D/2}$, which implies $G(0, 0) \leq c \sum_{n=1}^{\infty} n^{-D/2} < \infty$ (see, e.g., [53, p. 361] or [59, Example 6.31]). Here, however, we follow a different route.

Things would be easy if the individual coordinates of the chain were independent one-dimensional random walks. In this case, the probability that at time $2n$ all coordinates are zero would be the D th power of the probability that the first coordinate is zero. For one coordinate, however, which moves only with probability $1/D$ and thus has variance $1/D$, the probability of being back at the origin at time $2n$ is approximately $(n\pi/D)^{-1/2}$. Up to a factor, we would thus get (17.18) without using the multidimensional local central limit theorem.

An elegant way to decouple the coordinates is to pass from discrete time to continuous time in such a way that the individual coordinates become independent but such that the Green function remains unchanged.

We give the details. Let $(T_t^i)_{t \geq 0}$, $i = 1, \dots, D$ be independent Poisson processes with rate $1/D$. Let $Z^1, \dots, Z^D$ be independent (and independent of the Poisson processes) symmetric simple random walks on $\mathbb{Z}$. Define $T := T^1 + \dots + T^D$, $Y_t^i := Z_{T_t^i}^i$ for $i = 1, \dots, D$ and let $Y_t = (Y_t^1, \dots, Y_t^D)$. Then Y is a Markov chain in continuous time with Q -matrix $q(x, y) = p(x, y) - \mathbb{1}_{\{x=y\}}$. As T is a Poisson process with rate 1, $(X_{T_t})_{t \geq 0}$ is also a Markov process with Q -matrix q . It follows that $(X_{T_t})_{t \geq 0} \stackrel{\mathcal{D}}{=} (Y_t)_{t \geq 0}$. We now compute

$$\begin{aligned}
G_Y &:= \int_0^\infty \mathbf{P}_0[Y_t = 0] dt = \int_0^\infty \sum_{n=0}^\infty \mathbf{P}_0[X_{2n} = 0, T_t = 2n] dt \\
&= \sum_{n=0}^\infty p^{2n}(0, 0) \int_0^\infty e^{-t} \frac{t^{2n}}{(2n)!} dt = G(0, 0)
\end{aligned}$$

The two processes $(X_n)_{n \in \mathbb{N}_0}$ and $(Y_t)_{t \in [0, \infty)}$ thus have the same Green function. As the coordinates of Y are independent, we have

$$G_Y = \int_0^\infty \mathbf{P}_0[Y_t^1 = 0]^D dt$$

Hence we only have to compute the asymptotics of $\mathbf{P}_0[Y_t^1 = 0]$ for large t . We can argue as follows. By the law of large numbers, we have $T_t^1 \approx t/D$ for large t . Furthermore, $\mathbf{P}_0[Y_t^1 \text{ is even}] \approx \frac{1}{2}$. Hence we have, with $n_t = \lfloor t/2D \rfloor$ for $t \rightarrow \infty$ (compare Exercise 17.5.2),

$$\mathbf{P}_0[Y_t^1 = 0] \sim \frac{1}{2} \mathbf{P}[Z_{2n_t}^1 = 0] = \frac{1}{2} \binom{2n_t}{n_t} 4^{-n_t} \sim (2\pi/D)^{-1/2} t^{-1/2}. \quad (17.21)$$

Since $\int_1^\infty t^{-\alpha} dt < \infty$ if and only if $\alpha > 1$, we also have $G_Y < \infty$ if and only if $D > 2$. However, this is the statement of Pólya's theorem.

Finally, we present a third method of studying recurrence and transience of random walks that does not rely on the Euclidean properties of the integer lattice but rather on the Fourier inversion formula.

First consider a general (discrete time) irreducible random walk with transition matrix p on $\mathbb{Z}^D$. By $\phi(t) = \sum_{x \in \mathbb{Z}^D} e^{i\langle t, x \rangle} p(0, x)$ denote the characteristic function of a single transition. The convolution of the transition probabilities translates into powers of the characteristic function; hence

$$\phi^n(t) = \sum_{x \in \mathbb{Z}^D} e^{i\langle t, x \rangle} p^n(0, x)$$

By the Fourier inversion formula (Theorem 15.10), we recover the n -step transition probabilities from ϕ^n by

$$p^n(0, x) = (2\pi)^{-D} \int_{[-\pi, \pi]^D} e^{-i\langle t, x \rangle} \phi^n(t) dt$$

In particular, for $\lambda \in (0, 1)$,

$$\begin{aligned}
R_\lambda &:= \sum_{n=0}^\infty \lambda^n p^n(0, 0) \\
&= (2\pi)^{-D} \sum_{n=0}^\infty \int_{[-\pi, \pi]^D} \lambda^n \phi^n(t) dt \\
&= (2\pi)^{-D} \int_{[-\pi, \pi]^D} \frac{1}{1 - \lambda \phi(t)} dt \\
&= (2\pi)^{-D} \int_{[-\pi, \pi]^D} \operatorname{Re} \left(\frac{1}{1 - \lambda \phi(t)} \right) dt
\end{aligned}$$

Now $G(0, 0) = \lim_{\lambda \uparrow 1} R_\lambda$ and hence

$$X \text{ is recurrent} \iff \lim_{\lambda \uparrow 1} \int_{[-\pi, \pi]^D} \operatorname{Re} \left(\frac{1}{1 - \lambda \phi(t)} \right) dt = \infty \quad (17.22)$$

If we had $\phi(t) = 1$ for some $t \in (-2\pi, 2\pi)^D \setminus \{0\}$, then we would have $\phi^n(t) = 1$ for every $n \in \mathbb{N}$ and hence, by Exercise 15.2.1, $\mathbf{P}_0[\langle X_n, t/(2\pi) \rangle \in \mathbb{Z}] = 1$. Thus X would not be irreducible contradicting the assumption. Due to the continuity of ϕ for all $\varepsilon > 0$, we thus have

$$\inf \{ |\phi(t) - 1| : t \in [-\pi, \pi)^D \setminus (-\varepsilon, \varepsilon)^D \} > 0$$

We summarize the discussion in a theorem due to Chung and Fuchs [26].

Theorem 17.41 (Chung-Fuchs (1951) [26]) An irreducible random walk on $\mathbb{Z}^D$ with characteristic function ϕ is recurrent if and only if, for every $\varepsilon > 0$,

$$\lim_{\lambda \uparrow 1} \int_{(-\varepsilon, \varepsilon)^D} \operatorname{Re} \left(\frac{1}{1 - \lambda \phi(t)} \right) dt = \infty \quad (17.23)$$

Now consider symmetric simple random walk. Here $\phi(t) = \frac{1}{D} \sum_{i=1}^D \cos(t_i)$. Expanding the cosine function in a Taylor series, we get $\cos(t_i) = 1 - \frac{1}{2}t_i^2 + O(t_i^4)$; hence $1 - \phi(t) = \frac{1}{2D} \|t\|_2^2 + O(\|t\|_2^4)$. We infer that X is recurrent if and only if $\int_{\|t\|_2 < \varepsilon} \|t\|_2^{-2} dt = \infty$. We compute this integral in polar coordinates (with C_D the surface of the unit sphere in $\mathbb{R}^D$):

$$\int_{\|t\|_2 < \varepsilon} \|t\|_2^{-2} dt = C_D \int_0^\varepsilon r^{D-1} r^{-2} dr = \infty \iff D \leq 2$$

Hence, X is recurrent if and only if $D \leq 2$.

In Section 19.3, we will encounter a further method of proving Pólya's theorem that has a completely different structure and that is based on the connection between Markov chains and electrical networks.

In fact, the Chung-Fuchs theorem can be used to compute the numerical values of the Green function $G_D(0, 0)$ of symmetric simple random walk on $\mathbb{Z}^D$ if we compute numerically the so-called Watson integral

$$G_D(0, 0) = (2\pi)^{-D} \int_{[-\pi, \pi)^D} \frac{D}{D - (\cos(x_1) + \dots + \cos(x_D))} dx \quad (17.24)$$

For this purpose, we follow [80] (where there are further refinements of the method) to transform the D -fold integral into a double integral. Denote by

$$I_0(t) := \frac{1}{\pi} \int_0^\pi e^{t \cos(\theta)} d\theta$$

the so-called modified Bessel function of the first kind. Using the identity $\frac{1}{\lambda} = \int_0^\infty e^{-\lambda t} dt$ for the integrand and applying Fubini's theorem, we get

$$G_D(0, 0) = \frac{D}{(2\pi)^D} \int_0^\infty e^{-Dt} \left(\int_{[-\pi, \pi)^D} e^{t(\cos(x_1) + \dots + \cos(x_D))} dx \right) dt$$

and thus

$$G_D(0, 0) = D \int_0^\infty e^{-Dt} I_0(t)^D dt \quad (17.25)$$

The right-hand side of (17.25) can quickly be computed numerically with great accuracy (see Table 17.1).

For the case $D = 3$, Watson [168] found the expression

$$G_3(0, 0) = 12 \frac{18 + 12\sqrt{2} - 10\sqrt{3} - 7\sqrt{6}}{\pi^2} K((2 - \sqrt{3})(\sqrt{3} - \sqrt{2}))^2,$$

where $K(m) = \int_0^1 ((1-t^2)(1-mt^2))^{-1/2} dt$ is the complete elliptic integral of the first kind with modulus $m \in (-1, 1)$. This in turn can be expressed as a (quickly convergent) series

$$K(m) = \frac{\pi}{2} \left(1 + \sum_{n=1}^{\infty} \left(\frac{(2n)!}{4^n (n!)^2} \right)^2 m^2 \right)$$

Glasser and Zucker [61] found an expression as a product of four Gamma functions,

$$G_3(0,0) = \frac{\sqrt{6}}{32\pi^3} \Gamma\left(\frac{1}{24}\right) \Gamma\left(\frac{5}{24}\right) \Gamma\left(\frac{7}{24}\right) \Gamma\left(\frac{11}{24}\right) = 1.5163860591519780181\dots$$

Exercise 17.5.1 For $n \in \mathbb{N}_0$, let p^n be the matrix of n -step transition probabilities of simple symmetric random walk on $\mathbb{Z}^D$. For $n \in \mathbb{N}$, derive the formula (see Theorem 15.10)

$$p^{2n}(0,0) = (2\pi)^{-D} \int_{[-\pi,\pi]^D} D^{-2n} (\cos(t_1) + \dots + \cos(t_D))^{2n} dt$$

Table 17.1 Green function $G_D(0,0)$

$(0,0)$ and return probability $F_D(0,0)$ of simple symmetric random walk on $\mathbb{Z}^D$. The numerical computations are based on (17.25)}

D	$G_D(0,0)$	$F_D(0,0)$
2	∞	1
3	1.51638605915	0.34053732955
4	1.23946712185	0.19320167322
5	1.15630812484	0.13517860982
6	1.11696337322	0.10471549562
7	1.09390631559	0.08584493411
8	1.07864701202	0.07291264996
9	1.06774608638	0.06344774965
10	1.05954374789	0.05619753597
11	1.05313615291	0.05045515982
12	1.04798637482	0.04578912090
13	1.04375406289	0.04191989708
14	1.04021240323	0.03865787709
15	1.03720412092	0.03586962312
16	1.03461657857	0.03345836447
17	1.03236691238	0.03135214040
18	1.03039276285	0.02949628913
19	1.02864627888	0.02784852234
20	1.02709011674	0.02637559869

By a suitable bound for the integral, conclude the convergence $n^{D/2} p^{2n}(0,0) \xrightarrow{n \rightarrow \infty} 2(4\pi/D)^{-D/2}$ (see (17.18)).

Exercise 17.5.2 Show (17.21) formally.

Exercise 17.5.3 Use Theorem 17.41 to show that a random walk on $\mathbb{Z}^2$ with $\sum_{x \in \mathbb{Z}^2} x p(0,x) = 0$ is recurrent if $\sum_{x \in \mathbb{Z}^2} \|x\|_2^2 p(0,x) < \infty$.

Exercise 17.5.4 Use Theorem 17.41 to show that, for $D \geq 3$ every irreducible random walk on $\mathbb{Z}^D$ is transient.

Exercise 17.5.5 Show (17.25) for $G_D(0,0)$ directly with the $p^{2n}(0,0)$ from (17.20) and using the representation of $I_0(t)$ as the series $I_0(t) = \sum_{k=0}^{\infty} (k!)^{-2} (t/2)^k$.

5.17.6 17.6 Invariant Distributions

In the following, let p be a stochastic matrix on the discrete space E and let $(X_n)_{n \in \mathbb{N}_0}$ be a corresponding Markov chain.

This section is devoted to the question: Which distributions are preserved under the dynamics of the Markov chain? Of course, often the chain will not stay put in a specific state but the distribution of the random state of the chain might nevertheless be the same for all times. If such an invariant distribution exists, we will see in Chapter 18 that under rather weak conditions, the distribution of a Markov chain (started in an arbitrary state) converges in the large time limit to such an invariant distribution.

Definition 17.42 If μ is a measure on E and $f : E \rightarrow \mathbb{R}$ is a map, then we write $\mu p(\{x\}) = \sum_{y \in E} \mu(\{y\})p(y, x)$ and $pf(x) = \sum_{y \in E} p(x, y)f(y)$ if the sums converge.

Definition 17.43 (i) A σ -finite measure μ on E is called an invariant measure if

$$\mu p = \mu.$$

A probability measure that is an invariant measure is called an invariant distribution. Denote by $\mathcal{I}$ the set of invariant distributions.

(ii) A function $f : E \rightarrow \mathbb{R}$ is called subharmonic if pf exists and if $f \leq pf$. f is called superharmonic if $f \geq pf$ and harmonic if $f = pf$.

Remark 17.44 In the terminology of linear algebra, an invariant measure is a left eigenvector of p corresponding to the eigenvalue 1. A harmonic function is a right eigenvector corresponding to the eigenvalue 1.

Lemma 17.45 If f is bounded and (sub-, super-) harmonic, then $(f(X_n))_{n \in \mathbb{N}_0}$ is a (sub-, super-) martingale with respect to the filtration $\mathbb{F} = \sigma(X)$ generated by X .

Proof Let f be bounded and subharmonic. Then

$$\begin{aligned} \mathbf{E}_x[f(X_n) \mid \mathcal{F}_{n-1}] &= \mathbf{E}_{X_{n-1}}[f(X_1)] = \sum_{y \in E} p(X_{n-1}, y) f(y) \\ &= pf(X_{n-1}) \geq f(X_{n-1}) \end{aligned}$$

Theorem 17.46 If any point is transient, then an invariant distribution does not exist.

Proof By assumption, $G(x, y) = \sum_{n=0}^{\infty} p^n(x, y) < \infty$ for all $x, y \in E$; hence $p^n(x, y) \xrightarrow{n \rightarrow \infty} 0$. For every probability measure μ on E , we thus have that $\mu p^n(\{x\}) \xrightarrow{n \rightarrow \infty} 0$. If μ was invariant, however, then we would have $\mu p^n(\{x\}) = \mu(\{x\})$ for all $n \in \mathbb{N}$.

Theorem 17.47 Let x be a recurrent state and let $\tau_x^1 = \inf\{n \geq 1 : X_n = x\}$. Then one invariant measure μ_x is defined by

$$\mu_x(\{y\}) = \mathbf{E}_x \left[\sum_{n=0}^{\tau_x^1 - 1} \mathbb{1}_{\{X_n = y\}} \right] = \sum_{n=0}^{\infty} \mathbf{P}_x[X_n = y, \tau_x^1 > n].$$

Proof First we have to show that $\mu_x(\{y\}) < \infty$ for all $y \in E$. For $y = x$, clearly $\mu_x(\{x\}) = 1$. For $y \neq x$ and $F(x, y) = 0$, we have $\mu_x(\{y\}) = 0$. Now let $y \neq x$ and $F(x, y) > 0$. As x is recurrent, we have $F(x, y) = F(y, x) = 1$ and y is recurrent (Theorem 17.35). Let

$$\hat{F}(x, y) = \mathbf{P}_x[\tau_x^1 > \tau_y^1].$$

Then $\hat{F}(x, y) > 0$ (otherwise y would not be visited). Changing the roles of x and y , we also get $\hat{F}(y, x) > 0$.

By the strong Markov property (Theorem 17.14), we have

$$\begin{aligned}
\mathbf{E}_y \left[\sum_{n=0}^{\tau_x^1-1} \mathbb{1}_{\{X_n=y\}} \right] &= 1 + \mathbf{E}_y \left[\sum_{n=\tau_y^1}^{\tau_x^1-1} \mathbb{1}_{\{X_n=y\}}; \tau_x^1 > \tau_y^1 \right] \\
&= 1 + (1 - \widehat{F}(y, x)) \mathbf{E}_y \left[\sum_{n=0}^{\tau_x^1-1} \mathbb{1}_{\{X_n=y\}} \right].
\end{aligned}$$

Hence,

$$\mathbf{E}_y \left[\sum_{n=0}^{\tau_x^1-1} \mathbb{1}_{\{X_n=y\}} \right] = \frac{1}{\widehat{F}(y, x)}.$$

Therefore,

$$\mu_x(\{y\}) = \mathbf{E}_x \left[\sum_{n=0}^{\tau_x^1-1} \mathbb{1}_{\{X_n=y\}} \right] = \mathbf{E}_x \left[\sum_{n=\tau_y^1}^{\tau_x^1-1} \mathbb{1}_{\{X_n=y\}}; \tau_x^1 > \tau_y^1 \right] = \frac{\widehat{F}(x, y)}{\widehat{F}(y, x)} < \infty.$$

Define $\bar{p}_n(x, y) = \mathbf{P}_x [X_n = y; \tau_x^1 > n]$. Then, for every $z \in E$,

$$\mu_x p(\{z\}) = \sum_{y \in E} \mu_x(\{y\}) p(y, z) = \sum_{n=0}^{\infty} \sum_{y \in E} \bar{p}_n(x, y) p(y, z).$$

Case 1: $x \neq z$. In this case,

$$\begin{aligned}
\sum_{y \in E} \bar{p}_n(x, y) p(y, z) &= \sum_{y \in E} \mathbf{P}_x [X_n = y, \tau_x^1 > n, X_{n+1} = z] \\
&= \mathbf{P}_x [\tau_x^1 > n+1; X_{n+1} = z] = \bar{p}_{n+1}(x, z).
\end{aligned}$$

Hence (since $\bar{p}_0(x, z) = 0$)

$$\mu_x p(\{z\}) = \sum_{n=0}^{\infty} \bar{p}_{n+1}(x, z) = \sum_{n=1}^{\infty} \bar{p}_n(x, z) = \sum_{n=0}^{\infty} \bar{p}_n(x, z) = \mu_x(\{z\})$$

Case 2: $x = z$. In this case, we have

$$\sum_{y \in E} \bar{p}_n(x, y) p(y, x) = \sum_{y \in E} \mathbf{P}_x [X_n = y; \tau_x^1 > n; X_{n+1} = x] = \mathbf{P}_x [\tau_x^1 = n+1]$$

Thus (since $\mathbf{P}_x [\tau_x^1 = 0] = 0$)

$$\mu_x p(\{x\}) = \sum_{n=0}^{\infty} \mathbf{P}_x [\tau_x^1 = n+1] = 1 = \mu_x(\{x\})$$

Corollary 17.48 If X is positive recurrent, then $\pi := \frac{\mu_x}{\mathbf{E}_x[\tau_x^1]}$ is an invariant distribution for any $x \in E$.

Theorem 17.49 If X is irreducible, then X has at most one invariant distribution.

Remark 17.50

(i) One could in fact show that if X is irreducible and recurrent, then an invariant measure of X is unique up to a multiplicative factor. However, the proof is a little more involved. Since we will not need the statement here, we leave its proof as an exercise (compare Exercise 17.6.6; see also [39, Section 6.5]).

(ii) For transient X , there can be more than one invariant measure. For example, consider the asymmetric random walk on $\mathbb{Z}$ that jumps one step to the right with probability r and one step to the left with probability $1 - r$ (for some $r \in (0, 1)$). The invariant measures are the nonnegative linear combinations of the measures μ_1 and μ_2 given by $\mu_1(\{x\}) \equiv 1$ and $\mu_2(\{x\}) = (r/(1 - r))^x$, $x \in \mathbb{Z}$. X is transient if and only if $r \neq 1/2$, in which case we have $\mu_1 \neq \mu_2$. $\diamond$

Proof Let π and ν be invariant distributions. Choose an arbitrary probability vector $(g_n)_{n \in \mathbb{N}}$ with $g_n > 0$ for all $n \in \mathbb{N}$. Define the stochastic matrix $\tilde{p}(x, y) = \sum_{n=1}^{\infty} g_n p^n(x, y)$. Then $\tilde{p}(x, y) > 0$ for all $x, y \in E$ and $\pi \tilde{p} = \pi$ as well as $\nu \tilde{p} = \nu$.

Consider now the signed measure $\mu = \pi - \nu$. We have $\mu \tilde{p} = \mu$. If we had $\mu \neq 0$, then there would exist (since $\mu(E) = 0$) points $x_1, x_2 \in E$ with $\mu(\{x_1\}) > 0$ and $\mu(\{x_2\}) < 0$. Clearly, for every $y \in E$, this would imply $|\mu(\{x_1\}) \tilde{p}(x_1, y) + \mu(\{x_2\}) \tilde{p}(x_2, y)| < |\mu(\{x_1\}) \tilde{p}(x_1, y)| + |\mu(\{x_2\}) \tilde{p}(x_2, y)|$; hence

$$\begin{aligned} \|\mu \tilde{p}\|_{TV} &= \sum_{y \in E} \left| \sum_{x \in E} \mu(\{x\}) \tilde{p}(x, y) \right| \\ &< \sum_{y \in E} \sum_{x \in E} |\mu(\{x\})| \tilde{p}(x, y) = \sum_{x \in E} |\mu(\{x\})| = \|\mu\|_{TV}. \end{aligned}$$

Since this is a contradiction, we conclude that $\mu = 0$.

Recall that $\mathcal{I}$ is the set of invariant distributions of X .

Theorem 17.51 Let X be irreducible. X is positive recurrent if and only if $\mathcal{I} \neq \emptyset$. In this case, $\mathcal{I} = \{\pi\}$ with

$$\pi(\{x\}) = \frac{1}{\mathbf{E}_x[\tau_x^1]} > 0 \quad \text{for all } x \in E$$

Proof If X is positive recurrent, then $\mathcal{I} \neq \emptyset$ by Corollary 17.48. Now let $\mathcal{I} \neq \emptyset$ and $\pi \in \mathcal{I}$. As X is irreducible, we have $\pi(\{x\}) > 0$ for all $x \in E$. Let $\mathbf{P}_\pi = \sum_{x \in E} \pi(\{x\}) \mathbf{P}_x$. Fix an $x \in E$ and for $n \in \mathbb{N}_0$, let

$$\sigma_x^n = \sup \{m \leq n : X_m = x\} \in \mathbb{N}_0 \cup \{-\infty\}$$

be the time of last entrance in x before time n . (Note that this is not a stopping time.) By the Markov property, for all $k \leq n$,

$$\begin{aligned} \mathbf{P}_\pi[\sigma_x^n = k] &= \mathbf{P}_\pi[X_k = x, X_{k+1} \neq x, \dots, X_n \neq x] \\ &= \mathbf{P}_\pi[X_{k+1} \neq x, \dots, X_n \neq x \mid X_k = x] \mathbf{P}_\pi[X_k = x] \\ &= \pi(\{x\}) \mathbf{P}_x[X_1, \dots, X_{n-k} \neq x] \\ &= \pi(\{x\}) \mathbf{P}_x[\tau_x^1 \geq n - k + 1]. \end{aligned}$$

Hence, for every $n \in \mathbb{N}_0$ (since $\mathbf{P}_y[\tau_x^1 < \infty] = 1$ for all $y \in E$),

$$\begin{aligned}
1 &= \sum_{k=0}^n \mathbf{P}_\pi [\sigma_x^n = k] + \mathbf{P}_\pi [\sigma_x^n = -\infty] \\
&= \pi(\{x\}) \sum_{k=0}^n \mathbf{P}_x [\tau_x^1 \geq n - k + 1] + \mathbf{P}_\pi [\tau_x^1 \geq n + 1] \\
&\xrightarrow{n \rightarrow \infty} \pi(\{x\}) \sum_{k=1}^{\infty} \mathbf{P}_x [\tau_x^1 \geq k] = \pi(\{x\}) \mathbf{E}_x [\tau_x^1].
\end{aligned}$$

Therefore, $\mathbf{E}_x [\tau_x^1] = \frac{1}{\pi(\{x\})} < \infty$, and thus X is positive recurrent.

Example 17.52 Let $(p_x)_{x \in \mathbb{N}_0}$ be numbers in $(0, 1]$ and let X be an irreducible Markov chain on $\mathbb{N}_0$ with transition matrix

$$p(x, y) = \begin{cases} p_x, & \text{if } y = x + 1 \\ 1 - p_x, & \text{if } y = 0 \\ 0, & \text{else.} \end{cases}$$

If μ is an invariant measure, then the equations for $\mu p = \mu$ read

$$\begin{aligned}
\mu(\{n\}) &= p_{n-1} \mu(\{n-1\}) \quad \text{for } n \in \mathbb{N}, \\
\mu(\{0\}) &= \sum_{n=0}^{\infty} \mu(\{n\}) (1 - p_n).
\end{aligned}$$

Hence we get

$$\mu(\{n\}) = \mu(\{0\}) \prod_{k=0}^{n-1} p_k$$

and (note that the sum is a telescope sum)

$$\mu(\{0\}) = \mu(\{0\}) \sum_{n=0}^{\infty} (1 - p_n) \prod_{k=0}^{n-1} p_k = \mu(\{0\}) \left(1 - \prod_{n=0}^{\infty} p_n \right).$$

Hence there exists a nontrivial invariant measure μ (that is, $\mu(\{0\})$ can be chosen strictly positive) if and only if $\prod_{n=0}^{\infty} p_n = 0$. This, however, is true if and only if $\sum_{n=0}^{\infty} (1 - p_n) = \infty$. Using a Borel-Cantelli argument, it is not hard to show that this is exactly the condition for recurrence of X .

If $\mu \neq 0$, then μ is a finite measure if and only if

$$M := \sum_{n=0}^{\infty} \prod_{k=0}^{n-1} p_k < \infty$$

Hence X is positive recurrent if and only if $M < \infty$. In fact, it is not hard to show that M is the expected time to return to 0; hence the criterion for positive recurrence could also be deduced by Theorem 17.51.

A necessary condition for $M < \infty$ is of course that the series $\sum_{n=0}^{\infty} (1 - p_n)$ diverge; that is, that X is recurrent. One sufficient condition for $M < \infty$ is

$$\sum_{n=0}^{\infty} \exp \left(- \sum_{k=0}^{n-1} (1 - p_k) \right) < \infty$$

Exercise 17.6.1 Consider the Markov chain from Fig. 17.1 (p. 368). Determine the set of all invariant distributions. Show that the states 6, 7 and 8 are positive recurrent and compute the expected first entrance times

$$\mathbf{E}_6[\tau_6] = \frac{17}{4}, \quad \mathbf{E}_7[\tau_7] = \frac{17}{5} \quad \text{and} \quad \mathbf{E}_8[\tau_8] = \frac{17}{8}$$

Exercise 17.6.2 Let $X = (X_t)_{t \geq 0}$ be a Markov chain on E in continuous time with Q -matrix q . Show that a probability measure π on E is an invariant distribution for X if and only if $\sum_{x \in E} \pi(\{x\})q(x, y) = 0$ for all $y \in E$.

Exercise 17.6.3 Let G be a countable Abelian group and let p be the transition matrix of an irreducible random walk X on G . That is, we have $p(hg, hf) = p(g, f)$ for all $h, g, f \in G$. (This generalizes the notion of a random walk on $\mathbb{Z}^D$.) Use Theorem 17.51 to show that X is positive recurrent if and only if G is finite.

Exercise 17.6.4 Let $r \in [0, 1]$ and let X be the Markov chain on $\mathbb{N}_0$ with transition matrix (see Fig. 17.2 on p. 369)

$$p(x, y) = \begin{cases} 1, & \text{if } x = 0 \text{ and } y = 1, \\ r, & \text{if } y = x + 1 \geq 2, \\ 1 - r, & \text{if } y = x - 1, \\ 0, & \text{else.} \end{cases}$$

Compute the invariant measure and show the following using Theorem 17.51:

(i) If $r \in (0, \frac{1}{2})$, then X is positive recurrent.

(ii) If $r = \frac{1}{2}$, then X is null recurrent.

(iii) If $r \in \{0\} \cup (\frac{1}{2}, 1]$, then X is transient.

Exercise 17.6.5 (i) Use a direct argument to show that the Markov chain in Example 17.52 is recurrent if and only if $\sum_{n=0}^{\infty} (1 - p_n) = \infty$.

(ii) Show that the expected time to return to 0 is M and infer that the chain is positive recurrent if and only if $M < \infty$.

(iii) Give examples of sequences $(p_x)_{x \in \mathbb{N}_0}$ such that the chain is (a) transient, (b) null recurrent, (c) positive recurrent, and (d) positive recurrent but

$$\sum_{n=0}^{\infty} \exp \left(- \sum_{k=0}^{n-1} (1 - p_k) \right) = \infty.$$

Exercise 17.6.6 Let X be irreducible and recurrent. Show that, as claimed in Remark 17.50, the invariant measure is unique up to constant multiples.

Hint: Let $\pi \neq 0$ be an invariant measure for X and abbreviate

$$\mathbf{P}_\pi = \sum_{x \in E} \pi(\{x\}) \mathbf{P}_x$$

(note that, in general, this need not be a finite measure). Let $x, y \in E$ with $x \neq y$ and deduce by induction that

$$\pi(\{y\}) = \mathbf{P}_\pi [\tau_x^1 \geq n, X_0 \neq x, X_n = y] + \sum_{k=1}^n \mathbf{P}_\pi [\tau_x^1 \geq k, X_0 = x, X_k = y].$$

Infer that

$$\pi(\{y\}) \geq \sum_{k=1}^{\infty} \mathbf{P}_\pi [\tau_x^1 \geq k, X_0 = x, X_k = y] = \pi(\{x\})\mu_x(\{y\})$$

where μ_x is the invariant measure defined in Theorem 17.47. Now use the fact that $\pi p^n = \pi$ and $\mu_x p^n = \mu_x$ for all $n \in \mathbb{N}$ to conclude that even $\pi(\{y\}) = \pi(\{x\})\mu_x(\{y\})$ holds.

5.17.7 17.7 Stochastic Ordering and Coupling

In many situations, for the comparison of two distributions, it is helpful to construct a product space such that the two distributions are the marginal distributions but are not necessarily independent. We first introduce the abstract principle of such couplings and then give some examples.

There are many concepts to order probability measures on $\mathbb{R}$ or $\mathbb{R}^d$ such that the "larger" one has a greater preference for large values than the "smaller" one. As one of the most prominent orders we present here the so-called stochastic order and illustrate its connection with couplings. As an excuse for presenting this section in a chapter on Markov chains, we finally use a simple Markov chain in order to prove a theorem on the stochastic order of binomial distributions.

Definition 17.53 Let $(E_1, \mathcal{E}_1, \mu_1)$ and $(E_2, \mathcal{E}_2, \mu_2)$ be probability spaces. A probability measure μ on $(E_1 \times E_2, \mathcal{E}_1 \otimes \mathcal{E}_2)$ with $\mu(\cdot \times E_2) = \mu_1$ and $\mu(E_1 \times \cdot) = \mu_2$ is called a coupling of μ_1 and μ_2 .

Clearly, the product measure $\mu = \mu_1 \otimes \mu_2$ is a coupling, but in many situations there are more interesting ones.

Example 17.54 Let X be a real random variable and let $f, g : \mathbb{R} \rightarrow \mathbb{R}$ be monotone increasing functions with $\mathbf{E}[f(X)^2] < \infty$ and $\mathbf{E}[g(X)^2] < \infty$. We want to show that the random variables $f(X)$ and $g(X)$ are nonnegatively correlated.

To this end, let Y be an independent copy of X ; that is, a random variable with $\mathbf{P}_Y = \mathbf{P}_X$ that is independent of X . Note that $\mathbf{E}[f(X)] = \mathbf{E}[f(Y)]$ and $\mathbf{E}[g(X)] = \mathbf{E}[g(Y)]$. For all numbers $x, y \in \mathbb{R}$, we have $(f(x) - f(y))(g(x) - g(y)) \geq 0$. Hence

$$\begin{aligned} 0 &\leq \mathbf{E}[(f(X) - f(Y))(g(X) - g(Y))] \\ &= \mathbf{E}[f(X)g(X)] - \mathbf{E}[f(X)]\mathbf{E}[g(Y)] + \mathbf{E}[f(Y)g(Y)] - \mathbf{E}[f(Y)]\mathbf{E}[g(X)] \\ &= 2 \operatorname{Cov}[f(X), g(X)] \end{aligned}$$

Example 17.55 Let (E, ϱ) be a Polish space. For two probability measures P and Q on $(E, \mathcal{B}(E))$, denote by $K(P, Q) \subset \mathcal{M}_1(E \times E)$ the set of all couplings of P and Q . The so-called Wasserstein metric on $\mathcal{M}_1(E)$ is defined by

$$d_W(P, Q) := \inf \left\{ \int \varrho(x, y) \varphi(d(x, y)) : \varphi \in K(P, Q) \right\}. \quad (17.26)$$

It can be shown that (this is the Kantorovich-Rubinstein theorem [84]; see also [37, pp. 420ff])

$$d_W(P, Q) = \sup \left\{ \int f d(P - Q) : f \in \operatorname{Lip}_1(E; \mathbb{R}) \right\} \quad (17.27)$$

Compare this representation of the Wasserstein metric with that of the total variation norm,

$$\|P - Q\|_{TV} = \sup \left\{ \int f d(P - Q) : f \in \mathcal{L}^\infty(E) \text{ with } \|f\|_\infty \leq 1 \right\} \quad (17.28)$$

In fact, we can also give a definition for the total variation in terms of a coupling: Let $D := \{(x, x) : x \in E\}$ be the diagonal in $E \times E$. Then

$$\|P - Q\|_{TV} = \inf \{ \varphi((E \times E) \setminus D) : \varphi \in K(P, Q) \}. \quad (17.29)$$

See [60] for a comparison of different metrics on $\mathcal{M}_1(E)$.

As an example of a more involved coupling, we quote the following theorem that is due to Skorohod.

Theorem 17.56 (Skorohod coupling) Let $\mu, \mu_1, \mu_2, \dots$ be probability measures on a Polish space E with $\mu_n \xrightarrow{n \rightarrow \infty} \mu$. Then there exists a probability space $(\Omega, \mathcal{A}, \mathbf{P})$ with random variables $X, X_1, X_2, \dots$ with $\mathbf{P}_X = \mu$ and $\mathbf{P}_{X_n} = \mu_n$ for every $n \in \mathbb{N}$ such that $X_n \xrightarrow{n \rightarrow \infty} X$ almost surely.

Proof See, e.g., [83, p. 79].

We now come to the concept of stochastic order.

Definition 17.57 Let $\mu_1, \mu_2 \in \mathcal{M}_1(\mathbb{R}^d)$. We write $\mu_1 \leq_{st} \mu_2$ if

$$\int f d\mu_1 \leq \int f d\mu_2$$

for every monotone increasing bounded function $f : \mathbb{R}^d \rightarrow \mathbb{R}$. In this case, we say that μ_2 is stochastically larger than μ_1 .

Evidently, $\leq_{st}$ is a partial order on $\mathcal{M}_1(\mathbb{R}^d)$. The stochastic order belongs to the class of so-called integral orders that are defined by the requirement that the integrals with respect to a certain class of functions (here: monotone increasing and bounded) are ordered. Other classes of functions that are often considered are convex functions or indicator functions on lower or upper orthants.

Let F_1 and F_2 be the distribution functions of μ_1 and μ_2 . Clearly, $\mu_1 \leq_{st} \mu_2$ implies $F_1(x) \geq F_2(x)$ for all $x \in \mathbb{R}^d$. If $d = 1$, then both statements are equivalent. However, for $d \geq 2$, the condition $F_1 \geq F_2$ is weaker than $\mu_1 \leq_{st} \mu_2$. For example, consider $d = 2$ and

$$\mu_1 = \frac{1}{2}\delta_{(0,0)} + \frac{1}{2}\delta_{(1,1)} \quad \text{and} \quad \mu_2 = \frac{1}{2}\delta_{(1,0)} + \frac{1}{2}\delta_{(0,1)}.$$

The partial order defined by the comparison of the distribution functions is called (lower) orthant order.

For a survey on different orders of probability measures, see, e.g., [120].

The following theorem was shown by Strassen [160] in larger generality for integral orders.

Theorem 17.58 (Strassen's theorem) Let

$$L := \{(x_1, x_2) \in \mathbb{R}^d \times \mathbb{R}^d : x_1 \leq x_2\}.$$

Then $\mu_1 \leq_{st} \mu_2$ if and only if there is a coupling φ of μ_1 and μ_2 with $\varphi(L) = 1$.

Proof Let φ be such a coupling. For monotone increasing bounded $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we have $f(x_1) - f(x_2) \leq 0$ for every $x = (x_1, x_2) \in L$; hence $\int f d\mu_1 - \int f d\mu_2 = \int_L (f(x_1) - f(x_2)) \varphi(dx) \leq 0$ and thus $\mu_1 \leq_{st} \mu_2$.

Now assume $\mu_1 \leq_{st} \mu_2$. We only consider the case $d = 1$ (see [120, Theorem 3.3.5] for $d \geq 2$). Here $F((x_1, x_2)) := \min(F_1(x_1), F_2(x_2))$ defines a distribution function on $\mathbb{R} \times \mathbb{R}$ (see Exercise 1.5.5) that corresponds to a coupling φ with $\varphi(L) = 1$. A somewhat more explicit representation can be obtained using random variables. Let U be a random variable that is uniformly distributed on $(0, 1)$. Then

$$X_i := F_i^{-1}(U) := \inf \{x \in \mathbb{R} : F_i(x) \geq U\}$$

is a real random variable with distribution μ_i (see proof of Theorem 1.104). Clearly, we have $X_1 \leq X_2$ almost surely; that is, $\mathbf{P}[(X_1, X_2) \in L] = 1$. Evidently, the distribution function of (X_1, X_2) is F .

While Strassen's theorem yields the existence of an abstract coupling, in many examples a natural coupling can be established and used as a tool for proving, e.g., stochastic orders.

Example 17.59 Let $n \in \mathbb{N}$ and $0 \leq p_1 \leq p_2 \leq 1$. Let $Y_1, \dots, Y_n$ be independent random variables that are uniformly distributed on $[0, 1]$. Define $X_i = \#\{k \leq n : Y_k \leq p_i\}$, $i = 1, 2$. Then $X_i \sim b_{n, p_i}$ and $X_1 \leq X_2$ almost surely. This coupling shows that $b_{n, p_1} \leq_{\text{st}} b_{n, p_2}$.

An even simpler coupling can be used to show that $b_{m, p} \leq_{\text{st}} b_{n, p}$ for $m \leq n$ and $p \in [0, 1]$.

Theorem 17.60 Let $n_1, n_2 \in \mathbb{N}$ and $p_1, p_2 \in (0, 1)$. We have $b_{n_1, p_1} \leq_{\text{st}} b_{n_2, p_2}$ if and only if

$$(1 - p_1)^{n_1} \geq (1 - p_2)^{n_2} \quad (17.30)$$

and

$$n_1 \leq n_2. \quad (17.31)$$

Proof (The proof follows the exposition in [100, Section 3].)

Since $b_{n_i, p_i}(\{0\}) = (1 - p_i)^{n_i}$, conditions (17.30) and (17.31) are clearly necessary for $b_{n_1, p_1} \leq_{\text{st}} b_{n_2, p_2}$. Hence we only have to show sufficiency of the two conditions.

Assume that (17.30) and (17.31) hold. By Example 17.59, it is enough to consider the smallest p_2 that fulfills (17.30). Hence we assume $(1 - p_1)^{n_1} = (1 - p_2)^{n_2}$. Define $\lambda := -n_1 \log(1 - p_1) = -n_2 \log(1 - p_2)$. We will construct a binomially distributed random variable by throwing a Poi_λ -distributed number T of balls in n_i boxes and count the number of nonempty boxes. More precisely, let $T \sim \text{Poi}_\lambda$ and let $X_1, X_2, \dots$ be independent and uniformly distributed on $[0, 1]$ and independent of T . For $n \in \mathbb{N}$, $t \in \mathbb{N}_0$ and $l = 1, \dots, n$, define

$$M_{n, t, l} = \#\{s \leq t : X_s \in ((l-1)/n, l/n]\}$$

and the number of nonempty boxes after t balls are thrown:

$$N_{n, t} := \sum_{l=1}^n \mathbb{1}_{\{M_{n, t, l} > 0\}}$$

By Theorem 5.35, the random variables $M_{n, T, 1}, \dots, M_{n, T, n}$ are independent and $\text{Poi}_{\lambda/n}$ -distributed. In particular, we have

$$\mathbf{P}[M_{n_i, T, l} > 0] = 1 - e^{-\lambda/n_i} = p_i$$

and thus $N_{n_i, T} \sim b_{n_i, p_i}$, $i = 1, 2$. Hence it suffices to show that $N_{n_1, T} \leq_{\text{st}} N_{n_2, T}$. For this in turn it is enough to show

$$N_{n_1, t} \leq_{\text{st}} N_{n_2, t} \quad \text{for all } t \in \mathbb{N}_0. \quad (17.32)$$

In fact, let $f : \{0, \dots, n\} \rightarrow \mathbb{R}$ be monotone increasing. Then

$$\begin{aligned} \mathbf{E}[f(N_{n_1, T})] &= \sum_{t=0}^{\infty} \mathbf{E}[f(N_{n_1, t})] \mathbf{P}[T = t] \\ &\leq \sum_{t=0}^{\infty} \mathbf{E}[f(N_{n_2, t})] \mathbf{P}[T = t] = \mathbf{E}[f(N_{n_2, T})] \end{aligned}$$

We use an induction argument to show (17.32). For $t = 0$, the claim holds trivially. Now assume that (17.32) holds for some given $t \in \mathbb{N}_0$. We are now at the point to use a Markov chain. Note that (for fixed n), $(N_{n, t})_{t=0, 1, \dots}$ is a Markov chain with state space $\{0, \dots, n\}$ and transition matrix

$$p_n(k, l) = \begin{cases} k/n, & \text{if } l = k \\ 1 - k/n, & \text{if } l = k + 1 \\ 0, & \text{otherwise} \end{cases}$$

We define for $k, l = 0, \dots, n$

$$h_{n,l}(k) = \sum_{j=l}^n p_n(k, j) = \begin{cases} 0, & \text{if } k < l - 1 \\ 1 - k/n, & \text{if } k = l - 1 \\ 1, & \text{if } k > l - 1 \end{cases}$$

Then $\mathbf{P}[N_{n,t+1} \geq l] = \mathbf{E}[h_{n,l}(N_{n,t})]$ and $h_{n,l}(k)$ is monotone increasing both in k and in n . Hence by the induction hypothesis, we have

$$\begin{aligned} \mathbf{P}[N_{n_1,t+1} \geq l] &= \mathbf{E}[h_{n_1,l}(N_{n_1,t})] \\ &\leq \mathbf{E}[h_{n_1,l}(N_{n_2,t})] \\ &\leq \mathbf{1}(N_{n_2,t}) = \mathbf{P}[N_{n_2,t+1} \geq l]. \end{aligned}$$

We conclude that $N_{n_1,t+1} \leq_{\text{st}} N_{n_2,t+1}$ which completes the induction and the proof of the theorem.

Exercise 17.7.1 Use an elementary direct coupling argument to show the claim of Theorem 17.60 for the case $n_2/n_1 \in \mathbb{N}$.

Exercise 17.7.2 For the Poisson distribution, show that

$$\text{Poi}_{\lambda_1} \leq_{\text{st}} \text{Poi}_{\lambda_2} \iff \lambda_1 \leq \lambda_2.$$

Exercise 17.7.3 Let $n \in \mathbb{N}$, $p \in (0, 1)$ and $\lambda > 0$. Show that

$$b_{n,p} \leq_{\text{st}} \text{Poi}_{\lambda} \iff (1-p)^n \geq e^{-\lambda}.$$

5.18 18 Convergence of Markov Chains

We consider a Markov chain X with invariant distribution π and investigate conditions under which the distribution of X_n converges to π for $n \rightarrow \infty$. Essentially it is necessary and sufficient that the state space of the chain cannot be decomposed into subspaces

- that the chain does not leave
- or that are visited by the chain periodically; e.g., only for odd n or only for even n .

In the first case, the chain would be called reducible, and in the second case, it would be periodic.

We study periodicity of Markov chains in the first section. In the second section, we prove the convergence theorem. The third section is devoted to applications of the convergence theorem to computer simulations with the so-called Monte Carlo method. In the last section, we describe the speed of convergence to the equilibrium by means of the spectrum of the transition matrix.

5.18.1 18.1 Periodicity of Markov Chains

We study the conditions under which a positive recurrent Markov chain X on the countable space E (and with transition matrix p), started in an arbitrary $\mu \in \mathcal{M}_1(E)$, converges in distribution to an invariant distribution π ; that is, $\mu p^n \xrightarrow{n \rightarrow \infty} \pi$. Clearly, it is necessary that π be the unique invariant distribution; that is, up to a factor π it is the unique left eigenvector of p for the eigenvalue 1. As shown in Theorem 17.49, for this uniqueness it is sufficient that the chain be irreducible.

In order for $\mu p^n \xrightarrow{n \rightarrow \infty} \pi$ to hold for every $\mu \in \mathcal{M}_1(E)$, a certain contraction property of p is necessary. Manifestly, 1 is the largest (absolute value of an) eigenvalue of p . However, p is

sufficiently contractive only if the multiplicity of the eigenvalue 1 is exactly 1 and if there are no further (possibly complex-valued) eigenvalues of modulus 1.

For the latter property, it is not sufficient that the chain be irreducible. For example, consider on $E = \{0, \dots, N-1\}$ the Markov chain with transition matrix




Fig. 18.1 The left Markov chain is periodic with period 2, and the right Markov chain is aperiodic

$p(x, y) = \mathbb{1}_{\{y=x+1 \pmod{N}\}}$. The eigenvalue 1 has the multiplicity 1. However, all complex N th roots of unity $e^{2\pi i k/N}$, $k = 0, \dots, N-1$, are eigenvalues of modulus 1. Clearly, the uniform distribution on E is invariant but $\lim_{n \rightarrow \infty} \delta_x p^n$ does not exist for any $x \in E$. In fact, every point is visited periodically after N steps. In order to obtain criteria for the convergence of Markov chains, we thus have to understand periodicity first. Thereafter, for irreducible aperiodic chains, we state the convergence theorem.

If $m, n \in \mathbb{N}$, then write $m \mid n$ if m is a divisor of n ; that is, if $\frac{n}{m} \in \mathbb{N}$. If $M \subset \mathbb{N}$, then denote by $\gcd(M)$ the greatest common divisor of all $n \in M$. In the following, let X be a Markov chain on the countable space E with transition matrix p .

Definition 18.1 (i) For $x, y \in E$, define

$$N(x, y) := \{n \in \mathbb{N}_0 : p^n(x, y) > 0\}.$$

For any $x \in E$, $d_x := \gcd(N(x, x))$ is called the period of the state x .

(ii) If $d_x = d_y$ for all $x, y \in E$, then $d := d_x$ is called the period of X .

(iii) If $d_x = 1$ for all $x \in E$, then X is called aperiodic.

See Figs. 18.1 and 18.2 for illustrations of aperiodic and periodic Markov chains.

Lemma 18.2 For any $x \in E$, there exists an $n_x \in \mathbb{N}$ with

$$p^{nd_x}(x, x) > 0 \quad \text{for all } n \geq n_x. \quad (18.1)$$

Proof Let $k_1, \dots, k_r \in N(x, x)$ with $\gcd(\{k_1, \dots, k_r\}) = d_x$. Then, for all $m_1, \dots, m_r \in \mathbb{N}_0$, we also have $\sum_{i=1}^r k_i m_i \in N(x, x)$. Basic number theory then yields that, for every $n \geq n_x := r \cdot \prod_{i=1}^r (k_i/d_x)$, there are numbers $m_1, \dots, m_r \in \mathbb{N}_0$ with $nd_x = \sum_{i=1}^r k_i m_i$. Hence (18.1) holds. $\square$

The problem of finding the smallest number N such that any $nd_x, n \geq N$ can be written as a nonnegative integer linear combination of $k_1, \dots, k_r$ is called the Frobenius problem. The general solution is unknown; however, for the case $r = 2$,




Fig. 18.2 Here $N(8,8) = \{6, 10, 12, 14, 16, \dots\}$; hence $d_8 := \gcd(\{6, 10, 12, \dots\}) = 2$ and $n_8 = 5$. The chain thus has period 2. However, $n_1 = 2$ and $n_4 = 4$

Sylvester [162] showed that $N = (k_1/d_x - 1)(k_2/d_x - 1)$ is minimal. In the general case, for N , the upper bound $2 \max \{k_i : i = 1, \dots, r\}^2 / (rd_x^2)$ is known; see, e.g., [45].

Lemma 18.3 Let X be irreducible. Then the following statements hold.

- (i) $d := d_x = d_y$ for all $x, y \in E$.
- (ii) For all $x, y \in E$, there exist $n_{x,y} \in \mathbb{N}$ and $L_{x,y} \in \{0, \dots, d - 1\}$ such that

$$nd + L_{x,y} \in N(x, y) \quad \text{for all } n \geq n_{x,y}. \quad (18.2)$$

$L_{x,y}$ is uniquely determined, and we have

$$L_{x,y} + L_{y,z} + L_{z,x} = 0 \pmod{d} \quad \text{for all } x, y, z \in E. \quad (18.3)$$

Proof (i) Let $m, n \in \mathbb{N}_0$ with $p^m(x, y) > 0$ and $p^n(y, z) > 0$. Then

$$p^{m+n}(x, z) \geq p^m(x, y)p^n(y, z) > 0.$$

Hence we have

$$N(x, y) + N(y, z) := \{m + n : m \in N(x, y), n \in N(y, z)\} \subset N(x, z). \quad (18.4)$$


If, in particular, $m \in N(x, y), n \in N(y, x)$ and $k \geq n_y$, then $kd_y \in N(y, y)$; hence $m + kd_y \in N(x, y)$ and $m + n + kd_y \in N(x, x)$. Therefore, $d_x \mid (m + n + kd_y)$ for every $k \geq n_y$; hence $d_x \mid d_y$. Similarly, we get $d_y \mid d_x$; hence $d_x = d_y$.

(ii) Let $m \in N(x, y)$. Then $m + kd \in N(x, y)$ for every $k \geq n_x$. Hence (18.2) holds with

$$n_{x,y} := n_x + \left\lfloor \frac{m}{d} \right\rfloor \quad \text{and} \quad L_{x,y} := m - d \left\lfloor \frac{m}{d} \right\rfloor.$$

Owing to (18.4), we have

$$(n_{x,y} + n_{y,z})d + L_{x,y} + L_{y,z} \in N(x, z).$$

Fig. 18.3 State space decomposition of a Markov chain with period $d = 3$

Together with $z = x$, it follows that $d \mid (L_{x,y} + L_{y,x})$. Hence the value of $L_{x,y}$ is unique in $\{0, \dots, d-1\}$ and $L_{x,y} = -L_{y,x} \pmod{d}$. For general z , we infer that $d \mid (L_{x,y} + L_{y,z} + L_{z,x})$; hence (18.3). $\square$

Theorem 18.4 Let X be irreducible with period d . Then there exists a disjoint decomposition of the state space

$$E = \bigsqcup_{i=0}^{d-1} E_i \quad (18.5)$$

with the property

$$p(x, y) > 0 \quad \text{and} \quad x \in E_i \implies y \in E_{i+1 \pmod{d}}. \quad (18.6)$$

This decomposition is unique up to cyclic permutations.

See Fig. 18.3 for an illustration of the state space decomposition of a periodic Markov chain.

Property (18.6) says that X visits the E_i one after the other (see Fig. 18.3 or Fig. 18.2, where $d = 2$, $E_0 = \{1, 3, 5, 7\}$ and $E_1 = \{2, 4, 6, 8\}$). Somewhat more formally, we could write: If $x \in E_i$ for some i , then $\mathbf{P}_x[X_n \in E_{i+n \pmod{d}}] = 1$.

Proof Existence. Fix an arbitrary $x_0 \in E$ and let

$$E_i := \{y \in E : L_{x_0, y} = i\} \quad \text{for } i = 0, \dots, d-1.$$

Clearly, (18.5) holds. Let $i \in \{0, \dots, d-1\}$ and $x \in E_i$. If $y \in E$ with $p(x, y) > 0$, then $L_{x, y} = 1$ and hence $L_{x_0, y} = L_{x_0, x} + L_{x, y} = i + 1 \pmod{d}$.

Uniqueness. Let $(\tilde{E}_i, i = 0, \dots, d-1)$ be another decomposition that satisfies (18.5) and (18.6). Without loss of generality, assume $E_0 \cap \tilde{E}_0 \neq \emptyset$ (otherwise permute the $\tilde{E}_i$ cyclically until this holds). Fix an arbitrary $x_0 \in E_0 \cap \tilde{E}_0$. By assumption, $p(x_0, y) > 0$ now implies $y \in E_1$ and $y \in \tilde{E}_1$; hence $y \in E_1 \cap \tilde{E}_1$. Inductively, we get that $p^{nd+i}(x_0, y) > 0$ implies $y \in E_i \cap \tilde{E}_i$ (for all $n \in \mathbb{N}$ and $i = 0, \dots, d-1$).

However, since the chain is irreducible, for every $y \in E$, there exist numbers $n(y)$ and $i(y)$ such that $p^{n(y)d+i(y)}(x_0, y) > 0$; hence $y \in E_{i(y)} \cap \tilde{E}_{i(y)}$. Therefore, we have $E_i = \tilde{E}_i$ for every $i = 0, \dots, d-1$.

5.18.2 18.2 Coupling and Convergence Theorem

Our goal is to use a coupling of two discrete Markov chains that are started in different distributions μ and ν in order to show the convergence theorem for Markov chains.

In the following, let E be a countable space and let p be a stochastic matrix on E . Recall the definition of a general coupling of two probability measures from Definition 17.53.

Definition 18.5 A bivariate process $((X_n, Y_n))_{n \in \mathbb{N}_0}$ with values in $E \times E$ is called a coupling if $(X_n)_{n \in \mathbb{N}_0}$ and $(Y_n)_{n \in \mathbb{N}_0}$ are Markov chains, each with transition matrix p .

A coupling is called successful if $\mathbf{P}_{(x,y)} [\bigcup_{m \geq n} \{X_m \neq Y_m\}] \xrightarrow{n \rightarrow \infty} 0$ for all $x, y \in E$.

Of course, two independent chains form a coupling, though maybe not the most interesting one.

Example 18.6 (Independent coalescence) The most important coupling is Markov chains that run independently until they coalesce: Let X and Y be independent chains with transition matrix p until they first meet. After that, the chains run together. We call this coupling the independent coalescent. The transition matrix is

$$\bar{p}((x_1, y_1), (x_2, y_2)) = \begin{cases} p(x_1, x_2) \cdot p(y_1, y_2), & \text{if } x_1 \neq y_1, \\ p(x_1, x_2), & \text{if } x_1 = y_1, x_2 = y_2, \\ 0, & \text{if } x_1 = y_1, x_2 \neq y_2. \end{cases}$$

Denote by $\tau := \inf \{n \in \mathbb{N}_0 : X_n = Y_n\}$ the time of coalescence. We can construct the coupling using two independent chains $\tilde{X}$ and $\tilde{Y}$ by defining $X := \tilde{X}$, $\tilde{\tau} := \inf \{n \in \mathbb{N}_0 : \tilde{X}_n = \tilde{Y}_n\}$ and

$$Y_n := \begin{cases} \tilde{Y}_n, & \text{if } n < \tilde{\tau} \\ X_n, & \text{if } n \geq \tilde{\tau} \end{cases}$$

Instead of checking by a direct computation that this process (X, Y) is indeed a coupling with transition matrix $\bar{p}$, consider the construction of Markov chains from Theorem 17.17: Let $(R_n(x) : n \in \mathbb{N}_0, x \in E)$ be independent random variables with distribution $\mathbf{P}[R_n(x_1) = x_2] = p(x_1, x_2)$, and let $\tilde{R}_n((x_1, y_1)) = (R_n(x_1), R_n(y_1))$. Then $(\tilde{R}_n)_{n \in \mathbb{N}_0}$ is independent and we have

$$\mathbf{P}[\tilde{R}_n((x_1, y_1)) = (x_2, y_2)] = \bar{p}((x_1, y_1), (x_2, y_2))$$

As we saw in Theorem 17.17, by $X_{n+1} := R_n(X_n)$ and $Y_{n+1} := R_n(Y_n)$, two Markov chains X and Y are defined with transition matrix p . On the other hand, we have $(X_{n+1}, Y_{n+1}) = \tilde{R}_n((X_n, Y_n))$. Hence the bivariate process is indeed a coupling with transition matrix $\bar{p}$.

Example 18.7 Let $E = \mathbb{Z}$ and $p(x, y) = 1/3$ if $|x - y| \leq 1$ and 0 otherwise. Clearly, p is the transition matrix of an aperiodic recurrent random walk on $\mathbb{Z}$. We will show that we can obtain a successful coupling by coalescing independent chains.

Accordingly, let $\tilde{X}$ and $\tilde{Y}$ be independent random walks with transition matrix p . Then the difference chain $(Z_n)_{n \in \mathbb{N}_0} := (\tilde{X}_n - \tilde{Y}_n)_{n \in \mathbb{N}_0}$ is a symmetric random walk with finite expectation and hence recurrent. Furthermore, Z is irreducible. For any two points $x, y \in \mathbb{Z}$, we thus have

$$\mathbf{P}_{(x,y)}[\tilde{\tau} < \infty] = \mathbf{P}_{x-y}[Z_n = 0 \text{ for some } n \in \mathbb{N}_0] = 1.$$

Therefore, X and Y coalesce almost surely.

Recurrence, irreducibility and aperiodicity alone are not sufficient for the independent coalescence coupling to be successful. In Exercise 18.2.4, an example is studied that shows that spacial homogeneity cannot easily be dropped if we want to have a successful coupling. Dropping the assumption of recurrence is easier, as the following theorem shows.

Theorem 18.8 Let X be an arbitrary aperiodic and irreducible random walk on $\mathbb{Z}^d$ with transition matrix p . Then there exists a successful coupling (X, Y) .

Proof Step 1. First, consider the case where $p(0, x) = 3^{-d}$ for all $x \in \{-1, 0, 1\}^d$. The individual coordinates $X^{(1)}, \dots, X^{(d)}$ of X are independent random walks on $\mathbb{Z}$ with transition probabilities $\mathbf{P}_0[X_1^{(i)} = x_i] = 1/3$ for $x_i = -1, 0, 1$. By Example 18.7, we can construct independent successful couplings $(X^{(i)}, Y^{(i)})$, $i = 1, \dots, d$, with merging times $\tau^{(i)}$. Define $Y = (Y^{(1)}, \dots, Y^{(d)})$ and $\tau = \max\{\tau^{(1)}, \dots, \tau^{(d)}\} < \infty$. Then (X, Y) is a successful coupling and $X_n = Y_n$ for $n \geq \tau$.

Step 2. Now, consider the case where

$$\lambda := 3^d \min\{p(0, x) : x \in \{-1, 0, 1\}^d\} > 0.$$

If $\lambda = 1$, then the condition of Step 1 is fulfilled and we are done. Hence, we assume that $\lambda \in (0, 1)$. We define the transition matrix $\hat{p}$ on $\mathbb{Z}^d$ by $\hat{p}(x, y) = 3^{-d}$ for $y - x \in \{-1, 0, 1\}^d$. Note that also $\check{p} := (p - \lambda\hat{p})/(1 - \lambda)$ is the transition matrix of a random walk on $\mathbb{Z}^d$ and that

$$p = \lambda\hat{p} + (1 - \lambda)\check{p}.$$

Let $\hat{X}$ and $\check{X}$ be independent random walks with transition matrices $\hat{p}$ and $\check{p}$, respectively. Assume that $\hat{X}_0 = X_0$ and $\check{X}_0 = 0$. Furthermore, let $Z_1, Z_2, \dots$ be i.i.d.

Bernoulli random variables with parameter λ that are independent of $\hat{X}$ and $\check{X}$. Define $S_n := Z_1 + \dots + Z_n$ for $n \in \mathbb{N}$ and

$$X_n := \hat{X}_{S_n} + \check{X}_{n-S_n}.$$

That is, in each time step, a coin flip decides whether X makes a jump according to the matrix $\hat{p}$ or $\check{p}$. Hence X is a random walk with transition matrix p .

By Step 1, there exists a successful coupling $(\hat{X}, \hat{Y})$ such that $\hat{Y}$ is independent of $\check{X}$ and $Z_1, Z_2, \dots$. Consequently,

$$Y_n := \hat{Y}_{S_n} + \check{X}_{n-S_n}, \quad n \in \mathbb{N},$$

is also a random walk with transition matrix p . Since we have $S_n \rightarrow \infty$ almost surely, the coupling (X, Y) is successful.

Step 3. Finally, we consider the general situation. Since X is irreducible and aperiodic, by Lemma 18.3(ii), there exists an $N \in \mathbb{N}$, such that the N -step transition matrix fulfills

$$p^N(0, x) > 0 \quad \text{for all } x \in \{-1, 0, 1\}^d$$

Hence, the random walk $X' = (X'_n)_{n \in \mathbb{N}} := (X_{nN})_{n \in \mathbb{N}}$ fulfills the condition from Step 2. Let (X', Y') be the coupling that was constructed in Step 2 and let

$$\tau := \inf\{n \in \mathbb{N}_0 : X'_m = Y'_m \text{ for all } m \geq n\}.$$

Then Y' is a random walk with transition matrix p^N . For $n \in \mathbb{N}_0$, define $Y_{nN} := Y'_n$. It remains to close the gaps between the points $\{0, N, 2N, \dots\}$ in such a way that Y is a random walk and (X, Y) is a successful coupling.

Let $(U^{x,y,n} : x, y \in \mathbb{Z}^d, n \in \mathbb{N}_0)$ be an independent family of $(\mathbb{Z}^d)^{N-1}$ -valued random variables $U^{x,y,n} = (U_1^{x,y,n}, \dots, U_{N-1}^{x,y,n})$ such that

$$\mathbf{P}[(X_1, \dots, X_{N-1}) \in \cdot \mid X_0 = x, X_N = y] = \mathbf{P}_{U^{x,y,n}}$$

for all $x, y \in \mathbb{Z}^d$ with $p^N(x, y) > 0$ and for all $n \in \mathbb{N}_0$. We further assume that the $U^{x,y,n}$ are independent of X and Y' . For $k \in \{nN + 1, \dots, (n+1)N - 1\}$, define

$$Y_k := \begin{cases} U_{k-nN}^{Y'_n, Y'_{n+1}, n}, & \text{if } n < \tau \\ X_k, & \text{else.} \end{cases}$$

It is easy to check that Y is indeed a random walk with transition matrix p . By construction, the coupling (X, Y) is successful.

Theorem 18.9 Let X be a Markov chain on E with transition matrix p . If there exists a successful coupling, then every bounded harmonic function is constant.

Proof Let $f : E \rightarrow \mathbb{R}$ be bounded and harmonic; hence $pf = f$. Let $x, y \in E$, and let (X, Y) be a successful coupling. By Lemma 17.45, $(f(X_n))_{n \in \mathbb{N}_0}$ and $(f(Y_n))_{n \in \mathbb{N}_0}$ are martingales; hence we have

$$|f(x) - f(y)| = |\mathbf{E}_{(x,y)}[f(X_n) - f(Y_n)]| \leq 2\|f\|_\infty \mathbf{P}_{(x,y)}[X_n \neq Y_n] \xrightarrow{n \rightarrow \infty} 0.$$

□

Corollary 18.10 If X is an irreducible random walk on $\mathbb{Z}^d$, then every bounded harmonic function is constant.

This statement holds more generally if we replace $\mathbb{Z}^d$ by a locally compact Abelian group. In that form, the theorem goes back to Choquet and Deny [24], see also [143].

Proof Let p be the transition matrix of X . Let $\bar{X}$ be a Markov chain with transition matrix $\bar{p}(x, y) = \frac{1}{2}p(x, y) + \frac{1}{2}\mathbb{1}_{\{x\}}(y)$. Clearly, X and $\bar{X}$ have the same harmonic functions. Now $\bar{X}$ is an aperiodic irreducible random walk; hence, by Theorem 18.8, there is a successful coupling for all initial states. □

Theorem 18.11 Let p be the transition matrix of an irreducible, positive recurrent, aperiodic Markov chain on E . Then the independent coalescent chain is a successful coupling.

Proof Let $\tilde{X}$ and $\tilde{Y}$ be two independent Markov chains on E , each with transition matrix p . Then the bivariate Markov chain $Z := \left((\tilde{X}_n, \tilde{Y}_n) \right)_{n \in \mathbb{N}_0}$ has the transition matrix $\tilde{p}$ defined by

$$\tilde{p}((x_1, y_1), (x_2, y_2)) = p(x_1, x_2) \cdot p(y_1, y_2)$$

We first show that the matrix $\tilde{p}$ is irreducible. Only here do we need aperiodicity of p . Accordingly, fix $(x_1, y_1), (x_2, y_2) \in E \times E$. Then, by Lemma 18.2, there exists an $m_0 \in \mathbb{N}$ such that

$$p^n(x_1, x_2) > 0 \quad \text{and} \quad p^n(y_1, y_2) > 0 \quad \text{for all } n \geq m_0.$$

For $n \geq m_0$, we thus have $\tilde{p}^n((x_1, y_1), (x_2, y_2)) > 0$. Hence $\tilde{p}$ is irreducible. Now define the stopping time τ of the first entrance of $(\tilde{X}, \tilde{Y})$ into the diagonal $D := \{(x, x) : x \in E\}$ by $\tau := \inf \{n \in \mathbb{N}_0 : \tilde{X}_n = \tilde{Y}_n\}$. Let π be the invariant distribution of $\tilde{X}$. Then, clearly, the product measure $\pi \otimes \pi \in \mathcal{M}_1(E \times E)$ is an (and then the) invariant distribution of $(\tilde{X}, \tilde{Y})$. Thus $(\tilde{X}, \tilde{Y})$ is positive recurrent (hence, in particular, recurrent) by Theorem 17.51. Therefore, $\mathbf{P}_{(x,y)}[\tau < \infty] = 1$ for all initial points $(x, y) \in E \times E$ of Z . □

Theorem 18.12 Let X be a Markov chain with transition matrix p such that there exists a successful coupling. Then $\|(\mu - v)p^n\|_{TV} \xrightarrow{n \rightarrow \infty} 0$ for all $\mu, v \in \mathcal{M}_1(E)$.

If X is aperiodic and positive recurrent with invariant distribution π , then we have $\|\mathcal{L}_\mu[X_n] - \pi\|_{TV} \xrightarrow{n \rightarrow \infty} 0$ for all $\mu \in \mathcal{M}_1(E)$.

Proof It is enough to consider the case $\mu = \delta_x, v = \delta_y$ for some $x, y \in E$. Summation over x and y yields the general case. Let $(X_n, Y_n)_{n \in \mathbb{N}_0}$ be a successful coupling. Then

$$\|(\delta_x - \delta_y)p^n\|_{TV} \leq 2\mathbf{P}_{(x,y)}[X_n \neq Y_n] \xrightarrow{n \rightarrow \infty} 0.$$

We summarize the connection between aperiodicity and convergence of distributions of X in the following theorem.

Theorem 18.13 (Convergence of Markov chains) Let X be an irreducible, positive recurrent Markov chain on E with invariant distribution π . Then the following are equivalent:

- (i) X is aperiodic.
- (ii) For every $x \in E$, we have

$$\|\mathcal{L}_x[X_n] - \pi\|_{TV} \xrightarrow{n \rightarrow \infty} 0. \quad (18.7)$$

- (iii) Equation (18.7) holds for some $x \in E$.
- (iv) For every $\mu \in \mathcal{M}_1(E)$, we have $\|\mu p^n - \pi\|_{TV} \xrightarrow{n \rightarrow \infty} 0$.

Proof The implications (iv) $\implies$ (ii) $\implies$ (iii) are evident. The implication (i) $\implies$ (ii) was shown in Theorem 18.12. Hence we only show (iii) $\implies$ (i).

"(iii) $\implies$ (i)" Assume that (i) does not hold. If X has period $d \geq 2$, and if $n \in \mathbb{N}$ is not a multiple of d , then, by Theorem 17.51,

$$\|\delta_x p^n - \pi\|_{TV} \geq |p^n(x, x) - \pi(\{x\})| = \pi(\{x\}) > 0.$$

Thus, for every $x \in E$, we have $\limsup_{n \rightarrow \infty} \|\delta_x p^n - \pi\|_{TV} > 0$. Therefore, (iii) does not hold.

Exercise 18.2.1 Let d_P be the Prohorov metric (see (13.3) and Exercise 13.2.1). Show that $d_P(P, Q) \leq \sqrt{d_W(P, Q)}$ for all $P, Q \in \mathcal{M}_1(E)$. If E has a finite diameter $\text{diam}(E)$, then $d_W(P, Q) \leq (\text{diam}(E) + 1)d_P(P, Q)$ for all $P, Q \in \mathcal{M}_1(E)$.

Exercise 18.2.2 Consider the bivariate process (X, Y) that was constructed from $\tilde{X}$ and $\tilde{Y}$ in Example 18.6. Show that (X, Y) is a coupling with transition matrix $\bar{p}$.

Exercise 18.2.3 Let X be an arbitrary aperiodic irreducible recurrent random walk on $\mathbb{Z}^d$. Show that, for any two starting points, the independent coalescent coupling is successful.

Hint: Show that the difference of two independent recurrent random walks is a recurrent random walk.

Exercise 18.2.4 Let X be a Markov chain on $\mathbb{Z}^2$ with transition matrix

$$p((x_1, x_2), (y_1, y_2)) = \begin{cases} \frac{1}{4}, & \text{if } x_1 = 0, \|y - x\|_2 = 1, \\ \frac{1}{4}, & \text{if } x_1 \neq 0 \text{ and } y_1 = x_1 \pm 1, x_2 = y_2, \\ \frac{1}{2}, & \text{if } x_1 \neq 0 \text{ and } y_1 = x_1, x_2 = y_2, \\ 0, & \text{else.} \end{cases}$$

Intuitively, this is the symmetric simple random walk whose vertical transitions are all blocked away from the vertical axis. Show that X is null recurrent, irreducible and aperiodic and that independent coalescence does not give a successful coupling.

5.18.3 18.3 Markov Chain Monte Carlo Method

Let E be a finite set and let $\pi \in \mathcal{M}_1(E)$ with $\pi(x) := \pi(\{x\}) > 0$ for every $x \in E$. We consider the problem of sampling a random variable Y with distribution π on a computer. For example, this is a relevant problem if E is a very large set and if sums of the type $\sum_{x \in E} f(x)\pi(x)$ have to be approximated numerically by the estimator $n^{-1} \sum_{i=1}^n f(Y_i)$ (see Example 5.21).

Assume that our computer has a random number generator that provides realizations of i.i.d. random variables $U_1, U_2, \dots$ that are uniformly distributed on $[0, 1]$. In order for the problem to be interesting, assume also that the distribution π cannot be constructed directly too easily.

Metropolis Algorithm

We have seen already in Example 17.19 how to simulate a Markov chain on a computer. Now the idea is to construct a Markov chain X whose distribution converges to π in the long run. If we simulate such a chain and let it run long enough this should give a sample that is distributed approximately like π . The chain should be designed so that at each step, only a small number of transitions are possible in order to ensure that the procedure described in Example 17.19 works efficiently. (Of course, the chain with transition matrix $p(x, y) = \pi(y)$ converges to π , but this does not help a lot.) This method of producing (approximately) π -distributed samples and using them to estimate expected values of functions of interest is called the Markov chain Monte Carlo method or, briefly, MCMC (see [15, 112, 119]).

Let q be the transition matrix of an arbitrary irreducible Markov chain on E (with $q(x, y) = 0$ for most $y \in E$). We use this to construct the Metropolis matrix (see [70, 114]).

Definition 18.14 Define a stochastic matrix p on E by

$$p(x, y) = \begin{cases} q(x, y) \min \left(1, \frac{\pi(y)q(y, x)}{\pi(x)q(x, y)} \right), & \text{if } x \neq y, q(x, y) > 0, \\ 0, & \text{if } x \neq y, q(x, y) = 0, \\ 1 - \sum_{z \neq x} p(x, z), & \text{if } x = y. \end{cases}$$

p is called the Metropolis matrix of q and π .

Note that p is reversible (see Section 19.2); that is, for all $x, y \in E$, we have

$$\pi(x)p(x, y) = \pi(y)p(y, x). \quad (18.8)$$

In particular, π is invariant (check this!). We thus obtain the following theorem.

Theorem 18.15 Assume that q is irreducible and that for any $x, y \in E$, we have $q(x, y) > 0$ if and only if $q(y, x) > 0$. Then the Metropolis matrix p of q and π is irreducible with unique invariant distribution π . If, in addition, q is aperiodic, or q is not reversible with respect to π , then p is aperiodic.

In order to simulate a chain X that converges to π , we take a reference chain with transition matrix q and use the Metropolis algorithm: If the chain with transition matrix q proposes a transition from the present state x to state y , then we accept this proposal with probability

$$\frac{\pi(y)q(y, x)}{\pi(x)q(x, y)} \wedge 1.$$

Otherwise the chain X stays at x .

In the definition of p , the distribution π appears only in terms of the quotients $\pi(y)/\pi(x)$. In many cases of interest, these quotients are easy to compute even though $\pi(x)$ and $\pi(y)$ are not. We illustrate this with an example.

Example 18.16 (Ising model) The Ising model (pronounced like the English word "easing") is a thermodynamical (and quantum mechanical) model for ferromagnetism in crystals. It makes the following assumptions:

- Atoms are placed at the sites of a lattice Λ (for example, $\Lambda = \{0, \dots, N-1\}^2$).
- Each atom $i \in \Lambda$ has a magnetic spin $x(i) \in \{-1, 1\}$ that either points upwards ($x(i) = +1$) or downwards ($x(i) = -1$).
- Neighboring atoms interact.
- Due to thermic fluctuations, the state of the system is random and distributed according to the so-called Boltzmann distribution π on the state space $E := \{-1, 1\}^\Lambda$. A parameter of this distribution is the inverse temperature $\beta = \frac{1}{T} \geq 0$ (with T the absolute temperature).

Define the local energy that describes the energy level of a single atom at $i \in \Lambda$ as a function H^i of the state x of the whole system,

$$H^i(x) = \frac{1}{2} \sum_{j \in \Lambda: i \sim j} \mathbb{1}_{\{x(i) \neq x(j)\}}$$

Here $i \sim j$ indicates that i and j are neighbors in Λ (that is, coordinate-wise mod N , we also speak of periodic boundary conditions). The total energy (or Hamilton function) of the system in state x is the sum of the individual energies,

$$H(x) = \sum_{i \in \Lambda} H^i(x) = \sum_{i \sim j} \mathbb{1}_{\{x(i) \neq x(j)\}}$$

The Boltzmann distribution π on $E := \{-1, 1\}^\Lambda$ for the inverse temperature $\beta \geq 0$ is defined by

$$\pi(x) = Z_\beta^{-1} \exp(-\beta H(x)),$$

where the partition sum $Z_\beta = \sum_{x \in E} \exp(-\beta H(x))$ is the normalising constant such that π is a probability measure.

Macroscopically, the individual spins cannot be observed but the average magnetization can; that is, the modulus of the average of all spins,

$$m_\Lambda(\beta) = \sum_{x \in E} \pi(x) \left| \frac{1}{\#\Lambda} \sum_{i \in \Lambda} x(i) \right|$$

If we consider a very large system, then we are close to the so-called thermodynamic limit

$$m(\beta) := \lim_{\Lambda \uparrow \mathbb{Z}^d} m_\Lambda(\beta)$$

Using a contour argument, as for percolation (see [127]), one can show that (for $d \geq 2$) there exists a critical value $\beta_c = \beta_c(d) \in (0, \infty)$ such that

$$m(\beta) \begin{cases} > 0, & \text{if } \beta > \beta_c \\ = 0, & \text{if } \beta < \beta_c \end{cases} \quad (18.9)$$

See Fig. 18.4 for a computer simulation of the curve $\beta \mapsto m(\beta)$.

For a similar model, the Weiss ferromagnet, we will prove in Example 23.20 the existence of such a phase transition. In the physical literature, $T_c := 1/\beta_c$ is called the Curie temperature for spontaneous magnetization. This is a material-dependent constant (chromium bromide (CrBr) 37 Kelvin, nickel 645 K, iron 1017 K, cobalt 1404 K). Below the Curie temperature, these materials are magnetic, and above it they are not. Below the critical temperature, the magnetization increases with decreasing temperature. We will see in a computer simulation that the Ising model displays this critical temperature effect.




Fig. 18.4 Computer simulation of the magnetization curve of the Ising model on a 1000×1000 grid. The dashed vertical line indicates the critical inverse temperature

If $x \in E$, then denote by $x^{i,\sigma}$ the state in which at site i the spin is changed to $\sigma \in \{-1, +1\}$; that is,

$$x^{i,\sigma}(j) = \begin{cases} \sigma, & \text{if } j = i, \\ x(j), & \text{if } j \neq i. \end{cases}$$

Furthermore, define the state x^i in which the spin at i is reversed, $x^i := x^{i,-x(i)}$. As reference chain, we choose a chain with transition probabilities

$$q(x, y) = \begin{cases} \frac{1}{\#\Lambda}, & \text{if } y = x^i \text{ for some } i \in \Lambda, \\ 0, & \text{else.} \end{cases}$$

In words, we choose a random site $i \in \Lambda$ (uniformly on Λ) and invert the spin at that site. Clearly, q is irreducible.

The Metropolis algorithm for this chain accepts the proposal of the reference chain with probability 1 if $\pi(x^i) \geq \pi(x)$. Otherwise the proposal is accepted only with probability $\pi(x^i)/\pi(x)$. However, now

$$\begin{aligned} H(x^i) - H(x) &= \sum_{j:j \sim i} \mathbb{1}_{\{x(j) \neq -x(i)\}} - \sum_{j:j \sim i} \mathbb{1}_{\{x(j) \neq x(i)\}} \\ &= -2 \sum_{j:j \sim i} \left(\mathbb{1}_{\{x(j) \neq x(i)\}} - \frac{1}{2} \right) \end{aligned}$$




Fig. 18.5 Equilibrium states of the Ising model on an 800×800 grid (black dot $=$ spin $+1$). Left side: below the critical temperature ($\beta > \beta_c$); Right side: above the critical temperature




Fig. 18.6 Ising model (150×150 grid) below the critical temperature. Even after a long time, the computer simulation does not produce the equilibrium state but rather so-called metastable states, in which the Weiss domains are clearly visible

Hence $\pi(x^i) / \pi(x) = \exp \left(-2\beta \sum_{j \sim i} \left(\mathbb{1}_{\{x(j)=x(i)\}} - \frac{1}{2} \right) \right)$, and this expression is easy to compute as it depends only on the $2d$ neighboring spins and, in particular, does not require knowledge of the value of Z_β . We thus obtain the Metropolis transition matrix

$$p(x, y) = \begin{cases} \frac{1}{\#\Lambda} \left(1 \wedge \exp \left[2\beta \sum_{j: j \sim i} \left(\mathbb{1}_{\{x(j) \neq x(i)\}} - \frac{1}{2} \right) \right] \right), & \text{if } y = x^i \text{ for some } i \in \Lambda, \\ 1 - \sum_{i \in \Lambda} p(x, x^i), & \text{if } x = y, \\ 0, & \text{else.} \end{cases}$$

For a practical simulation use the computer's random number generator to produce independent random variables $I_1, I_2, \dots$ and $U_1, U_2, \dots$ with $I_n \sim \mathcal{U}_\Lambda$ and $U_n \sim \mathcal{U}_{[0,1]}$. Then define

$$F_n(x) = \begin{cases} x^{I_n}, & \text{if } U_n \leq \exp \left[2\beta \sum_{j: j \sim i} \left(\mathbb{1}_{\{x(j) \neq x(i)\}} - \frac{1}{2} \right) \right], \\ x, & \text{else,} \end{cases}$$

and define the Markov chain $(X_n)_{n \in \mathbb{N}}$ by $X_n = F_n(X_{n-1})$ for $n \in \mathbb{N}$. See Figs. 18.5 and 18.6 for computer simulations of equilibrium states and metastable states of the Ising model. $\square$

Gibbs Sampler

We consider a situation where, as in the above example, a state consists of many components $x = (x_i)_{i \in \Lambda} \in E$ and where Λ is a finite set. As an alternative to the Metropolis chain, we consider a different procedure to establish a Markov chain with a given invariant distribution. For the so-called Gibbs sampler or heat bath algorithm, the idea is to adapt the state locally to the stationary distribution. If x is a state and $i \in \Lambda$, then define

$$x_{-i} := \{y \in E : y(j) = x(j) \text{ for } j \neq i\}.$$

Definition 18.17 (Gibbs sampler) Let $q \in \mathcal{M}_1(\Lambda)$ with $q(i) > 0$ for every $i \in \Lambda$. The transition matrix p on E with

$$p(x, y) = \begin{cases} q_i \frac{\pi(x^{i,\sigma})}{\pi(x_{-i})}, & \text{if } y = x^{i,\sigma} \text{ for some } i \in \Lambda, \\ 0, & \text{else,} \end{cases}$$

is called a Gibbs sampler for the invariant distribution π .

Verbally, each step of the chain with transition matrix p can be described by the following instructions.

- (1) Choose a random coordinate I according to some distribution $(q_i)_{i \in \Lambda}$.
- (2) With probability $\pi(x^{I,\sigma}) / \pi(x_{-I})$, replace x by $x^{I,\sigma}$.

If $I = i$, then the new state has the distribution $\mathcal{L}(X \mid X_{-i} = x_{-i})$, where X is a random variable with distribution π . Note that, for the Gibbs sampler also it is enough to know the values of the distribution π only up to the normalising constant. (In a more general framework, the Gibbs sampler and the Metropolis algorithm can be understood as special cases of one and the same method.) For states x and y that differ only in the i th coordinate, we have (since $x_{-i} = y_{-i}$)

$$\pi(x)p(x, y) = \pi(x)q_i \frac{\pi(y)}{\pi(x_{-i})} = \pi(y)q_i \frac{\pi(x)}{\pi(y_{-i})} = \pi(y)p(y, x).$$

Thus the Gibbs sampler is a reversible Markov chain with invariant measure π . Irreducibility of the Gibbs sampler, however, has to be checked for each case.

Example 18.18 (Ising model) In the Ising model described above, we have $x_{-i} = \{x^{i,-1}, x^{i,+1}\}$. Hence, for $i \in \Lambda$ and $\sigma \in \{-1, +1\}$,

$$\begin{aligned} \pi(x^{i,\sigma} \mid x_{-i}) &= \frac{\pi(x^{i,\sigma})}{\pi(\{x^{i,-1}, x^{i,+1}\})} \\ &= \frac{e^{-\beta H(x^{i,\sigma})}}{e^{-\beta H(x^{i,-1})} + e^{-\beta H(x^{i,+1})}} \\ &= (1 + \exp[\beta(H(x^{i,\sigma}) - H(x^{i,-\sigma}))])^{-1} \\ &= \left(1 + \exp\left[2\beta \sum_{j:j \sim i} \left(\mathbb{1}_{\{x(j) \neq \sigma\}} - \frac{1}{2}\right)\right]\right)^{-1} \end{aligned}$$

The Gibbs sampler for the Ising model is thus the Markov chain $(X_n)_{n \in \mathbb{N}_0}$ with values in $E = \{-1, 1\}^\Lambda$ and with transition matrix

$$p(x, y) = \begin{cases} \frac{1}{\#\Lambda} \left(1 + \exp \left[2\beta \sum_{j:j \sim i} (\mathbb{1}_{\{x(j) \neq x(i)\}} - \frac{1}{2}) \right] \right)^{-1}, & \text{if } y = x^i \text{ for some } i \in \Lambda, \\ 0, & \text{otherwise.} \end{cases}$$

Perfect Sampling

The MCMC method as described above is based on hope: We let the chain run for a long time and hope that its distribution is close to the invariant distribution. Even if we can compute the speed of convergence (and in many cases, this is not trivial, we come back to this point in Section 18.4), the distribution will never be exactly the invariant distribution.

Although this flaw might seem inevitable in the MCMC method, it is in fact, at least theoretically, possible to use a very similar method that allows perfect sampling according to the invariant distribution π , even if we do not know anything about the speed of convergence. The idea is simple. Assume that $F_1, F_2, \dots$ are i.i.d. random maps $E \rightarrow E$ with $\mathbf{P}[F(x) = y] = p(x, y)$ for all $x, y \in E$. We have seen how to construct the Markov chain X with initial value $X_0 = x$ by defining $X_n = F_n \circ F_{n-1} \circ \dots \circ F_1(x)$.

Note that $F_1^n(x) := F_1 \circ \dots \circ F_n(x) \stackrel{\mathcal{D}}{=} F_n \circ \dots \circ F_1(x)$. Hence we have

$$\mathbf{P}[F_1^n(x) = y] \xrightarrow{n \rightarrow \infty} \pi(y) \quad \text{for every } y.$$

However, if F_1^n turns out to be a constant map (e.g., $F_1^n \equiv x^*$ for some random x^*), then we will also have $F_1^m \equiv x^*$ for all $m \geq n$. If by some clever choice of the distribution of F_n one can ensure that the stopping time $T := \inf\{n \in \mathbb{N} : F_1^n \text{ is constant}\}$ is almost surely finite (and this is always possible), then we will have $\mathbf{P}[F_1^T(x) = y] = \pi(y)$ for all $x, y \in E$. A simple algorithm for this method is the following.

- (1) Let $F \leftarrow \text{id}_E$ and $n \leftarrow 0$.
- (2) Let $n \leftarrow n + 1$. Generate F_n and let $F \leftarrow F \circ F_n$.
- (3) If F is not a constant map, then go to (2).
- (4) Output $F(*)$.

This method is called coupling from the past and goes back to Propp and Wilson [138] (see also [55, 56, 92, 137, 139, 170]). David Wilson has nice simulations and a survey of the current research on his web site <http://www.dbwilson.com/>. A nice survey on MCMC methods including coupling from the past is [66].

For a practical implementation, there are two main problems: (1) The full map F_n has to be generated and has to be composed with F . The computer time needed for this is at least of the order of the size of the space E . (2) Checking if F is constant needs computer time of the same order of magnitude. Consequently, the method can be efficiently implemented only if there is more structure. For example, assume that E is partially ordered with a smallest element $\underline{0}$ and a largest element $\underline{1}$ (like the Ising model). Further, assume that the maps F_n can be chosen to be almost surely monotone increasing. In this case, it is enough to compute at each step $F(\underline{0})$ and $F(\underline{1})$ since F is constant if the values coincide.

5.18.4 18.4 Speed of Convergence

So far we have ignored the question of the speed of convergence of the distribution $\mathbf{P}_{X_n}$ to π . For practical purposes, however, this is often the most interesting question. We do not intend to go into the details and we only briefly touch upon the topic. Without loss of generality, assume $E = \{1, \dots, N\}$. If p is reversible (Eq. (18.8)), then $f \mapsto pf$ defines a symmetric linear operator on $L^2(E, \pi)$ (exercise!). All eigenvalues $\lambda_1, \dots, \lambda_N$ (listed according to the corresponding multiplicity) are real and have modulus at most 1 since p is stochastic.

Thus we can arrange the eigenvalues by decreasing modulus $\lambda_1 = 1 \geq |\lambda_2| \geq \dots \geq |\lambda_N|$. If p is irreducible and aperiodic, then $|\lambda_2| < 1$. Let $\mu_1 = \pi, \mu_2, \dots, \mu_N$ be an orthonormal basis of left eigenvectors for the eigenvalues $\lambda_1, \dots, \lambda_N$. Then, for every $\mu = \alpha_1 \mu_1 + \dots + \alpha_N \mu_N$, we have $\mu p^n = \sum_{i=1}^N \lambda_i^n \alpha_i \mu_i$ and hence

$$\|\mu p^n - \pi\|_{TV} \leq C |\lambda_2|^n \quad (18.10)$$

for a constant C (that does not depend on μ). A similar formula holds if p is not reversible; however, with a correction term of order at most n^{V-1} . Here, V is the size of the largest Jordan block square matrix for the eigenvalue λ_2 in the Jordan canonical form of p . In particular, V is no larger than the multiplicity of the eigenvalue with second largest modulus.

The speed of convergence is thus exponential with a rate that is determined by the spectral gap $1 - |\lambda_2|$ of the second largest eigenvalue of p . In practice, for a large space E , computing the spectral gap is often extremely difficult.

Example 18.19 Let $r \in (0, 1)$ and $N \in \mathbb{N}, N \geq 2$. Further, let $E = \{0, \dots, N-1\}$. We consider the transition matrix

$$p(i, j) = \begin{cases} r, & \text{if } j = i + 1 \pmod{N} \\ 1 - r, & \text{if } j = i - 1 \pmod{N} \\ 0, & \text{else.} \end{cases}$$

p is the transition matrix of simple (asymmetric) random walk on the discrete torus $\mathbb{Z}/(N)$, which with probability r makes a jump to the right and with probability $1 - r$ makes a jump to the left. Clearly, p is irreducible, and p is aperiodic if and only if N is odd. Furthermore, the uniform distribution $\mathcal{U}_E$ is the unique invariant distribution.

Case 1: N odd. Let $\theta_k = e^{2\pi i k/N}, k = 0, \dots, N-1$, be the N th roots of unity and let the corresponding (right) eigenvectors be

$$x^k := (\theta_k^0, \theta_k^1, \dots, \theta_k^{N-1})$$

It is easy to check that p has the eigenvalues

$$\lambda_k := r\theta_k + (1-r)\bar{\theta}_k = \cos\left(\frac{2\pi k}{N}\right) + (2r-1)i \sin\left(\frac{2\pi k}{N}\right)$$

$$k = 0, \dots, N-1$$

The moduli of the eigenvalues are given by $|\lambda_k| = f(2\pi k/N)$, where

$$f(\vartheta) = \sqrt{1 - 4r(1-r)\sin(\vartheta)^2} \quad \text{for } \vartheta \in \mathbb{R}$$

Since N is odd, $|\lambda_k|$ is maximal (except for $k = 0$) for $k = \frac{N-1}{2}$ and for $k = \frac{N+1}{2}$. For these k , $|\lambda_k|$ equals $\gamma := \sqrt{1 - 4r(1-r)\sin(\pi/N)^2}$. Since all eigenvalues are different, every eigenvalue has multiplicity 1. Hence there is a constant $C < \infty$ such that

$$\|\mu p^n - \mathcal{U}_E\|_{TV} \leq C\gamma^n \quad \text{for all } n \in \mathbb{N}, \mu \in \mathcal{M}_1(E)$$

Case 2: N even. In this case, p is not aperiodic. Nevertheless, the eigenvalues and eigenvectors are of the same form as in Case 1. In order to get an aperiodic chain, for $\varepsilon > 0$, define the transition matrix

$$p_\varepsilon := (1 - \varepsilon)p + \varepsilon I,$$

where I is the unit matrix on E . p_ε describes the random walk on E that with probability ε does not move and with probability $1 - \varepsilon$ makes a jump according to p . Clearly, p_ε is irreducible and aperiodic. The eigenvalues are

$$\lambda_{\varepsilon,k} = (1 - \varepsilon)\lambda_k + \varepsilon, \quad k = 0, \dots, N-1,$$

and the corresponding eigenvectors are the x^k from above. Evidently, $\lambda_{\varepsilon,0} = 1$, and if $\varepsilon > 0$ is very small, then $\lambda_{\varepsilon,N/2} = 2\varepsilon - 1$ is the eigenvalue with the second largest modulus. For larger values of ε , we have $|\lambda_{\varepsilon,1}| > |\lambda_{\varepsilon,N/2}|$. More precisely, if we let

$$\varepsilon_0 := \frac{(1 - (2r - 1)^2) \sin(2\pi/N)^2}{(1 - (2r - 1)^2) \sin(2\pi/N)^2 + 2 \cos(2\pi/N)},$$

then the eigenvalue with the second largest modulus has modulus

$$\gamma_\varepsilon = |\lambda_{\varepsilon,N/2}| = 1 - 2\varepsilon, \quad \text{if } \varepsilon \leq \varepsilon_0,$$

or

$$\begin{aligned} \gamma_\varepsilon &= |\lambda_{\varepsilon,1}| \\ &= \sqrt{\left((1 - \varepsilon) \cos\left(\frac{2\pi}{N}\right) + \varepsilon\right)^2 + \left((1 - \varepsilon)(2r - 1) \sin\left(\frac{2\pi}{N}\right)\right)^2}, \quad \text{if } \varepsilon \geq \varepsilon_0 \end{aligned}$$

It is easy to check that $\varepsilon \mapsto |\lambda_{\varepsilon,N/2}|$ is monotone decreasing and that $\varepsilon \mapsto |\lambda_{\varepsilon,1}|$ is monotone increasing. Hence γ_ε is minimal for $\varepsilon = \varepsilon_0$.

Hence there is a $C < \infty$ with

$$\|\mu p_\varepsilon^n - \mathcal{U}_E\|_{TV} \leq C \gamma_\varepsilon^n \quad \text{for all } n \in \mathbb{N}, \mu \in \mathcal{M}_1(E),$$

and the best speed of convergence (in this class of transition matrices) can be obtained by choosing $\varepsilon = \varepsilon_0$.

Example 18.20 (Gambler's ruin) We consider the gambler's ruin problem from Example 10.19 with the probability of a gain $r \in (0, 1)$. Here the state space is $E = \{0, \dots, N\}$, and the transition matrix is of the form

$$p(i, j) = \begin{cases} r, & \text{if } j = i + 1 \in \{2, \dots, N\}, \\ 1 - r, & \text{if } j = i - 1 \in \{0, \dots, N - 2\}, \\ 1, & \text{if } j = i \in \{0, N\}, \\ 0, & \text{else.} \end{cases}$$

This transition matrix is not irreducible; rather it has two absorbing states 0 and N . In Example 10.19 (Eq. (10.5)) for the case $r \neq \frac{1}{2}$, and Example 10.16 for the case $r = \frac{1}{2}$, it was shown that, for every $\mu \in \mathcal{M}_1(E)$,

$$\mu p^n \xrightarrow{n \rightarrow \infty} (1 - m(\mu))\delta_0 + m(\mu)\delta_N. \quad (18.11)$$

Here $m(\mu) = \int p_N(x)\mu(dx)$, where the probability $p_N(x)$ that the chain, if started at x , hits N is given by

$$p_N(x) = \begin{cases} \frac{1 - \left(\frac{1-r}{r}\right)^x}{1 - \left(\frac{1-r}{r}\right)^N}, & \text{if } r \neq \frac{1}{2} \\ \frac{x}{N}, & \text{if } r = \frac{1}{2} \end{cases}$$

How quick is the convergence in (18.11)? Here also the convergence has exponential speed and the rate is determined by the second largest eigenvalue of p .

Hence we have to compute the spectrum of p . Clearly, $x^0 = (1, 0, \dots, 0)$ and $x^N = (0, \dots, 0, 1)$ are left eigenvectors for the eigenvalue 1. In order for $x = (x_0, \dots, x_N)$ to be a left eigenvector for the eigenvalue λ , the following equations have to hold:

$$\lambda x_k = r x_{k-1} + (1-r)x_{k+1} \quad \text{for } k = 2, \dots, N-2 \quad (18.12)$$

and

$$\lambda x_{N-1} = r x_{N-2} \quad (18.13)$$

If (18.12) and (18.13) hold for $x_1, \dots, x_{N-1}$, then we define $x_0 := \frac{1-p}{\lambda-1}x_1$ and $x_N := \frac{p}{\lambda-1}x_{N-1}$ and get that in fact $xp = \lambda x$. We make the ansatz

$$\lambda = (1-r)\rho(\theta + \bar{\theta}) \quad \text{and} \quad x_k = \varrho^k (\theta^k - \bar{\theta}^k) \quad \text{for } k = 1, \dots, N-1,$$

where

$$\rho = \sqrt{r/(1-r)} \quad \text{and} \quad \theta \in \mathbb{C} \setminus \{-1, +1\} \quad \text{with } |\theta| = 1.$$

Thus we have $\theta\bar{\theta} = 1$ and $(1-r)\rho^{k+1} = r\rho^{k-1}$. Therefore, for every $k = 2, \dots, N-1$,

$$\begin{aligned} \lambda x_k &= (1-r)\rho^{k+1} (\theta^k - \bar{\theta}^k) (\theta + \bar{\theta}) \\ &= (1-r)\rho^{k+1} [(\theta^{k+1} - \bar{\theta}^{k+1}) + \theta\bar{\theta} (\theta^{k-1} - \bar{\theta}^{k-1})] \\ &= r\rho^{k-1} (\theta^{k-1} - \bar{\theta}^{k-1}) + (1-r)\rho^{k+1} (\theta^{k+1} - \bar{\theta}^{k+1}) \\ &= r x_{k-1} + (1-r)x_{k+1} \end{aligned}$$

That is, (18.12) holds. The same computation with $k = N-1$ shows that (18.13) holds if and only if $\theta^N - \bar{\theta}^N = 0$; that is, if $\theta^{2N} = 1$. In all, then, for θ , we get $N-1$ different values (note that the complex conjugates of the values considered here lead to the same values λ_n),

$$\theta_n = e^{(n/N)\pi i} \quad \text{for } n = 1, \dots, N-1$$

The corresponding eigenvalues are

$$\lambda_n = \sigma \cos\left(\frac{n\pi}{N}\right) \quad \text{for } n = 1, \dots, N-1.$$

Here the variance of the individual random walk step is

$$\sigma^2 := 4r(1-r). \quad (18.14)$$

As all eigenvalues are real, the corresponding eigenvectors are given by

$$x_k^n = 2 \left(\frac{r}{1-r} \right)^{n/2} \sin\left(\frac{n\pi}{N}\right), \quad k = 1, \dots, N-1.$$

The second largest modulus of an eigenvalue is $|\lambda_n| = \sigma \cos\left(\frac{\pi}{N}\right)$ if $n = 1$ or $n = N-1$. Thus there exists a $C > 0$ such that, for every $\mu \in \mathcal{M}_1(E)$, we have

$$\mu p^n(\{1, \dots, N-1\}) \leq C \left(\sigma \cos\left(\frac{\pi}{N}\right) \right)^n \quad \text{for every } n \in \mathbb{N}.$$

In other words, the probability that the game has not finished up to the n th round is at most $C(\sigma \cos(\pi/N))^n$.

An alternative approach to the eigenvalues can be made via the roots of the characteristic polynomial

$$\chi_N(x) = \det(p - xI), \quad x \in \mathbb{R}.$$

Clearly, $\chi_1(x) = (1 - x)^2$ and $\chi_2(x) = -x(1 - x)^2$. Using Laplace's expansion formula for the determinant (elimination of rows and columns), we get the recursion

$$\chi_N(x) = -x\chi_{N-1}(x) - r(1 - r)\chi_{N-2}(x). \quad (18.15)$$

The solution is (check this!)

$$\chi_N(x) = (-1)^{N-1}(\sigma/2)^{N-1}(1 - x)^2 U_{N-1}(x/\sigma) \quad (18.16)$$

where

$$U_m(x) := \sum_{k=0}^{\lfloor m/2 \rfloor} (-1)^k \binom{m-k}{k} (2x)^{m-2k}$$

denotes the so-called m th Chebyshev polynomial of the second kind. Using de Moivre's formula, one can show that, for $x \in (-\sigma, \sigma)$,

$$\begin{aligned} \chi_N(x) &= (-1)^{N-1}(\sigma/2)^{N-1}(1 - x)^2 \frac{\sin(N \arccos(x/\sigma))}{\sqrt{1 - (x/\sigma)^2}} \\ &= (1 - x)^2 \prod_{k=1}^{N-1} \left(\sigma \cos\left(\frac{\pi k}{N}\right) - x \right) \end{aligned} \quad (18.17)$$

Apart from the double zero at 1, we get the zeros

$$\sigma \cos(\pi k/N), \quad k = 1, \dots, N - 1.$$

Exercise 18.4.1 Show (18.16).

Exercise 18.4.2 Show (18.17).

Exercise 18.4.3 Let $v(dx) = \frac{2}{\pi} \sqrt{1 - x^2} \mathbb{1}_{[-1,1]}(x) dx$. Show that the Chebyshev polynomials of the second kind are orthonormal with respect to v ; that is,

$$\int U_m U_n dv = \mathbb{1}_{\{m=n\}}$$

Exercise 18.4.4 Let

$$E = \{1, 2, 3\} \quad \text{and} \quad p = \begin{pmatrix} 1/2 & 1/3 & 1/6 \\ 1/3 & 1/3 & 1/3 \\ 0 & 3/4 & 1/4 \end{pmatrix}$$

Compute the invariant distribution and the exponential rate of convergence.

Exercise 18.4.5 Let $E = \{0, \dots, N - 1\}$, $r \in (0, 1)$ and

$$p(i, j) = \begin{cases} r, & \text{if } j = i + 1 \pmod{N} \\ 1 - r, & \text{if } j = i \pmod{N} \\ 0, & \text{else.} \end{cases}$$

Show that p is the transition matrix of an irreducible, aperiodic random walk and compute the invariant distribution and the exponential rate of convergence.

Exercise 18.4.6 Let $N \in \mathbb{N}$ and let $E = \{0, 1\}^N$ denote the N -dimensional hypercube. That is, two points $x, y \in E$ are connected by an edge if they differ in exactly one coordinate. Let p

be the transition matrix of the random walk on E that stays put with probability $\varepsilon > 0$ and that with probability $1 - \varepsilon$ makes a jump to a randomly (uniformly) chosen neighboring site.

Describe p formally and show that p is aperiodic and irreducible. Compute the invariant distribution and the exponential rate of convergence.

5.19 19 Markov Chains and Electrical Networks

We consider symmetric simple random walk on $\mathbb{Z}^2$. By Pólya's theorem (Theorem 17.39), this random walk is recurrent. However, is this still true if we remove a single edge from the lattice $\mathbb{L}^2$ of $\mathbb{Z}^2$? Intuitively, such a small local change should not make a difference for a global phenomenon such as recurrence. However, the computations used in Section 17.5 to prove recurrence are not very robust and would need a substantial improvement in order to cope with even a small change. The situation becomes even more puzzling if we restrict the random walk to, e.g., the upper half plane $\{(x, y) : x \in \mathbb{Z}, y \in \mathbb{N}_0\}$ of $\mathbb{Z}^2$. Is this random walk recurrent? Or consider bond percolation on $\mathbb{Z}^2$. Fix a parameter $p \in [0, 1]$ and independently declare any edge of $\mathbb{L}^2$ open with probability p and closed with probability $1 - p$. At a second stage, start a random walk on the random subgraph of open edges. At each step, the walker chooses one of the adjacent open edges at random (with equal probability) and traverses it. For $p > \frac{1}{2}$, there exists a unique infinite connected component of open edges (Theorem 2.47). The question that we answer at the end of this chapter is: Is a random walk on the infinite open cluster recurrent or transient?

The aim of this chapter is to establish a connection between certain Markov chains and electrical networks. This connection

- in some cases allows us to distinguish between recurrence and transience by means of easily computable quantities, and
- in other cases provides a comparison criterion that says that if a random walk on a graph is recurrent, then a random walk on any connected subgraph is recurrent. Any of the questions raised above can be answered using this comparison technique.

Some of the material of this chapter is taken from [36, 111].

5.19.1 19.1 Harmonic Functions

In this chapter, E is always a countable set and X is a discrete Markov chain on E with transition matrix p and Green function G . Recall that $F(x, y)$ is the probability of hitting y at least once when starting at x . Compare Section 17.4, in particular, Definitions 17.28 and 17.33.

Definition 19.1 Let $A \subset E$. A function $f : E \rightarrow \mathbb{R}$ is called harmonic on $E \setminus A$ if $pf(x) = \sum_{y \in E} p(x, y)f(y)$ exists and if $pf(x) = f(x)$ for all $x \in E \setminus A$.

Theorem 19.2 (Superposition principle) Assume f and g are harmonic on $E \setminus A$ and let $\alpha, \beta \in \mathbb{R}$. Then $\alpha f + \beta g$ is also harmonic on $E \setminus A$.

Proof This is trivial.

Example 19.3 Let X be transient and let $a \in E$ be a transient state (that is, a is not absorbing). Then $f(x) := G(x, a)$ is harmonic on $E \setminus \{a\}$: For $x \neq a$, we have

$$pf(x) = p \sum_{n=0}^{\infty} p^n(x, a) = \sum_{n=1}^{\infty} p^n(x, a) = G(x, a) - \mathbb{1}_{\{a\}}(x) = G(x, a)$$

Example 19.4 For $x \in E$, let $\tau_x := \inf \{n > 0 : X_n = x\}$. For $A \subset E$, let

$$\tau := \tau_A := \inf_{x \in A} \tau_x$$

be the stopping time of the first entrance to A . Assume that A is chosen so that $\mathbf{P}_x[\tau_A < \infty] = 1$ for every $x \in E$. Let $g : A \rightarrow \mathbb{R}$ be a bounded function. Define

$$f(x) := \begin{cases} g(x), & \text{if } x \in A, \\ \mathbf{E}_x[g(X_\tau)], & \text{if } x \in E \setminus A. \end{cases} \quad (19.1)$$

Then f is harmonic on $E \setminus A$. We give two proofs for this statement.

1st Proof By the Markov property, for $x \notin A$ and $y \in E$,

$$\mathbf{E}_x[g(X_\tau) \mid X_1 = y] = \begin{cases} g(y), & \text{if } y \in A \\ \mathbf{E}_y[g(X_\tau)], & \text{if } y \in E \setminus A \end{cases} = f(y).$$

Hence, for $x \in E \setminus A$,

$$\begin{aligned} f(x) &= \mathbf{E}_x[g(X_\tau)] = \sum_{y \in E} \mathbf{E}_x[g(X_\tau); X_1 = y] \\ &= \sum_{y \in E} p(x, y) \mathbf{E}_x[g(X_\tau) \mid X_1 = y] = \sum_{y \in E} p(x, y) f(y) = pf(x). \end{aligned}$$

2nd Proof We change the Markov chain by adjoining a cemetery state Δ . That is, the new state space is $\tilde{E} = E \cup \{\Delta\}$ and the transition matrix is

$$\tilde{p}(x, y) = \begin{cases} p(x, y), & \text{if } x \in E \setminus A, y \neq \Delta, \\ 0, & \text{if } x \in E \setminus A, y = \Delta, \\ 1, & \text{if } x \in A \cup \{\Delta\}, y = \Delta. \end{cases} \quad (19.2)$$

The corresponding Markov chain $\tilde{X}$ is transient, and Δ is the only absorbing state. Furthermore, we have $pf = f$ on $E \setminus A$ if and only if $\tilde{p}f = f$ on $E \setminus A$. Since $\tilde{G}(y, y) = 1$ for all $y \in A$, we have (compare Theorem 17.34)

$$\mathbf{P}_x[X_\tau = y] = \mathbf{P}_x[\tilde{\tau}_y < \infty] = \tilde{F}(x, y) = \tilde{G}(x, y) \quad \text{for all } x \in E \setminus A, y \in A.$$

Now $x \mapsto \tilde{G}(x, y)$ is harmonic on $E \setminus A$. Hence, by the superposition principle,

$$f(x) = \sum_{y \in A} \tilde{G}(x, y) g(y) \quad (19.3)$$

is harmonic on $E \setminus A$. Due to the analogy of (19.3) to Green's formula in continuous space potential theory, the function $\tilde{G}$ is called the Green function for the equation $(p - I)f = 0$ on $E \setminus A$.

Definition 19.5 The system of equations

$$\begin{aligned} (p - I)f(x) &= 0, & \text{for } x \in E \setminus A, \\ f(x) &= g(x), & \text{for } x \in A, \end{aligned} \quad (19.4)$$

is called the Dirichlet problem on $E \setminus A$ with respect to $p - I$ and with boundary value g on A .

We have shown the existence of solutions of the Dirichlet problem in Example 19.4. In order to show uniqueness (under certain conditions) we first derive the maximum principle for harmonic functions.

If $p = I$ then any function f that coincides with g on A is a solution of the Dirichlet problem. However, even in less extreme situations the solution of (19.4) may be ambiguous. This is the case if $E \setminus A$ decomposes into domains between which the chain that is stopped in A cannot change.

In order to describe formally the irreducibility condition that we have to impose, we introduce the transition matrix p_A of the chain stopped upon reaching A by

$$p_A(x, y) := \begin{cases} p(x, y), & \text{if } x \notin A, \\ \mathbb{1}_{\{x=y\}}, & \text{if } x \in A. \end{cases}$$

Further, define F_A for p_A similarly as F was defined for p . Finally, for $x \in E$ let

$$S_A^n(x) = \{y \in E : (p_A)^n(x, y) > 0\}, \quad \text{for } n \in \mathbb{N}_0$$

and

$$S_A(x) = \bigcup_{n=0}^{\infty} S_A^n(x) = \{y \in E : F_A(x, y) > 0\}$$

Theorem 19.6 (Maximum principle) Let f be a harmonic function on $E \setminus A$.

(i) If there exists an $x_0 \in E \setminus A$ such that

$$f(x_0) = \sup f(S_A(x_0)) \quad (19.5)$$

then $f(y) = f(x_0)$ for any $y \in S_A(x_0)$.

(ii) In particular, if $F_A(x, y) > 0$ for all $x, y \in E \setminus A$, and if there is an $x_0 \in E \setminus A$ such that $f(x_0) = \sup f(E)$, then $f(x_0) = f(y)$ for any $y \in E \setminus A$.

Proof (i) Let $m := \sup f(S_A(x_0))$. As f is harmonic on $E \setminus A$, we have $p_A f = f$ on E . Hence, for any $n \in \mathbb{N}$,

$$f(x_0) = (p_A)^n f(x_0) = \sum_{y \in S_A^n(x_0)} p_A^n(x_0, y) f(y) \leq m$$

with equality if and only if $f(y) = m$ for all $y \in S_A^n(x_0)$. Since (19.5) implies equality, we infer $f(x_0) = f(y)$ for all $y \in S_A(x_0)$.

(ii) This is a direct consequence of (i) since $S_A(x) \supset E \setminus A$ for any $x \in E \setminus A$.

Theorem 19.7 (Uniqueness of harmonic functions) Assume that $F(x, y) > 0$ for all $x, y \in E$. Let $A \subset E$ be such that $A \neq \emptyset$ and $E \setminus A$ is finite. Assume that f_1 and f_2 are harmonic on $E \setminus A$. If $f_1 = f_2$ on A , then $f_1 = f_2$.

In other words, the Dirichlet problem (19.4) has a unique solution given by (19.3) (or equivalently by (19.1)).

Proof By the superposition principle, $f := f_1 - f_2$ is harmonic on $E \setminus A$ with $f|_A \equiv 0$.

We will show $f \leq 0$. Then, by symmetry, also $f \geq 0$ and hence $f \equiv 0$. To this end, we assume that there exists an $x \in E$ such that $f(x) > 0$ and deduce a contradiction.

Since $f|_A \equiv 0$ and since $E \setminus A$ is finite, there is an $x_0 \in E \setminus A$ such that $f(x_0) = \max f(E) \geq f(x) > 0$.

Since $F(x, y) > 0$ for all $x, y \in E$, we have

$$n_0 := \min \{n \in \mathbb{N}_0 : p^n(x_0, y) > 0 \text{ for some } y \in A\} < \infty.$$

Clearly, we have $p^{n_0}(x_0, y) = (p_A)^{n_0}(x_0, y)$ for all $y \in A$. Hence, there exists a $y \in A$ such that $(p_A)^{n_0}(x_0, y) > 0$, i.e., $y \in S_A(x_0)$. By Theorem 19.6, this implies $f(x_0) = f(y) = 0$ contradicting the assumption.

Exercise 19.1.1 Let $\bar{p}$ be the substochastic $E \times E$ matrix that is given by $\bar{p}(x, y) = \tilde{p}(x, y)$, $x, y \in E$ (with $\tilde{p}$ as in (19.2)). Hence $\bar{p}(x, y) = p(x, y)\mathbb{1}_{x \in E \setminus A}$. Let I be the unit matrix on E .

(i) Show that $I - \bar{p}$ is invertible.

(ii) Define $\bar{G} := (I - \bar{p})^{-1}$. Show that $\bar{G}(x, y) = \tilde{G}(x, y)$ for all $x, y \in E \setminus A$ and that $\bar{G}(x, y) = \mathbb{1}_{\{x=y\}}$ if $x \in A$. In particular,

$$\bar{G}(x, y) = \mathbf{P}_x[X_{\tau_A} = y] \quad \text{for } x \in E \setminus A \text{ and } y \in A.$$

5.19.2 19.2 Reversible Markov Chains

Definition 19.8 The Markov chain X is called reversible with respect to the measure π if

$$\pi(\{x\})p(x, y) = \pi(\{y\})p(y, x) \quad \text{for all } x, y \in E. \quad (19.6)$$

Equation (19.6) is sometimes called the equation of detailed balance. X is called reversible if there is a π with respect to which X is reversible.

Remark 19.9 If X is reversible with respect to π , then π is an invariant measure for X since

$$\pi p(\{x\}) = \sum_{y \in E} \pi(\{y\})p(y, x) = \sum_{y \in E} \pi(\{x\})p(x, y) = \pi(\{x\}).$$

If X is irreducible and recurrent, then, by Remark 17.50, π is thus unique up to constant multiples.

Example 19.10 Let (E, K) be a graph with vertex set (or set of nodes) E and with edge set K (see p. 65). By $\langle x, y \rangle = \langle y, x \rangle \in K$, denote an (undirected) edge that connects x with y . Let $C := (C(x, y), x, y \in E)$ be a family of weights with $C(x, y) = C(y, x) \geq 0$ for all $x, y \in E$ and

$$C(x) := \sum_{y \in E} C(x, y) < \infty \quad \text{for all } x \in E.$$

If we define $p(x, y) := \frac{C(x, y)}{C(x)}$ for all $x, y \in E$, then X is reversible with respect to $\pi(\{x\}) = C(x)$. In fact,

$$\begin{aligned} \pi(\{x\})p(x, y) &= C(x) \frac{C(x, y)}{C(x)} = C(x, y) \\ &= C(y, x) = C(y) \frac{C(y, x)}{C(y)} = \pi(\{y\})p(y, x) \end{aligned}$$

Definition 19.11 Let (E, K) , C and X be as in Example 19.10. Then X is called a random walk on E with weights C . In particular, if $C(x, y) = \mathbb{1}_{\{\langle x, y \rangle \in K\}}$, then X is called a simple random walk on (E, K) .

Thus the random walk with weights C is reversible. However, the converse is also true.

Theorem 19.12 If X is a reversible Markov chain and if π is an invariant measure, then X is a random walk on E with weights $C(x, y) = p(x, y)\pi(\{x\})$. If X is irreducible and recurrent, then π and hence C are unique up to a factor.

Proof This is obvious.

Exercise 19.2.1 Show that p is reversible with respect to π if and only if the linear map $L^2(\pi) \rightarrow L^2(\pi)$, $f \mapsto pf$ is self-adjoint.

Exercise 19.2.2 Let $\beta > 0$, $K \in \mathbb{N}$ and $W_1, \dots, W_K \in \mathbb{R}$. Define

$$p(i, j) := \frac{1}{Z} \exp(-\beta W_j) \quad \text{for all } i, j = 1, \dots, K,$$

where $Z := \sum_{j=1}^K \exp(-\beta W_j)$ is the normalising constant.

Assume that in K (enumerated) urns there are a total of N indistinguishable balls. At each step, choose one of the N balls uniformly at random. If i is the number of the urn from which the ball is drawn, then with probability $p(i, j)$ move the ball to the urn with number j .

(i) Give a formal description of this process as a Markov chain.

(ii) Determine the invariant distribution π and show that the chain is reversible with respect to π .

5.19.3 19.3 Finite Electrical Networks

An electrical network (E, C) consists of a set E of sites (the electrical contacts) and wires between pairs of sites. The conductance of the wire that connects the points $x \in E$ and $y \in E \setminus \{x\}$ is denoted by $C(x, y) \in [0, \infty)$. If $C(x, y) = 0$, then we could just as well assume that there is no wire connecting x and y . By symmetry, we have $C(x, y) = C(y, x)$ for all x and y . Denote by

$$R(x, y) = \frac{1}{C(x, y)} \in (0, \infty]$$

the resistance of the connection $\langle x, y \rangle$. A particular case is that of a graph (E, K) where all edges have the same conductance, say 1; that is, $C(x, y) = \mathbb{1}_{\{\langle x, y \rangle \in K\}}$. The corresponding network (E, C) will be called the unit network on (E, K) .

In the remainder of this section, assume that (E, C) is a finite electrical network.

Now let $A \subset E$. At the points $x_0 \in A$, we apply the voltages $u(x_0)$ (e.g., using batteries). What is the voltage $u(x)$ at $x \in E \setminus A$?

Definition 19.13 A map $I : E \times E \rightarrow \mathbb{R}$ is called a flow on $E \setminus A$ if it is antisymmetric (that is, $I(x, y) = -I(y, x)$) and if it obeys Kirchhoff's rule:

$$\begin{aligned} I(x) &= 0, \quad \text{for } x \in E \setminus A, \\ I(A) &= 0. \end{aligned} \tag{19.7}$$

Here we denoted

$$I(x) := \sum_{y \in E} I(x, y) \quad \text{and} \quad I(A) := \sum_{x \in A} I(x).$$

Definition 19.14 A flow $I : E \times E \rightarrow \mathbb{R}$ on $E \setminus A$ is called a current flow if there exists a function $u : E \rightarrow \mathbb{R}$ with respect to which Ohm's rule is fulfilled:

$$I(x, y) = \frac{u(x) - u(y)}{R(x, y)} \quad \text{for all } x, y \in E, x \neq y.$$

In this case, $I(x, y)$ is called the flow from x to y and $u(x)$ is called the electrical potential (or voltage) at x .

Theorem 19.15 An electrical potential u in (E, C) is a harmonic function on $E \setminus A$:

$$u(x) = \sum_{y \in E} \frac{1}{C(x, y)} C(x, y) u(y) \quad \text{for all } x \in E \setminus A.$$

In particular, if the network is irreducible, an electrical potential is uniquely determined by the values on A .

Proof By Ohm's rule and Kirchhoff's rule,

$$u(x) - \sum_{y \in E} \frac{C(x, y)}{C(x)} u(y) = \sum_{y \in E} \frac{C(x, y)}{C(x)} (u(x) - u(y)) = \frac{1}{C(x)} \sum_{y \in E} I(x, y) = 0.$$

Hence u is harmonic for the stochastic matrix $p(x, y) = C(x, y)/C(x)$. The claim follows by the uniqueness theorem for harmonic functions (Theorem 19.7).

Corollary 19.16 Let X be a Markov chain on E with edge weights C . Then $u(x) = \mathbf{E}_x[u(X_{\tau_A})]$.

Assume $A = \{x_0, x_1\}$ where $x_0 \neq x_1$, and $u(x_0) = 0, u(x_1) = 1$. Then $I(x_1)$ is the total flow into the network and $-I(x_0)$ is the total flow out of the network. Kirchhoff's rule says that the flow is divergence-free and that the flows into and out of the network are equal. In other words, the net flow is $I(x_0) + I(x_1) = 0$.

Recall that, by Ohm's rule, the resistance of a wire is the quotient of the potential difference and the current flow. Hence we define the effective resistance between x_0 and x_1 as

$$R_{\text{eff}}(x_0 \leftrightarrow x_1) = \frac{u(x_1) - u(x_0)}{I(x_1)} = \frac{1}{I(x_1)} = -\frac{1}{I(x_0)}.$$

Correspondingly, the effective conductance is $C_{\text{eff}}(x_0 \leftrightarrow x_1) = R_{\text{eff}}(x_0 \leftrightarrow x_1)^{-1}$. As I and u are uniquely determined by x_0, x_1 and C , the quantities $C_{\text{eff}}(x_0 \leftrightarrow x_1)$ and $R_{\text{eff}}(x_0 \leftrightarrow x_1)$ are well-defined and can be computed from C .

Consider now two sets $A_0, A_1 \subset E$ with $A_0 \cap A_1 = \emptyset, A_0, A_1 \neq \emptyset$. Define $u(x) = 0$ for every $x \in A_0$ and $u(x) = 1$ for every $x \in A_1$. Let I be the corresponding current flow. In a manner similar to the above, we make the following definition.

Definition 19.17 We call $C_{\text{eff}}(A_0 \leftrightarrow A_1) := I(A_1)$ the effective conductance between A_0 and A_1 and $R_{\text{eff}}(A_0 \leftrightarrow A_1) := \frac{1}{I(A_1)}$ the effective resistance between A_0 and A_1 .

Example 19.18

(i) Let $E = \{0, 1, 2\}$ with $C(0, 2) = 0$, and $A_0 = \{x_0\} = \{0\}, A_1 = \{x_1\} = \{2\}$. Define $u(0) = 0$ and $u(2) = 1$. Then (with $p(x, y) = C(x, y)/C(x)$),

$$\begin{aligned} u(1) &= 1 \cdot p(1, 2) + 0 \cdot p(1, 0) \\ &= \frac{C(1, 2)}{C(1, 2) + C(1, 0)} = \frac{R(1, 0)}{R(1, 0) + R(1, 2)} \\ &= \frac{R_{\text{eff}}(1 \leftrightarrow 0)}{R_{\text{eff}}(1 \leftrightarrow 0) + R_{\text{eff}}(1 \leftrightarrow 2)} \end{aligned}$$

The total current flow is

$$I(\{2\}) = u(1)C(0, 1) = \frac{1}{R(0, 1) + R(1, 2)} = \frac{1}{\frac{1}{C(0, 1)} + \frac{1}{C(1, 2)}}.$$

Hence we have $R_{\text{eff}}(0 \leftrightarrow 2) = \frac{1}{I(\{2\})} = R(0, 1) + R(1, 2)$ and $C_{\text{eff}}(0 \leftrightarrow 2) = (C(0, 1)^{-1} + C(1, 2)^{-1})^{-1}$.

(ii) (Series connection) Let $n \in \mathbb{N}, n \geq 2$ and $E = \{0, \dots, n\}$ with conductances $C(k-1, k) > 0$ and $C(k, l) = 0$ if $|k-l| > 1$. By Kirchhoff's rule, we have $I(l, l+1) = -I(x_1)$ for any $l = 0, \dots, n-1$. By Ohm's rule, we get $u(1) = u(0) + I(x_1)R(0, 1), u(2) = u(1) + I(x_1)R(1, 2)$ and so on, yielding

$$u(k) - u(0) = I(x_1) \sum_{l=0}^{k-1} R(l, l+1)$$

Hence

$$R_{\text{eff}}(0 \leftrightarrow k) = \frac{u(k) - u(0)}{I(x_1)} = \sum_{l=0}^{k-1} R(l, l+1).$$




Fig. 19.1 Series connection of six resistors. The effective resistance is $R_{\text{eff}}(0 \leftrightarrow 6) = R(0,1) + \dots + R(5,6)$

By symmetry, we also have

$$R_{\text{eff}}(k \leftrightarrow n) = \sum_{l=k}^{n-1} R(l, l+1)$$

and thus $R_{\text{eff}}(0 \leftrightarrow n) = R_{\text{eff}}(0 \leftrightarrow k) + R_{\text{eff}}(k \leftrightarrow n)$.

Finally, for $k \in \{1, \dots, n-1\}$, we get

$$u(k) = \frac{R_{\text{eff}}(0 \leftrightarrow k)}{R_{\text{eff}}(0 \leftrightarrow k) + R_{\text{eff}}(k \leftrightarrow n)}.$$

Note that this yields the ruin probability of the corresponding Markov chain X on $\{0, \dots, n\}$,

$$\mathbf{P}_k[\tau_n < \tau_0] = u(k) = \frac{R_{\text{eff}}(0 \leftrightarrow k)}{R_{\text{eff}}(0 \leftrightarrow n)} = \frac{\sum_{l=0}^{k-1} R(l, l+1)}{\sum_{l=0}^{n-1} R(l, l+1)}. \quad (19.8)$$

(iii) (Parallel connection) Let $E = \{0, 1\}$. We extend the model a little by allowing for more than one wire to connect 0 and 1. Denote the conductances of these wires by $C_1, \dots, C_n$. Then, by Ohm's rule, the current flow along the i th wire is $I_i = \frac{u(1) - u(0)}{R_i} = \frac{1}{R_i}$. Hence the total current is $I = \sum_{i=1}^n \frac{1}{R_i}$ and thus we have

$$C_{\text{eff}}(0 \leftrightarrow 1) = \sum_{i=1}^n C_i \quad \text{and} \quad R_{\text{eff}}(0 \leftrightarrow 1) = \left(\sum_{i=1}^n \frac{1}{R_i} \right)^{-1}.$$




Fig. 19.2 Parallel connection of six resistors. The effective resistance is $R_{\text{eff}}(0 \leftrightarrow 1) = \left(R_{\text{eff}}^{-1}(0 \leftrightarrow 1) + \dots + R_{\text{eff}}^{-1}(0 \leftrightarrow 1) \right)^{-1}$

In each of the three preceding examples, the effective resistance is a monotone function of the individual resistances. This is more than just coincidence.

Theorem 19.19 (Rayleigh's monotonicity principle) Let (E, C) and (E, C') be electrical networks with $C(x, y) \geq C'(x, y)$ for all $x, y \in E$.

Then, for $A_0, A_1 \subset E$ with $A_0, A_1 \neq \emptyset$ and $A_0 \cap A_1 = \emptyset$,

$$C_{\text{eff}}(A_0 \leftrightarrow A_1) \geq C'_{\text{eff}}(A_0 \leftrightarrow A_1).$$

The remainder of this section is devoted to the proof of this theorem. We will need a theorem on conservation of energy and Thomson's principle (also called Dirichlet's principle) on the minimization of the energy dissipation.

Theorem 19.20 (Conservation of energy) Let $A = A_0 \cup A_1$, and let I be a flow on $E \setminus A$ (but not necessarily a current flow; that is, Kirchhoff's rule holds but Ohm's rule need not). Further, let $w : E \rightarrow \mathbb{R}$ be a function that is constant both on A_0 and on A_1 : $w|_{A_0} \equiv: w_0$ and $w|_{A_1} \equiv: w_1$. Then

$$(w_1 - w_0) I(A_1) = \frac{1}{2} \sum_{x, y \in E} (w(x) - w(y)) I(x, y)$$

Note that this is a discrete version of Gauß's integral theorem for (wI) . In fact, Kirchhoff's rule says that I is divergence-free on $E \setminus A$.

Proof We compute

$$\begin{aligned} & \sum_{x, y \in E} (w(x) - w(y)) I(x, y) \\ &= \sum_{x \in E} \left(w(x) \sum_{y \in E} I(x, y) \right) - \sum_{y \in E} \left(w(y) \sum_{x \in E} I(x, y) \right) \\ &= \sum_{x \in A} \left(w(x) \sum_{y \in E} I(x, y) \right) - \sum_{y \in A} \left(w(y) \sum_{x \in E} I(x, y) \right) \\ &= w_0 I(A_0) + w_1 I(A_1) - w_0 (-I(A_0)) - w_1 (-I(A_1)) \\ &= 2(w_1 - w_0) I(A_1) \end{aligned}$$

Definition 19.21 Let I be a flow on $E \setminus A$. Denote by

$$L_I := L_I^C := \frac{1}{2} \sum_{x, y \in E} I(x, y)^2 R(x, y)$$

the energy dissipation of I in the network (E, C) .

Theorem 19.22 (Thomson's (or Dirichlet's) principle of minimization of energy dissipation) Let I and J be unit flows from A_1 to A_0 (that is, $I(A_1) = J(A_1) = 1$). Assume in addition that I is a current flow (that is, it satisfies Ohm's rule with some potential u that is constant both on A_0 and on A_1). Then

$$L_I \leq L_J$$

with equality if and only if $I = J$. In particular, the unit current flow is uniquely determined.

Proof Let $D = J - I \neq 0$ be the difference of the flows. Then clearly $D(A_0) = D(A_1) = 0$. We infer

$$\begin{aligned}
& \sum_{x,y \in E} J(x,y)^2 R(x,y) \\
&= \sum_{x,y \in E} (I(x,y) + D(x,y))^2 R(x,y) \\
&= \sum_{x,y \in E} (I(x,y)^2 + D(x,y)^2) R(x,y) + 2 \sum_{x,y \in E} I(x,y) D(x,y) R(x,y) \\
&= \sum_{x,y \in E} (I(x,y)^2 + D(x,y)^2) R(x,y) + 2 \sum_{x,y \in E} (u(x) - u(y)) D(x,y)
\end{aligned}$$

By the principle of conservation of energy, the last term equals

$$2 \sum_{x,y \in E} (u(x) - u(y)) D(x,y) = 4D(A_1)(u_1 - u_0) = 0.$$

Therefore (since $D \not\equiv 0$),

$$L_J = L_I + \frac{1}{2} \sum_{x,y \in E} D(x,y)^2 R(x,y) > L_I.$$

□

Proof (Rayleigh's monotonicity principle, Theorem 19.19) Let I and I' be the unit current flows from A_1 to A_0 with respect to C and C' , respectively. By Thomson's principle, the principle of conservation of energy and the assumption $R(x,y) \leq R'(x,y)$ for all $x,y \in E$, we have

$$\begin{aligned}
R_{\text{eff}}(A_0 \leftrightarrow A_1) &= \frac{u(1) - u(0)}{I(A_1)} = u(1) - u(0) \\
&= \frac{1}{2} \sum_{x,y \in E} I(x,y)^2 R(x,y) \\
&\leq \frac{1}{2} \sum_{x,y \in E} I'(x,y)^2 R(x,y) \leq \frac{1}{2} \sum_{x,y \in E} I'(x,y)^2 R'(x,y) \\
&= u'(1) - u'(0) = R'_{\text{eff}}(A_0 \leftrightarrow A_1)
\end{aligned}$$

□

5.19.4 19.4 Recurrence and Transience

We consider the situation where E is countable and $A_1 = \{x_1\}$ for some $x_1 \in E$. Let X be a random walk on E with weights $C = (C(x,y), x,y \in E)$ and hence with transition probabilities $p(x,y) = C(x,y)/C(x)$ (compare Definition 19.11).

The main goal of this section is to express the probability $1 - F(x_1, x_1)$ that the random walk never returns to x_1 in terms of effective resistances in the network. In order to apply the results on finite electrical networks from the last section, we henceforth assume that $A_0 \subset E$ is such that $E \setminus A_0$ is finite. We will obtain $1 - F(x_1, x_1)$ as the limit of the probability that a random walk started at x_1 hits A_0 before returning to x_1 as $A_0 \downarrow \emptyset$.

Let $u = u_{x_1, A_0}$ be the unique potential function on E with $u(x_1) = 1$ and $u(x) = 0$ for any $x \in A_0$. By Theorem 19.7, u is harmonic and can be written as

$$\begin{aligned}
u_{x_1, A_0}(x) &= \mathbf{E}_x \left[\mathbb{1}_{\{X_{\tau_{A_0} \cup \{x_1\}} = x_1\}} \right] \\
&= \mathbf{P}_x [\tau_{x_1} < \tau_{A_0}] \quad \text{for every } x \in E \setminus (A_0 \cup \{x_1\}).
\end{aligned}$$

Hence the current flow I with respect to u satisfies

$$\begin{aligned}
-I(A_0) &= I(x_1) = \sum_{x \in E} I(x_1, x) = \sum_{x \in E} (u(x_1) - u(x)) C(x_1, x) \\
&= C(x_1) \sum_{x \in E} (1 - u(x)) p(x_1, x) \\
&= C(x_1) \left(\sum_{x \notin A_0 \cup \{x_1\}} p(x_1, x) \mathbf{P}_x[\tau_{A_0} < \tau_{x_1}] + \sum_{x \in A_0} p(x_1, x) \right) \\
&= C(x_1) \mathbf{P}_{x_1}[\tau_{A_0} < \tau_{x_1}].
\end{aligned}$$

Therefore,

$$\begin{aligned}
p_F(x_1, A_0) &:= \mathbf{P}_{x_1}[\tau_{A_0} < \tau_{x_1}] \\
&= \frac{C_{\text{eff}}(x_1 \leftrightarrow A_0)}{C(x_1)} = \frac{1}{C(x_1)} \frac{1}{R_{\text{eff}}(x_1 \leftrightarrow A_0)}.
\end{aligned} \tag{19.9}$$

Definition 19.23 We denote the escape probability

$$p_F(x_1) = \mathbf{P}_{x_1}[\tau_{x_1} = \infty] = 1 - F(x_1, x_1).$$

We denote the effective conductance from x_1 to ∞ by

$$C_{\text{eff}}(x_1 \leftrightarrow \infty) := C(x_1) \inf \{p_F(x_1, A_0) : A_0 \subset E \text{ with } |E \setminus A_0| < \infty, A_0 \notin x_1\}.$$

Lemma 19.24 For any decreasing sequence $A_0^n \downarrow \emptyset$ such that $|E \setminus A_0^n| < \infty$ and $x_1 \notin A_0^n$ for all $n \in \mathbb{N}$, we have

$$C_{\text{eff}}(x_1 \leftrightarrow \infty) = \lim_{n \rightarrow \infty} C_{\text{eff}}(x_1 \leftrightarrow A_0^n).$$

Proof This is obvious since

$$C_{\text{eff}}(x_1 \leftrightarrow \infty) = C(x_1) \inf \{p_F(x_1, A_0) : |E \setminus A_0| < \infty, A_0 \notin x_1\} \tag{19.10}$$

and since $p_F(x_1, A_0)$ is monotone decreasing in A_0 . $\square$

Theorem 19.25 We have

$$p_F(x_1) = \frac{1}{C(x_1)} C_{\text{eff}}(x_1 \leftrightarrow \infty). \tag{19.11}$$

In particular,

$$x_1 \text{ is recurrent} \iff C_{\text{eff}}(x_1 \leftrightarrow \infty) = 0 \iff R_{\text{eff}}(x_1 \leftrightarrow \infty) = \infty.$$

Proof Let $A_0^n \downarrow \emptyset$ be a decreasing sequence such that $|E \setminus A_0^n| < \infty$ and $x_1 \notin A_0^n$ for all $n \in \mathbb{N}$. Define $F_n := \{\tau_{A_0^n} < \tau_{x_1}\}$. For every $M \in \mathbb{N}$, we have

$$\mathbf{P}_{x_1}[\tau_{A_0^n} \leq M] \leq \sum_{k=0}^M \mathbf{P}_{x_1}[X_k \in A_0^n] \xrightarrow{n \rightarrow \infty} 0.$$

Hence $\tau_{A_0^n} \uparrow \infty$ almost surely, and thus $F_n \downarrow \{\tau_{x_1} = \infty\}$ (up to a null set). We conclude

$$\frac{1}{C(x_1)} C_{\text{eff}}(x_1 \leftrightarrow \infty) = \lim_{n \rightarrow \infty} \mathbf{P}_{x_1}[F_n] = \mathbf{P}_{x_1}[\tau_{x_1} = \infty] = p_F(x_1).$$

□

Example 19.26 Symmetric simple random walk on $E = \mathbb{Z}$ is recurrent. Here $C(x, y) = \mathbb{1}_{\{|x-y|=1\}}$. The effective resistance from 0 to ∞ can be computed by the formulas for parallel and sequence connections,

$$R_{\text{eff}}(0 \leftrightarrow \infty) = \frac{1}{2} \sum_{i=0}^{\infty} R(i, i+1) = \infty.$$

□

Example 19.27 Asymmetric simple random walk on $E = \mathbb{Z}$ with $p(x, x+1) = p \in (\frac{1}{2}, 1)$, $p(x, x-1) = 1-p$ is transient. Here one choice (and thus up to multiples the unique choice) for the conductances is

$$C(x, x+1) = \left(\frac{p}{1-p} \right)^x \quad \text{for } x \in \mathbb{Z},$$

and $C(x, y) = 0$ if $|x-y| > 1$. By the monotonicity principle, the effective resistance from 0 to ∞ can be bounded by

$$\begin{aligned} R_{\text{eff}}(0 \leftrightarrow \infty) &= \lim_{n \rightarrow \infty} R_{\text{eff}}(0 \leftrightarrow \{-n, n\}) \\ &\leq \lim_{n \rightarrow \infty} R_{\text{eff}}(0 \leftrightarrow n) \\ &= \sum_{n=0}^{\infty} \left(\frac{1-p}{p} \right)^n = \frac{p}{2p-1} < \infty. \end{aligned}$$

Example 19.28 Symmetric simple random walk on $E = \mathbb{Z}^2$ is recurrent. Here again $C(x, y) = \mathbb{1}_{\{|x-y|=1\}}$. Let $B_n = \{-n, \dots, n\}^2$ and $\partial B_n = B_n \setminus B_{n-1}$. We construct a network C' with greater conductances by adding ring-shaped superconductors along ∂B . (See Figs. 19.3 and 19.4 for illustrations.) That is, we replace $C(x, y)$ by

$$C'(x, y) = \begin{cases} \infty, & \text{if } x, y \in \partial B_n \text{ for some } n \in \mathbb{N}, \\ C(x, y), & \text{else.} \end{cases}$$




Fig. 19.3 Electrical network on $\mathbb{Z}^2$. The bold lines are superconductors. The n th and the $(n + 1)$ th superconductors are connected by $4(2n + 1)$ edges




Fig. 19.4 Effective network after adding superconductors to $\mathbb{Z}^2$. The ring-shaped superconductors have melted down to single points

Then $R'_{\text{eff}}(B_n \leftrightarrow B_n^c) = \frac{1}{4(2n+1)}$ (note that there are $4(2n + 1)$ edges that connect B_n with B_n^c), and thus

$$R'_{\text{eff}}(0 \leftrightarrow \infty) = \sum_{n=0}^{\infty} \frac{1}{4(2n+1)} = \infty$$

By the monotonicity principle, we thus have $R_{\text{eff}}(0 \leftrightarrow \infty) \geq R'_{\text{eff}}(0 \leftrightarrow \infty) = \infty$.

Example 19.29 Let (E, K) be an arbitrary connected subgraph of the square lattice $(\mathbb{Z}^2, \mathbb{L}^2)$. Then simple random walk on (E, K) (see Definition 19.11) is recurrent. Indeed, by the monotonicity principle, we have

$$R_{\text{eff}}^{(E,K)}(0 \leftrightarrow \infty) \geq R_{\text{eff}}^{(\mathbb{Z}^2, \mathbb{L}^2)}(0 \leftrightarrow \infty) = \infty.$$

We formulate the method used in the foregoing examples as a theorem.

Theorem 19.30 Let C and C' be edge weights on E with $C'(x, y) \leq C(x, y)$ for all $x, y \in E$. If the Markov chain X with weights C is recurrent, then the Markov chain X' with weights C' is also recurrent.

In particular, consider a graph (E, K) and a subgraph (E', K') . If simple random walk on (E, K) is recurrent, then so is simple random walk on (E', K') .

Proof This follows from Theorem 19.25 and Rayleigh's monotonicity principle (Theorem 19.19).

Example 19.31 Symmetric simple random walk on $\mathbb{Z}^3$ is transient. In order to prove this, we construct a subgraph for which we can compute $R'_{\text{eff}}(0 \leftrightarrow \infty) < \infty$.

Sketch. We consider the set of all infinite paths starting at 0 and that

- begin by taking one step in the x -direction, the y -direction or the z -direction,
- continue by choosing a possibly different direction x, y or z and make two steps in that direction, and
- at the n th stage choose a direction x, y or z and take 2^{n+1} steps in that direction.

For example, by $xyyxxxxzzzzzzzz...$ we denote the path that starts with one step in direction x , then chooses y , then x , then z and so on. Note that after two paths follow different directions for the first time, they will not have any common edge again, though some of the nodes can be visited by both paths.

Consider the electrical network with unit resistors. Apply a voltage of 1 at the origin and 0 at the endpoints of the paths at the n th stage. By symmetry, the potential at a given node depends only on the distance (length of the shortest path) from the origin. We thus obtain an equivalent network if we replace multiply used nodes by multiple nodes (see Fig. 19.5). Thus we obtain a tree-shaped network: For any $n \in \mathbb{N}_0$, after 2^n steps each path splits into three (see Fig. 19.6). The 3^n paths leading from the nodes of the n th generation to those of the $(n+1)$ th generation are disjoint




Fig. 19.5 Scheme of the first three steps (two stages) of the graph from Example 19.31. The left figure shows the actual edges where, e.g., xyy indicates that the first step is in direction x , the second step is in direction y and then the third step is necessarily also in direction y . In the right figure, the nodes at the ends of $xz/zx, xy/yx$ and yz/zy are split into two nodes and then connected by a superconductor (bold line). If we remove the superconductors from the network, we end up with the network of Fig. 19.6 whose effective resistance $R'_{\text{eff}}(0 \leftrightarrow \infty)$ is not smaller than that of $\mathbb{Z}^3$. (If at the root we apply a voltage of 1 and at the points to the right the voltage 0, then by symmetry no current flows through the superconductors. Thus, in fact, the network is equivalent to that in Fig. 19.6)

paths, each of length 2^{n-1} . If $B(n)$ denotes the set of points up to the n th generation, then

$$R'_{\text{eff}}(0 \leftrightarrow B(n+1)^c) = \sum_{k=0}^{n-1} R'_{\text{eff}}(B(k) \leftrightarrow B(k)^c) = \sum_{k=0}^{n-1} 2^k 3^{-k}.$$

Therefore, $R'_{\text{eff}}(0 \leftrightarrow \infty) = \frac{1}{3} \sum_{k=0}^{\infty} \left(\frac{2}{3}\right)^k = 1 < \infty$. On this tree, random walk is transient. Hence, by Theorem 19.30, random walk on $\mathbb{Z}^3$ is also transient. $\square$

Exercise 19.4.1 Consider the electrical network on $\mathbb{Z}^d$ with unit resistors between neighboring points. Let X be a symmetric simple random walk on $\mathbb{Z}^d$. Finally, fix two arbitrary neighboring points $x_0, x_1 \in \mathbb{Z}^d$. Show the following:

- (i) The effective conductance between x_0 and x_1 is $C_{\text{eff}}(x_0 \leftrightarrow x_1) = d$.
- (ii) If $d \leq 2$, then $\mathbf{P}_{x_0}[\tau_{x_1} < \tau_{x_0}] = \frac{1}{2}$.
- (iii) If $d \geq 3$, then $\mathbf{P}_{x_0}[\tau_{x_1} < \tau_{x_0} \mid \tau_{x_0} \wedge \tau_{x_1} < \infty] = \frac{1}{2}$.




Fig. 19.6 A tree as a subgraph of $\mathbb{Z}^3$ on which random walk is still transient

5.19.5 19.5 Network Reduction

Example 19.32 Consider a random walk on the graph in Fig. 19.7 that starts at x and at each step jumps to one of its neighbors at random with equal probability. What is the probability P that this Markov chain visits 1 before it visits 0?

We can regard the graph as an electrical network with unit resistors at each edge, voltage 0 at 0 and voltage 1 at 1. Then P equals the voltage at point x :

$$P = u(x).$$

In order to compute $u(x)$, we replace the network step by step by simpler networks such that the effective resistances between 0, 1, and x remain unchanged. Hence in each step the voltage $u(x)$ at point x does not change. $\square$




Fig. 19.7 Initial situation




Fig. 19.8 Reduced network with three nodes

Reduced Network

Assume that we have already reduced the network to a network with the three points 0, 1 and x and with resistors between these points $R'(0, 1)$, $R'(0, x)$ and $R'(1, x)$. See Fig. 19.8.

Clearly, we have

$$P = u(x) = \frac{R'(0, x)}{R'(0, x) + R'(1, x)}. \quad (19.12)$$

If we knew the effective resistances $R_{\text{eff}}(0 \leftrightarrow x)$, $R_{\text{eff}}(1 \leftrightarrow x)$ and $R_{\text{eff}}(0 \leftrightarrow 1)$, we could avoid the hassle of reducing the network and we could compute $u(x)$ directly. In order to derive the formula for $u(x)$, we make the following observations. In the reduced network, the effective resistances are easy to compute: If $\{a, b, c\} = \{0, 1, x\}$, then

$$R_{\text{eff}}(a \leftrightarrow b) = \left(\frac{1}{R'(a, b)} + \frac{1}{R'(a, c) + R'(b, c)} \right)^{-1}. \quad (19.13)$$

Solving these three equations for $R'(0, 1)$, $R'(0, x)$ and $R'(1, x)$ and plugging the values into (19.12) yields

$$P = u(x) = \frac{R_{\text{eff}}(0 \leftrightarrow 1) + R_{\text{eff}}(0 \leftrightarrow x) - R_{\text{eff}}(x \leftrightarrow 1)}{2R_{\text{eff}}(0 \leftrightarrow 1)}. \quad (19.14)$$

In particular, in the case $R'(0, 1) = \infty$ (or equivalently $R_{\text{eff}}(0 \leftrightarrow 1) = R_{\text{eff}}(0 \leftrightarrow x) + R_{\text{eff}}(x \leftrightarrow 1)$), we have $R_{\text{eff}}(0 \leftrightarrow x) = R'(0, x)$ and $R_{\text{eff}}(1 \leftrightarrow x) = R'(1, x)$, hence

$$u(x) = \frac{R_{\text{eff}}(0 \leftrightarrow x)}{R_{\text{eff}}(0 \leftrightarrow x) + R_{\text{eff}}(x \leftrightarrow 1)}. \quad (19.15)$$

Since we always have $u(x) \in [0, 1]$, rearranging the terms yields (again in the general situation)

$$R_{\text{eff}}(1 \leftrightarrow x) \leq R_{\text{eff}}(0 \leftrightarrow 1) + R_{\text{eff}}(0 \leftrightarrow x). \quad (19.16)$$

This is the triangle inequality for the effective resistances and it shows that the effective resistance is a metric in any electrical network.

Step-by-step Reduction of the Network

Having seen how to compute $u(x)$ from the effective resistances, we now turn to the systematic computation of these effective resistances. Later we will come back to the introductory example and make the computations explicit.

There are four elementary transformations for the reduction of an electrical network:

1. Deletion of loops. The three points on the very right of the graph form a loop that can be deleted from the network without changing any of the remaining voltages. In particular, any edge that directly connects 0 to 1 can be deleted.
2. Joining serial edges. If two (or more) edges are in a row such that the nodes along them do not have any further adjacent edges, this sequence of edges can be substituted by a single edge whose resistance is the sum of the resistances of the single edges (see Fig. 19.1).
3. Joining parallel edges. Two (or more) edges with resistances $R_1, \dots, R_n$ that connect the same two nodes can be replaced by a single edge with resistance $R = (R_1^{-1} + \dots + R_n^{-1})^{-1}$ (see Fig. 19.2).
4. Star-triangle transformation. (See Exercise 19.5.1.) The star-shaped part of a network (left in Fig. 19.9) is equivalent to the triangle-shaped part (right in Fig. 19.9) if the resistances $R_1, R_2, R_3, \tilde{R}_1, \tilde{R}_2, \tilde{R}_3$ satisfy the condition

$$R_i \tilde{R}_i = \delta \quad \text{for any } i = 1, 2, 3, \quad (19.17)$$

where

$$\delta = R_1 R_2 R_3 (R_1^{-1} + R_2^{-1} + R_3^{-1}) = \frac{\tilde{R}_1 \tilde{R}_2 \tilde{R}_3}{\tilde{R}_1 + \tilde{R}_2 + \tilde{R}_3}.$$




Fig. 19.9 Star-triangle transformation

Application to Example 19.32

With the four transformations at hand, we solve the problem of Example 19.32. Assume that initially all edges have resistance 1. In the figures we label each edge with its resistance if it differs (in the course of the reduction) from 1.

Step 1. Delete the loop at the right-hand side (left in Fig. 19.10).

Step 2. Replace the series on top, bottom and right by edges with resistance 2 (right in Fig. 19.10).

Step 3. Use the star-triangle transformation to remove the lower left node (left in Fig. 19.11). Here $R_1 = 1, R_2 = 2, R_3 = 1, \delta = 5, \tilde{R}_1 = \delta/R_1 = 5, \tilde{R}_2 = \delta/R_2 = 5/2$ and $\tilde{R}_3 = \delta/R_3 = 5$.

Step 4. Replace the parallel edges with resistances $R_1 = 5$ and $R_2 = 1$ by one edge with $R = \left(\frac{1}{5} + 1\right)^{-1} = \frac{5}{6}$ (right in Fig. 19.11).

Step 5. Use the star-triangle transformation to remove the lower right node (left in Fig. 19.12). Here $R_1 = 5, R_2 = 2, R_3 = \frac{5}{6}, \delta = 95/6, \tilde{R}_1 = \delta/R_1 = 19/6, \tilde{R}_2 = \delta/R_2 = 95/12$ and $\tilde{R}_3 = \delta/R_3 = 19$.




Fig. 19.10 Steps 1 and 2




Fig. 19.11 Steps 3 and 4




Fig. 19.12 Steps 5 and 6




Fig. 19.13 Steps 7 and 8

Step 6. Replace the parallel edges by edges with resistances $(\frac{12}{95} + \frac{2}{5})^{-1} = \frac{19}{10}$ and $(\frac{6}{19} + 1)^{-1} = \frac{19}{25}$, respectively (right in Fig. 19.12).

Step 7. Use the star-triangle transformation to remove the lower right node (left in Fig. 19.13). Here $R_1 = \frac{19}{10}, R_2 = \frac{19}{25}, R_3 = 1, \delta = \frac{513}{125}, \tilde{R}_1 = \delta/R_1 = \frac{54}{25}, \tilde{R}_2 = \delta/R_2 = \frac{27}{5}$ and $\tilde{R}_3 = \delta/R_3 = \frac{513}{125}$.

Step 8. Replace the three pairs of parallel edges by single edges with resistances $(\frac{5}{27} + 1)^{-1} = \frac{27}{32}, (\frac{25}{54} + \frac{1}{2})^{-1} = \frac{27}{26}$ and $(\frac{1}{19} + \frac{125}{513})^{-1} = \frac{27}{8}$, respectively.

In the reduced network, we have the resistances

$$R'(0, x) = \frac{27}{32} \quad \text{and} \quad R'(x, 1) = \frac{27}{26}.$$

Using (19.12), the probability that the random walk visits 1 before 0 is

$$P = \frac{\frac{27}{32}}{\frac{27}{32} + \frac{27}{26}} = \frac{13}{29}.$$

Using the values of $R'(0, x), R'(1, x)$ and $R'(0, 1)$ and Eq. (19.13), we compute the effective resistances in the reduced network (and hence in the original network):

$$\begin{aligned} R_{\text{eff}}(0 \leftrightarrow x) &= \left(\frac{32}{27} + \frac{1}{\frac{27}{8} + \frac{27}{26}} \right)^{-1} = \frac{17}{24} \\ R_{\text{eff}}(1 \leftrightarrow x) &= \left(\frac{26}{27} + \frac{1}{\frac{27}{32} + \frac{27}{8}} \right)^{-1} = \frac{5}{6} \\ R_{\text{eff}}(0 \leftrightarrow 1) &= \left(\frac{8}{27} + \frac{1}{\frac{27}{26} + \frac{27}{32}} \right)^{-1} = \frac{29}{24}. \end{aligned}$$

Using (19.14) we can use the values to compute $u(x)$:

$$P = u(x) = \frac{\frac{29}{24} + \frac{17}{24} - \frac{5}{6}}{2 \cdot \frac{29}{24}} = \frac{13}{29}.$$

Clearly, the latter computation is more complicated than using the resistances R' from the reduced network directly. However, it has the advantage that it can be performed without going through all the network reduction steps if, for some reason, we know the effective resistances already. For example, we could buy resistors in an electronic market, solder the network and measure the resistances with a multimeter.

Alternative Solution

A different approach to solving the problem of Example 19.32 is to use linear algebra instead of network reduction. It is a matter of taste as to which solution is preferable. First generate the transition matrix p of the Markov chain. To this end, enumerate the nodes of the graph from 1 to 12 as in Fig. 19.14.




Fig. 19.14 Graph with enumerated nodes

The chain starts at 2, and we want to compute the probability that it visits 3 before 5.

Generate the matrix $\bar{p}$ of the chain that is killed at 3 and at 5 and compute $\bar{G} = (I - \bar{p})^{-1}$. By Exercise 19.1.1 (with $A = \{3, 5\}$, $x = 2$ and $y = 3$), the probability of visiting 3 before 5 is $P = \bar{G}(2, 3) = \frac{13}{29}$.

$$\bar{p} := \begin{pmatrix} 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{1}{3} & 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & 0 & 0 & \frac{1}{4} & \frac{1}{4} & 0 & 0 \\ 0 & 0 & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & 0 & 0 & \frac{1}{4} & \frac{1}{4} & 0 \\ 0 & 0 & 0 & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} & 0 & 0 & 0 & 0 & 0 & \frac{1}{3} \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} & 0 & 0 & 0 & 0 & \frac{1}{3} \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 \end{pmatrix},$$

$$\bar{G} := (I - \bar{p})^{-1}$$

$$= \begin{pmatrix} \frac{143}{116} & \frac{81}{116} & \frac{21}{29} & \frac{3}{58} & \frac{8}{29} & \frac{19}{58} & \frac{3}{29} & \frac{3}{58} & \frac{15}{116} & \frac{9}{58} & \frac{3}{58} & \frac{11}{116} \\ \frac{116}{27} & \frac{81}{58} & \frac{29}{58} & \frac{58}{29} & \frac{29}{29} & \frac{58}{29} & \frac{29}{29} & \frac{58}{29} & \frac{116}{58} & \frac{58}{29} & \frac{58}{29} & \frac{116}{58} \\ \frac{58}{29} & \frac{58}{29} & \frac{29}{29} & \frac{58}{29} & \frac{29}{29} & \frac{58}{29} & \frac{29}{29} & \frac{58}{29} & \frac{116}{58} & \frac{58}{29} & \frac{58}{29} & \frac{116}{58} \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{3}{58} & \frac{9}{58} & \frac{24}{58} & \frac{165}{58} & \frac{5}{29} & \frac{15}{29} & \frac{78}{29} & \frac{68}{29} & \frac{21}{58} & \frac{30}{29} & \frac{107}{58} & \frac{27}{58} \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{19}{116} & \frac{57}{116} & \frac{18}{29} & \frac{15}{58} & \frac{11}{29} & \frac{95}{58} & \frac{15}{29} & \frac{15}{58} & \frac{75}{116} & \frac{45}{58} & \frac{15}{58} & \frac{55}{116} \\ \frac{116}{3} & \frac{116}{9} & \frac{29}{24} & \frac{58}{39} & \frac{29}{5} & \frac{58}{15} & \frac{29}{78} & \frac{58}{39} & \frac{116}{21} & \frac{58}{30} & \frac{58}{39} & \frac{116}{27} \\ \frac{58}{3} & \frac{58}{9} & \frac{29}{24} & \frac{29}{68} & \frac{29}{5} & \frac{29}{15} & \frac{29}{78} & \frac{29}{97} & \frac{58}{21} & \frac{29}{30} & \frac{29}{68} & \frac{58}{27} \\ \frac{58}{5} & \frac{58}{15} & \frac{29}{11} & \frac{29}{18} & \frac{29}{25} & \frac{29}{14} & \frac{29}{29} & \frac{58}{93} & \frac{29}{21} & \frac{29}{7} & \frac{29}{45} & \frac{58}{45} \\ \frac{58}{3} & \frac{58}{9} & \frac{29}{19} & \frac{29}{20} & \frac{29}{10} & \frac{29}{30} & \frac{29}{40} & \frac{29}{20} & \frac{58}{21} & \frac{29}{60} & \frac{29}{20} & \frac{58}{27} \\ \frac{29}{3} & \frac{29}{9} & \frac{24}{29} & \frac{107}{29} & \frac{5}{29} & \frac{15}{29} & \frac{78}{29} & \frac{68}{29} & \frac{21}{30} & \frac{30}{165} & \frac{29}{27} & \frac{29}{27} \\ \frac{58}{11} & \frac{58}{33} & \frac{29}{15} & \frac{58}{27} & \frac{29}{14} & \frac{29}{55} & \frac{29}{27} & \frac{29}{135} & \frac{58}{81} & \frac{29}{27} & \frac{58}{215} & \frac{58}{215} \end{pmatrix}.$$

Exercise 19.5.1 Show the validity of the star-triangle transformation.

Exercise 19.5.2 Consider a random walk on the honeycomb graph shown below. Show that if the walk starts at x , then the probability of visiting 1 before 0 is $\frac{8}{17}$ using

- (i) the method of network reduction, and
- (ii) the method of matrix inversion.




Fig. 19.15 Simple ladder graph

Exercise 19.5.3 Consider the graph of Fig. 19.15.

- (i) For the effective conductance between a and z , show that $C_{\text{eff}}(a \longleftrightarrow z) = \sqrt{3}$.
- (ii) For a random walk started at a , show that the probability $\mathbf{P}_a[\tau_z < \tau_a]$ of visiting z before returning to a is $\mathbf{P}_a[\tau_z < \tau_a] = 1/\sqrt{3}$.




Fig. 19.15 Simple ladder graph

Exercise 19.5.4 For the graph of Fig. 19.16, determine $C_{\text{eff}}(a \longleftrightarrow z)$ and $\mathbf{P}_a[\tau_z < \tau_a]$. (This is simpler than in Exercise 19.5.3!)




Fig. 19.16 Crossed ladder graph

Exercise 19.5.5 For a random walk on the graph of Fig. 19.17, determine the probability $\mathbf{P}_a[\tau_z < \tau_a]$.




Fig. 19.17 Random walk on a hypercube

5.19.6 19.6 Random Walk in a Random Environment

(Compare [174], [142] and [76, 77], [93].) Consider a Markov chain X on $\mathbb{Z}$ that at each step makes a jump either to the left (with probability w_i^-) or to the right (with probability w_i^+) if X is at $i \in \mathbb{Z}$. Hence, let $w_i^- \in (0, 1)$ and $w_i^+ := 1 - w_i^-$ for $i \in \mathbb{Z}$. Then X is the Markov chain with transition matrix

$$p_w(i, j) = \begin{cases} w_i^-, & \text{if } j = i - 1 \\ w_i^+, & \text{if } j = i + 1 \\ 0, & \text{else} \end{cases}$$

We consider $(w_i^-)_{i \in \mathbb{Z}}$ as an environment in which X walks and later choose the environment at random.

In order to describe X in terms of conductances of an electrical network, we define $\varrho_i := w_i^- / w_i^+$ for $i \in \mathbb{Z}$. Let $C_w(i, j) := 0$ if $|i - j| \neq 1$ and

$$C_w(i + 1, i) := C_w(i, i + 1) := \begin{cases} \prod_{k=0}^i \varrho_k^{-1}, & \text{if } i \geq 0 \\ \prod_{k=i}^{-1} \varrho_k, & \text{if } i < 0 \end{cases}$$

With this definition,

$$\frac{C_w(i, i + 1)}{C_w(i)} = \frac{1}{\varrho_i + 1} = w_i^+ \quad \text{and} \quad \frac{C_w(i, i - 1)}{C_w(i)} = \frac{\varrho_i}{\varrho_i + 1} = w_i^-.$$

Hence the transition probabilities p_w are indeed described by the C_w . Let

$$R_w^+ := \sum_{i=0}^{\infty} R_w(i, i + 1) = \sum_{i=0}^{\infty} \frac{1}{C_w(i, i + 1)} = \sum_{i=0}^{\infty} \prod_{k=0}^i \varrho_k$$

and

$$R_w^- := \sum_{i=0}^{\infty} R_w(-i, -i - 1) = \sum_{i=0}^{\infty} \frac{1}{C_w(-i, -i - 1)} = \sum_{i=1}^{\infty} \prod_{k=-i}^{-1} \varrho_k^{-1}.$$

Note that R_w^+ and R_w^- are the effective resistances from 0 to $+\infty$ and from 0 to $-\infty$, respectively. Hence

$$R_{w,\text{eff}}(0 \leftrightarrow \infty) = \frac{1}{\frac{1}{R_w^-} + \frac{1}{R_w^+}}$$

is finite if and only if $R_w^- < \infty$ or $R_w^+ < \infty$. Therefore, by Theorem 19.25,

$$X \text{ is transient} \iff R_w^- < \infty \text{ or } R_w^+ < \infty. \quad (19.18)$$

If X is transient, in which direction does it get lost?

Theorem 19.33

(i) If $R_w^- < \infty$ or $R_w^+ < \infty$, then (agreeing on $\frac{\infty}{\infty} = 1$)

$$\mathbf{P}_0 \left[X_n \xrightarrow{n \rightarrow \infty} -\infty \right] = \frac{R_w^+}{R_w^- + R_w^+} \quad \text{and} \quad \mathbf{P}_0 \left[X_n \xrightarrow{n \rightarrow \infty} +\infty \right] = \frac{R_w^-}{R_w^- + R_w^+}.$$

(ii) If $R_w^- = \infty$ and $R_w^+ = \infty$, then $\liminf_{n \rightarrow \infty} X_n = -\infty$ and $\limsup_{n \rightarrow \infty} X_n = \infty$ almost surely.

Proof (i) Let $\tau_N := \inf \{n \in \mathbb{N}_0 : X_n \in \{-N, N\}\}$. As X is transient, we have $\mathbf{P}_0[\tau_N < \infty] = 1$ and (as in (19.8))

$$\mathbf{P}_0[X_{\tau_N} = -N] = \frac{R_{w,\text{eff}}(0 \leftrightarrow N)}{R_{w,\text{eff}}(-N \leftrightarrow N)} = \frac{R_{w,\text{eff}}(0 \leftrightarrow N)}{R_{w,\text{eff}}(0 \leftrightarrow -N) + R_{w,\text{eff}}(0 \leftrightarrow N)}.$$

Again, since X is transient, we infer

$$\begin{aligned} \mathbf{P}_0 \left[X_n \xrightarrow{n \rightarrow \infty} -\infty \right] &= \mathbf{P}[\sup \{X_n : n \in \mathbb{N}_0\} < \infty] \\ &= \lim_{N \rightarrow \infty} \mathbf{P}[\sup \{X_n : n \in \mathbb{N}_0\} < N] \\ &\leq \limsup_{N \rightarrow \infty} \mathbf{P}[X_{\tau_N} = -N] \\ &= \frac{R_w^+}{R_w^- + R_w^+}. \end{aligned}$$

By symmetry (and since X is transient), we get

$$\mathbf{P}_0 \left[X_n \xrightarrow{n \rightarrow \infty} -\infty \right] = 1 - \mathbf{P}_0 \left[X_n \xrightarrow{n \rightarrow \infty} \infty \right] \geq 1 - \frac{R_w^-}{R_w^- + R_w^+} = \frac{R_w^+}{R_w^- + R_w^+}$$

(ii) If $R_w^- = R_w^+ = \infty$, then X is recurrent and hence every point is visited infinitely often. That is, $\limsup_{n \rightarrow \infty} X_n = \infty$ and $\liminf_{n \rightarrow \infty} X_n = -\infty$ a.s.

We now consider the situation where the sequence $w = (w_i^-)_{i \in \mathbb{Z}}$ is also random. That is, we consider a two-stage experiment: At the first stage we choose a realization of i.i.d. random variables $W = (W_i^-)_{i \in \mathbb{Z}}$ on $(0, 1)$ and let $W_i^+ := 1 - W_i^-$. At the second stage, given W , we construct a Markov chain X on $\mathbb{Z}$ with transition matrix

$$p_W(i, j) = \begin{cases} W_i^-, & \text{if } j = i - 1 \\ W_i^+, & \text{if } j = i + 1 \\ 0, & \text{else} \end{cases}$$

Note that X is a Markov chain only given W ; that is, under the probability measure $\mathbf{P}[X \in \cdot \mid W]$. However, it is not a Markov chain with respect to the so-called annealed

measure $\mathbf{P}[X \in \cdot]$. In fact, if W is unknown, observing X gives an increasing amount of information on the true realization of W . This is precisely what memory is and is thus in contrast with the Markov property of X .

Definition 19.34 The process X is called a random walk in the random environment W .

We are now in the position to prove a theorem of Solomon [157]. Let $\varrho_i := W_i^-/W_i^+$ for $i \in \mathbb{Z}$ and R_W^- and R_W^+ be defined as above.

Theorem 19.35 (Solomon (1975)) Assume that $\mathbf{E}[\log(\varrho_0)] < \infty$.

- (i) If $\mathbf{E}[\log(\varrho_0)] < 0$, then $X_n \xrightarrow{n \rightarrow \infty} \infty$ a.s.
- (ii) If $\mathbf{E}[\log(\varrho_0)] > 0$, then $X_n \xrightarrow{n \rightarrow \infty} -\infty$ a.s.
- (iii) If $\mathbf{E}[\log(\varrho_0)] = 0$, then $\liminf_{n \rightarrow \infty} X_n = -\infty$ and $\limsup_{n \rightarrow \infty} X_n = \infty$ a.s.

Proof (i) and (ii) By symmetry, it is enough to show (ii). Hence, let $c := \mathbf{E}[\log(\varrho_0)] > 0$. By the strong law of large numbers, there is an $n_0^- = n_0^-(\omega)$ with

$$\prod_{k=-n}^1 \varrho_k^{-1} = \exp\left(-\sum_{k=-n}^1 \log(\varrho_k)\right) < e^{-cn/2} \quad \text{for all } n \geq n_0^-.$$

Therefore,

$$R_W^- = \sum_{n=1}^{\infty} \prod_{k=-n}^1 \varrho_k^{-1} \leq \sum_{n=1}^{n_0^- - 1} \prod_{k=-n}^1 \varrho_k^{-1} + \sum_{n=n_0^-}^{\infty} e^{-cn/2} < \infty \quad \text{a.s.}$$

Similarly, there is an $n_0^+ = n_0^+(\omega)$ with

$$\prod_{k=0}^n \varrho_k > e^{cn/2} \quad \text{for all } n \geq n_0^+.$$

We conclude

$$R_W^+ = \sum_{n=0}^{\infty} \prod_{k=0}^n \varrho_k \geq \sum_{n=0}^{n_0^+ - 1} \prod_{k=0}^n \varrho_k + \sum_{n=n_0^+}^{\infty} e^{cn/2} = \infty \quad \text{a.s.}$$

Now, by Theorem 19.33, we get $X_n \xrightarrow{n \rightarrow \infty} -\infty$ almost surely.

(iii) In order to show $R_W^- = R_W^+ = \infty$ almost surely, it is enough to show $\limsup_{n \rightarrow \infty} \sum_{k=0}^n \log(\varrho_k) > -\infty$ and $\limsup_{n \rightarrow \infty} \sum_{k=-n}^1 \log(\varrho_k^{-1}) > -\infty$ almost surely if $\mathbf{E}[\log(\varrho_0)] = 0$. If $\log(\varrho_0)$ has a finite variance, this follows by the central limit theorem. In the general case, it follows by Theorem 20.21.

Exercise 19.6.1 Consider the situation of Theorem 19.35 but with the random walk restricted to $\mathbb{N}_0$. To this end, change the walk so that whenever it attempts to make a step from 0 to -1, it simply stays in 0. Show that this random walk in a random environment is

- a.s. transient if $\mathbf{E}[\log(\varrho_0)] < \infty$,
- a.s. null recurrent if $\mathbf{E}[\log(\varrho_0)] = \infty$, and
- a.s. positive recurrent if $\mathbf{E}[\log(\varrho_0)] > \infty$.

5.20 20 Ergodic Theory

Laws of large numbers, e.g., for i.i.d. random variables $X_1, X_2, \dots$, state that the sequence of averages converges a.s. to the expected value, $n^{-1} \sum_{i=1}^n X_i \xrightarrow{n \rightarrow \infty} \mathbf{E}[X_1]$. Hence averaging over one realization of many random variables is equivalent to averaging over all possible realizations

of one random variable. In the terminology of statistical physics this means that the time average, or path (Greek: odos) average, equals the space average. The "space" in "space average" is the probability space in mathematical terminology, and in physics it is considered the space of admissible states with a certain energy (Greek: ergon). Combining the Greek words gives rise to the name ergodic theory, which studies laws of large numbers for possibly dependent, but stationary, random variables.

For further reading, see, for example [103] or [88].

5.20.1 20.1 Definitions

Definition 20.1 Let $I \subset \mathbb{R}$ be a set that is closed under addition (for us the important examples are $I = \mathbb{N}_0, I = \mathbb{N}, I = \mathbb{Z}, I = \mathbb{R}, I = [0, \infty), I = \mathbb{Z}^d$ and so on). A stochastic process $X = (X_t)_{t \in I}$ is called stationary if

$$\mathcal{L}[(X_{t+s})_{t \in I}] = \mathcal{L}[(X_t)_{t \in I}] \quad \text{for all } s \in I. \quad (20.1)$$

Remark 20.2 If $I = \mathbb{N}_0, I = \mathbb{N}$ or $I = \mathbb{Z}$, then (20.1) is equivalent to

$$\mathcal{L}[(X_{n+1})_{n \in I}] = \mathcal{L}[(X_n)_{n \in I}]$$

Example 20.3

- (i) If $X = (X_t)_{t \in I}$ is i.i.d., then X is stationary. If only $\mathbf{P}_{X_t} = \mathbf{P}_{X_0}$ holds for every $t \in I$ (without the independence), then in general X is not stationary. For example, consider $I = \mathbb{N}_0$ and $X_1 = X_2 = X_3 = \dots$ but $X_0 \neq X_1$. Then X is not stationary.
- (ii) Let X be a Markov chain with invariant distribution π . If $\mathcal{L}[X_0] = \pi$, then X is stationary.
- (iii) Let $(Y_n)_{n \in \mathbb{Z}}$ be i.i.d. real random variables and let $c_1, \dots, c_k \in \mathbb{R}$. Then

$$X_n := \sum_{l=1}^k c_l Y_{n-l}, \quad n \in \mathbb{Z},$$

defines a stationary process X that is called the moving average with weights $(c_1, \dots, c_k)$. In fact, X is stationary if only Y is stationary.

Lemma 20.4 If $(X_n)_{n \in \mathbb{N}_0}$ is stationary, then X can be extended to a stationary process $(\tilde{X}_n)_{n \in \mathbb{Z}}$.

Proof Let $\tilde{X}$ be the canonical process on $\Omega = E^{\mathbb{Z}}$. For $n \in \mathbb{N}$, define a probability measure $\tilde{\mathbf{P}}^{\{-n, -n+1, \dots\}} \in \mathcal{M}_1(E^{\{-n, -n+1, \dots\}})$ by

$$\begin{aligned} \tilde{\mathbf{P}}^{\{-n, -n+1, \dots\}} \left[\tilde{X}_{-n} \in A_{-n}, \tilde{X}_{-n+1} \in A_{-n+1}, \dots \right] \\ = \mathbf{P} [X_0 \in A_{-n}, X_1 \in A_{-n+1}, \dots]. \end{aligned}$$

Then $\{-n, -n+1, \dots\} \uparrow \mathbb{Z}$ and $(\tilde{\mathbf{P}}^{\{-n, -n+1, \dots\}}, n \in \mathbb{N})$ is a consistent family. By the Ionescu-Tulcea theorem (Theorem 14.32), the projective limit $\tilde{\mathbf{P}} := \lim_{\leftarrow} \tilde{\mathbf{P}}^{\{-n, -n+1, \dots\}}$ exists. By construction, $\tilde{X}$ is stationary with respect to $\tilde{\mathbf{P}}$ and

$$\tilde{\mathbf{P}} \circ \left(\left(\tilde{X}_n \right)_{n \in \mathbb{N}_0} \right)^{-1} = \mathbf{P} \circ \left((X_n)_{n \in \mathbb{N}_0} \right)^{-1}$$

In the following, assume that $(\Omega, \mathcal{A}, \mathbf{P})$ is a probability space and $\tau : \Omega \rightarrow \Omega$ is a measurable map.

Definition 20.5 An event $A \in \mathcal{A}$ is called invariant if $\tau^{-1}(A) = A$ and quasiinvariant if $\mathbb{1}_{\tau^{-1}(A)} = \mathbb{1}_A$ $\mathbf{P}$ -a.s. Denote the σ -algebra of invariant events by

$$\mathcal{I} = \{A \in \mathcal{A} : \tau^{-1}(A) = A\}$$

Recall that a σ -algebra $\mathcal{I}$ is called $\mathbf{P}$ -trivial if $\mathbf{P}[A] \in \{0, 1\}$ for every $A \in \mathcal{I}$.

Definition 20.6

(i) τ is called measure-preserving if

$$\mathbf{P}[\tau^{-1}(A)] = \mathbf{P}[A] \quad \text{for all } A \in \mathcal{A}$$

In this case, $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ is called a measure-preserving dynamical system.

(ii) If τ is measure-preserving and $\mathcal{I}$ is $\mathbf{P}$ -trivial, then $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ is called ergodic.

Lemma 20.7

(i) A measurable map $f : (\Omega, \mathcal{A}) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$ is $\mathcal{I}$ -measurable if and only if $f \circ \tau = f$.

(ii) $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ is ergodic if and only if any $\mathcal{I}$ -measurable $f : (\Omega, \mathcal{I}) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$ is $\mathbf{P}$ -almost surely constant.

Proof (i) The statement is obvious if $f = \mathbb{1}_A$ is an indicator function. The general case, can be inferred by the usual approximation arguments (see Theorem 1.96(i)).

(ii) " $\implies$ " Assume that $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ is ergodic. Then, for any $c \in \mathbb{R}$, we have $f^{-1}((c, \infty)) \in \mathcal{I}$ and thus $\mathbf{P}[f^{-1}((c, \infty))] \in \{0, 1\}$. We conclude that

$$f = \inf \{c \in \mathbb{R} : \mathbf{P}[f^{-1}((c, \infty))] = 0\} \quad \mathbf{P}\text{-a.s.}$$

" $\Leftarrow$ " Assume any $\mathcal{I}$ -measurable map is $\mathbf{P}$ -a.s. constant. If $A \in \mathcal{I}$, then $\mathbb{1}_A$ is $\mathcal{I}$ -measurable and hence $\mathbf{P}$ -a.s. equals either 0 or 1. Thus $\mathbf{P}[A] \in \{0, 1\}$.

Example 20.8 Let $n \in \mathbb{N} \setminus \{1\}$, let $\Omega = \mathbb{Z}/(n)$, let $\mathcal{A} = 2^\Omega$ and let $\mathbf{P}$ be the uniform distribution on Ω . Let $r \in \{1, \dots, n\}$ and

$$\tau : \Omega \rightarrow \Omega, \quad x \mapsto x + r \pmod{n}.$$

Then τ is measure-preserving. If $d = \gcd(n, r)$ and

$$A_i = \{i, \tau(i), \tau^2(i), \dots, \tau^{n-1}(i)\} = i + \langle r \rangle \quad \text{for } i = 0, \dots, d-1,$$

then $A_0, \dots, A_{d-1}$ are the disjoint coset classes of the normal subgroup $\langle r \rangle \trianglelefteq \Omega$. Hence we have $A_i \in \mathcal{I}$ for $i = 0, \dots, d-1$, and each $A \in \mathcal{I}$ is a union of certain A_i 's. Hence we have

$$(\Omega, \mathcal{A}, \mathbf{P}, \tau) \text{ is ergodic} \iff \gcd(r, n) = 1.$$

Example 20.9 (Rotation) Let $\Omega = [0, 1)$, let $\mathcal{A} = \mathcal{B}(\Omega)$ and let $\mathbf{P} = \lambda$ be the Lebesgue measure. Let $r \in (0, 1)$ and $\tau_r(x) = x + r \pmod{1}$. Clearly, $(\Omega, \mathcal{A}, \mathbf{P}, \tau_r)$ is a measure-preserving dynamical system. We will show

$$(\Omega, \mathcal{A}, \mathbf{P}, \tau_r) \text{ is ergodic} \iff r \text{ is irrational.}$$

Let $f : [0, 1) \rightarrow \mathbb{R}$ be an $\mathcal{I}$ -measurable function. Without loss of generality, we assume that f is bounded and thus square integrable. Hence f can be expanded in a Fourier series

$$f(x) = \sum_{n=-\infty}^{\infty} c_n e^{2\pi i n x} \quad \text{for } \mathbf{P}\text{-a.a. } x.$$

This series converges in L^2 , and the sequence of square summable coefficients $(c_n)_{n \in \mathbb{Z}}$ is unique (compare Exercise 7.3.1 with $c_n = (-i/2)a_n + (1/2)b_n$ and $c_{-n} = (i/2)a_n + (1/2)b_n$ for $n \in \mathbb{N}$ as well as $c_0 = b_0$). Now we compute

$$(f \circ \tau_r)(x) = \sum_{n=-\infty}^{\infty} (c_n e^{2\pi i n r}) e^{2\pi i n x} \quad \text{a.e.}$$

By Lemma 20.7, f is $\mathcal{I}$ -measurable if and only if $f = f \circ \tau_r$; that is, if and only if

$$c_n = c_n e^{2\pi i n r} \quad \text{for all } n \in \mathbb{Z}$$

If r is irrational, this implies $c_n = 0$ for $n \neq 0$, and thus f is almost surely constant. Therefore, $(\Omega, \mathcal{A}, \mathbf{P}, \tau_r)$ is ergodic.

On the other hand, if r is rational, then there exists some $n \in \mathbb{Z} \setminus \{0\}$ with $e^{2\pi i n r} = e^{-2\pi i n r} = 1$. Hence $x \mapsto e^{2\pi i n x} + e^{-2\pi i n x} = 2 \cos(2\pi n x)$ is $\mathcal{I}$ -measurable but not a.s. constant. Thus, in this case $(\Omega, \mathcal{A}, \mathbf{P}, \tau_r)$ is not ergodic.

Example 20.10 Let $X = (X_n)_{n \in \mathbb{N}_0}$ be a stochastic process with values in a Polish space E . Without loss of generality, we may assume that X is the canonical process on the probability space $(\Omega, \mathcal{A}, \mathbf{P}) = (E^{\mathbb{N}_0}, \mathcal{B}(E)^{\otimes \mathbb{N}_0}, \mathbf{P})$. Define the shift operator

$$\tau : \Omega \rightarrow \Omega, \quad (\omega_n)_{n \in \mathbb{N}_0} \mapsto (\omega_{n+1})_{n \in \mathbb{N}_0}$$

Then $X_n(\omega) = X_0(\tau^n(\omega))$. Hence X is stationary if and only if $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ is a measure-preserving dynamical system.

Definition 20.11 The stochastic process X (from Example 20.10) is called ergodic if $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ is ergodic.

Example 20.12 Let $(X_n)_{n \in \mathbb{N}_0}$ be i.i.d. and let $X_n(\omega) = X_0(\tau^n(\omega))$. If $A \in \mathcal{I}$, then, for every $n \in \mathbb{N}$,

$$A = \tau^{-n}(A) = \{\omega : \tau^n(\omega) \in A\} \in \sigma(X_n, X_{n+1}, \dots)$$

Hence, if we let $\mathcal{T}$ be the tail σ -algebra of $(X_n)_{n \in \mathbb{N}}$ (see Definition 2.34), then

$$\mathcal{I} \subset \mathcal{T} = \bigcap_{n=1}^{\infty} \sigma(X_n, X_{n+1}, \dots)$$

By Kolmogorov's 0-1 law (Theorem 2.37), $\mathcal{T}$ is $\mathbf{P}$ -trivial. Hence $\mathcal{I}$ is also $\mathbf{P}$ -trivial and therefore $(X_n)_{n \in \mathbb{N}_0}$ is ergodic.

Exercise 20.1.1 Let G be a finite group of measure-preserving measurable maps on $(\Omega, \mathcal{A}, \mathbf{P})$ and let $\mathcal{A}_0 := \{A \in \mathcal{A} : g(A) = A \text{ for all } g \in G\}$.

Show that, for every $X \in \mathcal{L}^1(\mathbf{P})$, we have

$$\mathbf{E}[X \mid \mathcal{A}_0] = \frac{1}{\#G} \sum_{g \in G} X \circ g.$$

5.20.2 20.2 Ergodic Theorems

In this section, $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ always denotes a measure-preserving dynamical system. Further, let $f : \Omega \rightarrow \mathbb{R}$ be measurable and

$$X_n(\omega) = f \circ \tau^n(\omega) \quad \text{for all } n \in \mathbb{N}_0$$

Hence $X = (X_n)_{n \in \mathbb{N}_0}$ is a stationary real-valued stochastic process. Let

$$S_n = \sum_{k=0}^{n-1} X_k$$

denote the n th partial sum. Ergodic theorems are laws of large numbers for $(S_n)_{n \in \mathbb{N}}$. We start with a preliminary lemma.

Lemma 20.13 (Hopf's maximal-ergodic lemma) Let $X_0 \in \mathcal{L}^1(\mathbf{P})$. Define $M_n = \max \{0, S_1, \dots, S_n\}$, $n \in \mathbb{N}$. Then

$$\mathbf{E} [X_0 \mathbb{1}_{\{M_n > 0\}}] \geq 0 \quad \text{for every } n \in \mathbb{N}.$$

Proof For $k \leq n$, we have $M_n(\tau(\omega)) \geq S_k(\tau(\omega))$. Hence

$$X_0 + M_n \circ \tau \geq X_0 + S_k \circ \tau = S_{k+1}$$

Thus $X_0 \geq S_{k+1} - M_n \circ \tau$ for $k = 1, \dots, n$. Manifestly, $S_1 = X_0$ and $M_n \circ \tau \geq 0$ and hence also (for $k = 0$) $X_0 \geq S_1 - M_n \circ \tau$. Therefore,

$$X_0 \geq \max \{S_1, \dots, S_n\} - M_n \circ \tau \quad (20.2)$$

Furthermore, we have

$$\{M_n > 0\}^c \subset \{M_n = 0\} \cap \{M_n \circ \tau \geq 0\} \subset \{M_n - M_n \circ \tau \leq 0\}. \quad (20.3)$$

By (20.2) and (20.3), and since τ is measure-preserving, we conclude that

$$\begin{aligned} \mathbf{E} [X_0 \mathbb{1}_{\{M_n > 0\}}] &\geq \mathbf{E} [(\max \{S_1, \dots, S_n\} - M_n \circ \tau) \mathbb{1}_{\{M_n > 0\}}] \\ &= \mathbf{E} [(M_n - M_n \circ \tau) \mathbb{1}_{\{M_n > 0\}}] \\ &\geq \mathbf{E} [M_n - M_n \circ \tau] = \mathbf{E} [M_n] - \mathbf{E} [M_n] = 0. \end{aligned}$$

Theorem 20.14 (Individual ergodic theorem, Birkhoff (1931), [16]) Let $f = X_0 \in \mathcal{L}^1(\mathbf{P})$. Then

$$\frac{1}{n} \sum_{k=0}^{n-1} X_k = \frac{1}{n} \sum_{k=0}^{n-1} f \circ \tau^k \xrightarrow{n \rightarrow \infty} \mathbf{E} [X_0 \mid \mathcal{I}] \quad \mathbf{P}\text{-a.s.}$$

In particular, if τ is ergodic, then

$$\frac{1}{n} \sum_{k=0}^{n-1} X_k \xrightarrow{n \rightarrow \infty} \mathbf{E} [X_0] \quad \mathbf{P}\text{-a.s.}$$

Proof If τ is ergodic, then $\mathbf{E} [X_0 \mid \mathcal{I}] = \mathbf{E} [X_0]$ and the supplement is a consequence of the first statement.

Consider now the general case. By Lemma 20.7, we have $\mathbf{E} [X_0 \mid \mathcal{I}] \circ \tau = \mathbf{E} [X_0 \mid \mathcal{I}]$ $\mathbf{P}$ -a.s. Hence, by passing to $\tilde{X}_n := X_n - \mathbf{E} [X_0 \mid \mathcal{I}]$, without loss of generality, we can assume $\mathbf{E} [X_0 \mid \mathcal{I}] = 0$. Define

$$Z := \limsup_{n \rightarrow \infty} \frac{1}{n} S_n$$

Let $\varepsilon > 0$ and $F := \{Z > \varepsilon\}$. We have to show that $\mathbf{P}[F] = 0$. From this we infer $\mathbf{P}[Z > 0] = 0$ and similarly (with $-X$ instead of X) also $\liminf_{n \rightarrow \infty} \frac{1}{n} S_n \geq 0$ almost surely. Hence $\frac{1}{n} S_n \xrightarrow{n \rightarrow \infty} 0$ a.s.

Evidently, $Z \circ \tau = Z$; hence $F \in \mathcal{I}$. Define

$$X_n^\varepsilon := (X_n - \varepsilon) \mathbb{1}_F, \quad S_n^\varepsilon := X_0^\varepsilon + \dots + X_{n-1}^\varepsilon$$

$$M_n^\varepsilon := \max \{0, S_1^\varepsilon, \dots, S_n^\varepsilon\}, \quad F_n := \{M_n^\varepsilon > 0\}$$

Then $F_1 \subset F_2 \subset \dots$ and

$$\bigcup_{n=1}^{\infty} F_n = \left\{ \sup_{k \in \mathbb{N}} \frac{1}{k} S_k^\varepsilon > 0 \right\} = \left\{ \sup_{k \in \mathbb{N}} \frac{1}{k} S_k > \varepsilon \right\} \cap F = F$$

hence $F_n \uparrow F$. Dominated convergence yields $\mathbf{E}[X_0^\varepsilon \mathbb{1}_{F_n}] \xrightarrow{n \rightarrow \infty} \mathbf{E}[X_0^\varepsilon]$.

By the maximal-ergodic lemma (applied to X^ε), we have $\mathbf{E}[X_0^\varepsilon \mathbb{1}_{F_n}] \geq 0$; hence

$$0 \leq \mathbf{E}[X_0^\varepsilon] = \mathbf{E}[(X_0 - \varepsilon) \mathbb{1}_F]$$

$$= \mathbf{E}[\mathbf{E}[X_0 | \mathcal{I}] \mathbb{1}_F] - \varepsilon \mathbf{P}[F] = -\varepsilon \mathbf{P}[F]$$

We conclude that $\mathbf{P}[F] = 0$.

As a consequence, we obtain the statistical ergodic theorem, or L^p -ergodic theorem, that was found by von Neumann in 1931 right before Birkhoff proved his ergodic theorem, but was published only later in [122]. Before we formulate it, we state one more lemma.

Lemma 20.15 Let $p \geq 1$ and let $X_0, X_1, \dots$ be identically distributed, real random variables with $\mathbf{E}[|X_0|^p] < \infty$. Define $Y_n := \left| \frac{1}{n} \sum_{k=0}^{n-1} X_k \right|^p$ for $n \in \mathbb{N}$. Then $(Y_n)_{n \in \mathbb{N}}$ is uniformly integrable.

Proof Evidently, the singleton $\{|X_0|^p\}$ is uniformly integrable. Hence, by Theorem 6.19, there exists a monotone increasing convex map $f : [0, \infty) \rightarrow [0, \infty)$ with $\frac{f(x)}{x} \rightarrow \infty$ for $x \rightarrow \infty$ and $C := \mathbf{E}[f(|X_0|^p)] < \infty$. Again, by Theorem 6.19, it is enough to show that $\mathbf{E}[f(Y_n)] \leq C$ for every $n \in \mathbb{N}$. By Jensen's inequality (for $x \mapsto |x|^p$), we have

$$Y_n \leq \frac{1}{n} \sum_{k=0}^{n-1} |X_k|^p$$

Again, by Jensen's inequality (now applied to f), we get that

$$f(Y_n) \leq f\left(\frac{1}{n} \sum_{k=0}^{n-1} |X_k|^p\right) \leq \frac{1}{n} \sum_{k=0}^{n-1} f(|X_k|^p).$$

Hence $\mathbf{E}[f(Y_n)] \leq \frac{1}{n} \sum_{k=0}^{n-1} \mathbf{E}[f(|X_k|^p)] = C$.

Theorem 20.16 (L^p -ergodic theorem, von Neumann (1931)) Let $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ be a measure-preserving dynamical system, $p \geq 1$, $X_0 \in \mathcal{L}^p(\mathbf{P})$ and $X_n = X_0 \circ \tau^n$. Then

$$\frac{1}{n} \sum_{k=0}^{n-1} X_k \xrightarrow{n \rightarrow \infty} \mathbf{E}[X_0 | \mathcal{I}] \quad \text{in } L^p(\mathbf{P})$$

In particular, if τ is ergodic, then $\frac{1}{n} \sum_{k=0}^{n-1} X_k \xrightarrow{n \rightarrow \infty} \mathbf{E}[X_0]$ in $L^p(\mathbf{P})$.

Proof Define

$$Y_n := \left| \frac{1}{n} \sum_{k=0}^{n-1} X_k - \mathbf{E}[X_0 | \mathcal{I}] \right|^p \quad \text{for every } n \in \mathbb{N}.$$

By Lemma 20.15, $(Y_n)_{n \in \mathbb{N}}$ is uniformly integrable, and by Birkhoff's ergodic theorem, we have $Y_n \xrightarrow{n \rightarrow \infty} 0$ almost surely. By Theorem 6.25, we thus have $\lim_{n \rightarrow \infty} \mathbf{E}[Y_n] = 0$.

If τ is ergodic, then $\mathbf{E}[X_0 | \mathcal{I}] = \mathbf{E}[X_0]$.

5.20.3 20.3 Examples

Example 20.17 Let $(X, (\mathbf{P}_x)_{x \in E})$ be a positive recurrent, irreducible Markov chain on the countable space E . Let π be the invariant distribution of X . Then $\pi(\{x\}) > 0$ for every $x \in E$. Define $\mathbf{P}_\pi = \sum_{x \in E} \pi(\{x\}) \mathbf{P}_x$. Then X is stationary on $(\Omega, \mathcal{A}, \mathbf{P}_\pi)$. Denote the shift by τ ; that is, $X_n = X_0 \circ \tau^n$.

Now let $A \in \mathcal{I}$ be invariant. Then $A \in \mathcal{T} = \bigcap_{n=1}^{\infty} \sigma(X_n, X_{n+1}, \dots)$. By the strong Markov property, for every finite stopping time σ (recall that $\mathcal{F}_\sigma$ is the σ algebra of the σ -past),

$$\mathbf{P}_\pi[X \in A \mid \mathcal{F}_\sigma] = \mathbf{P}_{X_\sigma}[X \in A]. \quad (20.4)$$

Indeed, we have $\{X \in A\} = \{X \in \tau^{-n}(A)\} = \{(X_n, X_{n+1}, \dots) \in A\}$. For $B \in \mathcal{F}_\sigma$, using the Markov property (in the third line), we get

$$\begin{aligned} \mathbf{E}_\pi[\mathbb{1}_{\{X \in B\}} \mathbb{1}_{\{X \in A\}}] &= \sum_{n=0}^{\infty} \sum_{x \in E} \mathbf{P}_\pi[X \in B, \sigma = n, X_n = x, X \in A] \\ &= \sum_{n=0}^{\infty} \sum_{x \in E} \mathbf{P}_\pi[X \in B, \sigma = n, X_n = x, X \circ \tau^n \in A] \\ &= \sum_{n=0}^{\infty} \sum_{x \in E} \mathbf{P}_\pi[X \in B, \sigma = n, X_n = x] \mathbf{P}_x[X \in A] \\ &= \mathbf{E}_\pi[\mathbb{1}_{\{X \in B\}} \mathbf{P}_{X_\sigma}[X \in A]]. \end{aligned}$$

In particular, if $x \in E$ and $\sigma_x = \inf\{n \in \mathbb{N}_0 : X_n = x\}$, then $\sigma_x < \infty$ since X is recurrent and irreducible. By (20.4), we conclude that, for every $x \in E$,

$$\mathbf{P}_\pi[X \in A] = \mathbf{E}_\pi[\mathbf{P}_x[X \in A]] = \mathbf{P}_x[X \in A].$$

Hence $\mathbf{P}_{X_n}[X \in A] = \mathbf{P}_\pi[X \in A]$ almost surely and thus (with $\sigma = n$ in (20.4))

$$\mathbf{P}_\pi[X \in A \mid X_0, \dots, X_n] = \mathbf{P}_{X_n}[X \in A] = \mathbf{P}_\pi[X \in A].$$

Now $A \in \mathcal{I} \subset \sigma(X_1, X_2, \dots)$; hence

$$\mathbf{P}_\pi[X \in A \mid X_0, \dots, X_n] \xrightarrow{n \rightarrow \infty} \mathbf{P}_\pi[X \in A \mid \sigma(X_0, X_1, \dots)] = \mathbb{1}_{\{X \in A\}}.$$

This implies $\mathbf{P}_\pi[X \in A] \in \{0, 1\}$. Hence X is ergodic.

Birkhoff's ergodic theorem now implies that, for every $x \in E$,

$$\frac{1}{n} \sum_{k=0}^{n-1} \mathbb{1}_{\{X_k = x\}} \xrightarrow{n \rightarrow \infty} \pi(\{x\}) \quad \mathbf{P}_\pi - \text{a.s.}$$

In this sense, $\pi(\{x\})$ is the average time X spends in x in the long run.

Example 20.18 Let P and Q be probability measures on the measurable space $(\Omega, \mathcal{A})$, and let $(\Omega, \mathcal{A}, P, \tau)$ and $(\Omega, \mathcal{A}, Q, \tau)$ be ergodic. Then either $P = Q$ or $P \perp Q$. Indeed, if $P \neq Q$, then there exists an f with $|f| \leq 1$ and $\int f dP \neq \int f dQ$. However, by Birkhoff's ergodic theorem,

$$\frac{1}{n} \sum_{k=0}^{n-1} f \circ \tau^k \xrightarrow{n \rightarrow \infty} \begin{cases} \int f dP & P\text{-a.s.} \\ \int f dQ & Q\text{-a.s.} \end{cases}$$

If we define $A := \left\{ \frac{1}{n} \sum_{k=0}^{n-1} f \circ \tau^k \xrightarrow{n \rightarrow \infty} \int f dP \right\}$, then $P(A) = 1$ and $Q(A) = 0$. Thus $P \perp Q$.

Exercise 20.3.1 Let $(\Omega, \mathcal{A})$ be a measurable space and let $\tau : \Omega \rightarrow \Omega$ be a measurable map.

- (i) Show that the set $\mathcal{M} := \{\mu \in \mathcal{M}_1(\Omega) : \mu \circ \tau^{-1} = \mu\}$ of τ -invariant measures is convex.
 (ii) An element μ of $\mathcal{M}$ is called extremal if $\mu = \lambda\mu_1 + (1 - \lambda)\mu_2$ for some $\mu_1, \mu_2 \in \mathcal{M}$ and $\lambda \in (0, 1)$ implies $\mu = \mu_1 = \mu_2$. Show that $\mu \in \mathcal{M}$ is extremal if and only if τ is ergodic with respect to μ .

Exercise 20.3.2 Let $p = 2, 3, 5, 6, 7, 10, \dots$ be square-free (that is, there is no number $r = 2, 3, 4, \dots$, whose square is a divisor of p) and let $q \in \{2, 3, \dots, p-1\}$. For every $n \in \mathbb{N}$, let a_n be the leading digit of the p -adic expansion of q^n .

Show the following version of Benford's law: For every $d \in \{1, \dots, p-1\}$,

$$\frac{1}{n} \# \{i \leq n : a_i = d\} \xrightarrow{n \rightarrow \infty} \frac{\log(d+1) - \log(d)}{\log(p)}$$

5.20.4 20.4 Application: Recurrence of Random Walks

Let $(X_n)_{n \in \mathbb{N}}$ be a stationary process with values in $\mathbb{R}^d$. Define $S_n := \sum_{k=1}^n X_k$ for $n \in \mathbb{N}_0$. Further, let

$$R_n = \# \{S_1, \dots, S_n\}$$

denote the range of S ; that is, the number of distinct points visited by S up to time n . Finally, let $A := \{S_n \neq 0 \text{ for every } n \in \mathbb{N}\}$ be the event of an "escape" from 0.

Theorem 20.19 We have $\lim_{n \rightarrow \infty} \frac{1}{n} R_n = \mathbf{P}[A \mid \mathcal{I}]$ almost surely.

Proof Let X be the canonical process on $(\Omega, \mathcal{A}, \mathbf{P}) = \left((\mathbb{R}^d)^{\mathbb{N}}, \mathcal{B}(\mathbb{R}^d)^{\otimes \mathbb{N}}, \mathbf{P} \right)$ and let $\tau : \Omega \rightarrow \Omega$ be the shift; that is, $X_n = X_0 \circ \tau^n$.

Evidently,

$$\begin{aligned} R_n &= \# \{k \leq n : S_l \neq S_k \text{ for all } l \in \{k+1, \dots, n\}\} \\ &\geq \# \{k \leq n : S_l \neq S_k \text{ for all } l > k\} \\ &= \sum_{k=1}^n \mathbb{1}_A \circ \tau^k \end{aligned}$$

Birkhoff's ergodic theorem yields

$$\liminf_{n \rightarrow \infty} \frac{1}{n} R_n \geq \mathbf{P}[A \mid \mathcal{I}] \quad \text{a.s.} \quad (20.5)$$

For the converse inequality, consider $A_m = \{S_l \neq 0 \text{ for } l = 1, \dots, m\}$. Then, for every $n \geq m$,

$$\begin{aligned} R_n &\leq m + \# \{k \leq n - m : S_l \neq S_k \text{ for all } l \in \{k+1, \dots, n\}\} \\ &\leq m + \# \{k \leq n - m : S_l \neq S_k \text{ for all } l \in \{k+1, \dots, k+m\}\} \\ &= m + \sum_{k=1}^{n-m} \mathbb{1}_{A_m} \circ \tau^k \end{aligned}$$

Again, by the ergodic theorem,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} R_n \leq \mathbf{P}[A_m \mid \mathcal{I}] \quad \text{a.s.} \quad (20.6)$$

Since $A_m \downarrow A$ and $\mathbf{P}[A_m \mid \mathcal{I}] \xrightarrow{n \rightarrow \infty} \mathbf{P}[A \mid \mathcal{I}]$ almost surely (by Theorem 8.14(viii)), the claim follows from (20.5) and (20.6).

Theorem 20.20 Let $X = (X_n)_{n \in \mathbb{N}}$ be an integer-valued, integrable, stationary process with the property $\mathbf{E}[X_1 \mid \mathcal{I}] = 0$ a.s. Let $S_n = X_1 + \dots + X_n$, $n \in \mathbb{N}$. Then

$$\mathbf{P}[S_n = 0 \text{ for infinitely many } n \in \mathbb{N}] = 1.$$

In particular, a random walk on $\mathbb{Z}$ with centered increments is recurrent (ChungFuchs theorem, compare Theorem 17.40).

Proof Define $A = \{S_n \neq 0 \text{ for all } n \in \mathbb{N}\}$.

Step 1. We show $\mathbf{P}[A] = 0$. (If X is i.i.d., then S is a Markov chain, and this implies immediately that 0 is recurrent. Only for the more general case of stationary X do we need an additional argument.) By the ergodic theorem, we have

$\frac{1}{n} S_n \xrightarrow{n \rightarrow \infty} \mathbf{E}[X_1 | \mathcal{I}] = 0$ a.s. Thus, for every $m \in \mathbb{N}$,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \left(\frac{1}{n} \max_{k=1, \dots, n} |S_k| \right) &= \limsup_{n \rightarrow \infty} \left(\frac{1}{n} \max_{k=m, \dots, n} |S_k| \right) \\ &\leq \max_{k \geq m} \frac{|S_k|}{k} \xrightarrow{m \rightarrow \infty} 0 \end{aligned}$$

Therefore,

$$\lim_{n \rightarrow \infty} \left(\frac{1}{n} \max_{k=1, \dots, n} S_k \right) = \lim_{n \rightarrow \infty} \left(\frac{1}{n} \min_{k=1, \dots, n} S_k \right) = 0$$

Now $R_n \leq 1 + (\max_{k=1, \dots, n} S_k) - (\min_{k=1, \dots, n} S_k)$; hence $\frac{1}{n} R_n \xrightarrow{n \rightarrow \infty} 0$. By Theorem 20.19, this implies $\mathbf{P}[A] = 0$.

Step 2. Define $\sigma_n := \inf \{m \in \mathbb{N} : S_{m+n} = S_n\}$, $B_n := \{\sigma_n < \infty\}$ for $n \in \mathbb{N}_0$ and $B := \bigcap_{n=0}^{\infty} B_n$.

Since $\{\sigma_0 = \infty\} = A$, we have $\mathbf{P}[\sigma_0 < \infty] = 1$. By stationarity, $\mathbf{P}[\sigma_n < \infty] = 1$ for every $n \in \mathbb{N}_0$; hence $\mathbf{P}[B] = 1$.

Let $\tau_0 = 0$ and inductively define $\tau_{n+1} = \tau_n + \sigma_{\tau_n}$ for $n \in \mathbb{N}_0$. Then τ_n is the time of the n th return of S to 0. On B we have $\tau_n < \infty$ for every $n \in \mathbb{N}_0$ and hence

$$\mathbf{P}[S_n = 0 \text{ infinitely often}] = \mathbf{P}[\tau_n < \infty \text{ for all } n \in \mathbb{N}] \geq \mathbf{P}[B] = 1.$$

If in Theorem 20.20 the random variables X_n are not integer-valued, then there is no hope that $S_n = 0$ for any $n \in \mathbb{N}$ with positive probability. On the other hand, in this case, there is also some kind of recurrence property, namely $S_n/n \xrightarrow{n \rightarrow \infty} 0$ almost surely by the ergodic theorem. Note, however, that this does not exclude the possibility that $S_n \xrightarrow{n \rightarrow \infty} \infty$ with positive probability; for instance, if S_n grows like $\sqrt{n}$. The next theorem shows that if the X_n are integrable, then the process of partial sums can go to infinity only with a linear speed.

Theorem 20.21 Let $(X_n)_{n \in \mathbb{N}}$ be an integrable ergodic process and define $S_n = X_1 + \dots + X_n$ for $n \in \mathbb{N}_0$. Then the following statements are equivalent.

(i) $S_n \xrightarrow{n \rightarrow \infty} \infty$ almost surely.

(ii) $\mathbf{P}\left[S_n \xrightarrow{n \rightarrow \infty} \infty\right] > 0$.

(iii) $\lim_{n \rightarrow \infty} \frac{S_n}{n} = \mathbf{E}[X_1] > 0$ almost surely.

If the random variables $X_1, X_2, \dots$ are i.i.d. with $\mathbf{E}[X_1] = 0$ and $\mathbf{P}[X_1 = 0] < 1$, then $\liminf_{n \rightarrow \infty} S_n = -\infty$ and $\limsup_{n \rightarrow \infty} S_n = \infty$ almost surely.

Proof "(i) $\iff$ (ii)" Clearly, $\left\{S_n \xrightarrow{n \rightarrow \infty} \infty\right\}$ is an invariant event and thus has probability either 0 or 1.

"(iii) $\implies$ (i)" This is trivial.

"(i) $\implies$ (iii)" The equality follows by the individual ergodic theorem. Hence, it is enough to show that $\liminf_{n \rightarrow \infty} S_n/n > 0$ almost surely.

For $n \in \mathbb{N}_0$ and $\varepsilon > 0$, let

$$A_n^\varepsilon := \{S_m > S_n + \varepsilon \text{ for all } m \geq n+1\}.$$

Let $S^- := \inf \{S_n : n \in \mathbb{N}_0\}$. By assumption (i), we have $S^- > -\infty$ almost surely and $\tau := \sup \{n \in \mathbb{N}_0 : S_n = S^-\}$ is finite almost surely. Hence there is an $N \in \mathbb{N}$ with $\mathbf{P}[\tau < N] \geq \frac{1}{2}$. Therefore,

$$\mathbf{P} \left[\bigcup_{n=0}^{N-1} A_n^0 \right] = \mathbf{P}[\tau < N] \geq \frac{1}{2}$$

Since $A_n^\varepsilon \uparrow A_n^0$ for $\varepsilon \downarrow 0$, there is an $\varepsilon > 0$ with $p := \mathbf{P}[A_0^\varepsilon] \geq \frac{1}{4N} > 0$. As $(X_n)_{n \in \mathbb{N}}$ is ergodic, $(\mathbb{1}_{A_n^\varepsilon})_{n \in \mathbb{N}_0}$ is also ergodic. By the individual ergodic theorem, we conclude that $\frac{1}{n} \sum_{i=0}^{n-1} \mathbb{1}_{A_i^\varepsilon} \xrightarrow{n \rightarrow \infty} p$ almost surely. Hence there exists an $n_0 = n_0(\omega)$ such that $\sum_{i=0}^{n-1} \mathbb{1}_{A_i^\varepsilon} \geq \frac{pn}{2}$ for all $n \geq n_0$. This implies $S_n \geq \frac{pn\varepsilon}{2}$ for $n \geq n_0$ and hence $\liminf_{n \rightarrow \infty} S_n/n \geq \frac{pn\varepsilon}{2} > 0$.

The additional statement follows since $\liminf S_n$ and $\limsup S_n$ cannot assume any finite value and are thus measurable with respect to the tail σ -algebra, which implies that they are constantly $-\infty$ or $+\infty$. By what we have shown, we can exclude $S_n \xrightarrow{n \rightarrow \infty} \infty$; hence we have $\liminf_{n \rightarrow \infty} S_n = -\infty$. Similarly, we get $\limsup_{n \rightarrow \infty} S_n = \infty$.

Remark 20.22 It can be shown that Theorem 20.21 holds also without the assumption that the X_n are integrable. See [94].

5.20.5 20.5 Mixing

Ergodicity provides a weak notion of "independence" or "mixing". At the other end of the scale, the strongest notion is "i.i.d.". Here we are concerned with notions of mixing that lie between these two.

In the following, we always assume that $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ is a measure-preserving dynamical system and that $X_n := X_0 \circ \tau^n$. We start with a simple observation.

Theorem 20.23 $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ is ergodic if and only if, for all $A, B \in \mathcal{A}$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} \mathbf{P}[A \cap \tau^{-k}(B)] = \mathbf{P}[A]\mathbf{P}[B] \quad (20.7)$$

Proof " $\implies$ " Let $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ be ergodic. Define

$$Y_n := \frac{1}{n} \sum_{k=0}^{n-1} \mathbb{1}_{\tau^{-k}(B)} = \frac{1}{n} \sum_{k=0}^{n-1} \mathbb{1}_B \circ \tau^k$$

By Birkhoff's ergodic theorem, we have $Y_n \xrightarrow{n \rightarrow \infty} \mathbf{P}[B]$ almost surely. Hence $Y_n \mathbb{1}_A \xrightarrow{n \rightarrow \infty} \mathbb{1}_A \mathbf{P}[B]$ almost surely. Dominated convergence yields

$$\frac{1}{n} \sum_{k=0}^{n-1} \mathbf{P}[A \cap \tau^{-k}(B)] = \mathbf{E}[Y_n \mathbb{1}_A] \xrightarrow{n \rightarrow \infty} \mathbf{E}[\mathbb{1}_A \mathbf{P}[B]] = \mathbf{P}[A]\mathbf{P}[B]$$

" $\Leftarrow$ " Now assume that (20.7) holds. Let $A \in \mathcal{I}$ (recall that $\mathcal{I}$ is the invariant σ -algebra) and $B = A$. Evidently, $A \cap \tau^{-k}(A) = A$ for every $k \in \mathbb{N}_0$. Hence, by (20.7),

$$\mathbf{P}[A] = \frac{1}{n} \sum_{k=0}^{n-1} \mathbf{P}[A \cap \tau^{-k}(A)] \xrightarrow{n \rightarrow \infty} \mathbf{P}[A]^2$$

Thus $\mathbf{P}[A] \in \{0, 1\}$; hence $\mathcal{I}$ is trivial and therefore τ is ergodic.

We consider a strengthening of (20.7).

Definition 20.24 A measure-preserving dynamical system $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ is called mixing if

$$\lim_{n \rightarrow \infty} \mathbf{P}[A \cap \tau^{-n}(B)] = \mathbf{P}[A]\mathbf{P}[B] \quad \text{for all } A, B \in \mathcal{A} \quad (20.8)$$

Remark 20.25 Sometimes the mixing property of (20.8) is called strongly mixing, in contrast with a weakly mixing system $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$, for which we require only

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} |\mathbf{P}[A \cap \tau^{-i}(B)] - \mathbf{P}[A]\mathbf{P}[B]| = 0 \quad \text{for all } A, B \in \mathcal{A}$$

"Strongly mixing" implies "weakly mixing" (see Exercise 20.5.1). On the other hand, there exist weakly mixing systems that are not strongly mixing (see [81]).

Example 20.26 Let $I = \mathbb{N}_0$ or $I = \mathbb{Z}$, and let $(X_n)_{n \in I}$ be an i.i.d. sequence with values in the measurable space $(E, \mathcal{E})$. Hence τ is the shift on the product space $\Omega = E^I$, $\mathbf{P} = (\mathbf{P}_{X_0})^{\otimes I}$. Let $A, B \in \mathcal{E}^{\otimes I}$. For every $\varepsilon > 0$, there exist events A^ε and B^ε that depend on only finitely many coordinates and such that $\mathbf{P}[A \Delta A^\varepsilon] < \varepsilon$ and $\mathbf{P}[B \Delta B^\varepsilon] < \varepsilon$. Clearly, $\mathbf{P}[\tau^{-n}(A \Delta A^\varepsilon)] < \varepsilon$ and $\mathbf{P}[\tau^{-n}(B \Delta B^\varepsilon)] < \varepsilon$ for every $n \in \mathbb{Z}$. For sufficiently large $|n|$, the sets A^ε and $\tau^{-n}(B^\varepsilon)$ depend on different coordinates and are thus independent. This implies

$$\begin{aligned} & \limsup_{|n| \rightarrow \infty} |\mathbf{P}[A \cap \tau^{-n}(B)] - \mathbf{P}[A]\mathbf{P}[B]| \\ & \leq \limsup_{|n| \rightarrow \infty} |\mathbf{P}[A^\varepsilon \cap \tau^{-n}(B^\varepsilon)] - \mathbf{P}[A^\varepsilon]\mathbf{P}[B^\varepsilon]| + 4\varepsilon = 4\varepsilon \end{aligned}$$

Hence τ is mixing. Letting $A = B \in \mathcal{I}$, we obtain the 0-1 law for invariant events: $\mathbf{P}[A] \in \{0, 1\}$.

Remark 20.27 Clearly, (20.8) implies (20.7) and hence "mixing" implies "ergodic". The converse implication is false.

Example 20.28 Let $\Omega = [0, 1)$, $\mathcal{A} = \mathcal{B}([0, 1))$ and let $\mathbf{P} = \lambda$ be the Lebesgue measure on $([0, 1), \mathcal{B}([0, 1)))$. For $r \in [0, 1)$, define $\tau_r : \Omega \rightarrow \Omega$ by

$$\tau_r(x) = x + r - \lfloor x + r \rfloor = x + r \pmod{1}.$$

If r is irrational, then τ_r is ergodic (Example 20.9). However, τ_r is not mixing: Since r is irrational, there exists a sequence $k_n \uparrow \infty$ such that

$$\tau_r^{k_n}(0) \in \left(\frac{1}{4}, \frac{3}{4}\right) \quad \text{for } n \in \mathbb{N}.$$

Hence, for $A = [0, \frac{1}{4}]$, we have $A \cap \tau_r^{-k_n}(A) = \emptyset$. Therefore,

$$\liminf_{n \rightarrow \infty} \mathbf{P}[A \cap \tau_r^{-n}(A)] = 0 \neq \frac{1}{16} = \mathbf{P}[A]^2$$

Theorem 20.29 Let X be an irreducible, positive recurrent Markov chain on the countable space E and let π be its invariant distribution. Let $\mathbf{P}_\pi = \sum_{x \in E} \pi(x) \mathbf{P}_x$. Then:

- (i) X is ergodic (on $(\Omega, \mathcal{A}, \mathbf{P}_\pi)$).
- (ii) X is mixing if and only if X is aperiodic.

Proof (i) This has been shown already in Example 20.17.

(ii) As X is irreducible, by Theorem 17.51, we have $\pi(\{x\}) > 0$ for every $x \in E$.

" $\implies$ " Let X be periodic with period $d \geq 2$. If $n \in \mathbb{N}$ is not a multiple of d , then $p^n(x, x) = 0$. Hence, for $A = B = \{X_0 = x\}$,

$$\begin{aligned}\liminf_{n \rightarrow \infty} \mathbf{P}_\pi [X_0 = x, X_n = x] &= \liminf_{n \rightarrow \infty} \pi(\{x\}) p^n(x, x) \\ &= 0 \neq \pi(\{x\})^2 = \mathbf{P}_\pi [X_0 = x]^2\end{aligned}$$

Thus X is not mixing.

" $\Leftarrow$ " Let X be aperiodic. In order to simplify the notation, we may assume that X is the canonical process on $E^{\mathbb{N}_0}$. Let $A, B \subset \Omega = E^{\mathbb{N}_0}$ be measurable. For every $\varepsilon > 0$, there exists an $N \in \mathbb{N}$ and a $\tilde{A}^\varepsilon \in E^{\{0, \dots, N\}}$ such that, letting $A^\varepsilon = \tilde{A}^\varepsilon \times E^{\{N+1, N+2, \dots\}}$, we have $\mathbf{P}[A \Delta A^\varepsilon] < \varepsilon$. By the Markov property, for every $n \geq N$,

$$\begin{aligned}\mathbf{P}_\pi [A^\varepsilon \cap \tau^{-n}(B)] &= \mathbf{P}_\pi [(X_0, \dots, X_N) \in \tilde{A}^\varepsilon, (X_n, X_{n+1}, \dots) \in B] \\ &= \sum_{x, y \in E} \mathbf{E}_\pi [\mathbb{1}_{A^\varepsilon} \mathbb{1}_{\{X_N=x\}} \mathbb{1}_{\{X_n=y\}} (X_n, X_{n+1}, \dots) \in B] \\ &= \sum_{x, y \in E} \mathbf{E}_\pi [\mathbb{1}_{A^\varepsilon} \mathbb{1}_{\{X_N=x\}}] p^{n-N}(x, y) \mathbf{P}_y[B].\end{aligned}$$

By Theorem 18.13, we have $p^{n-N}(x, y) \xrightarrow{n \rightarrow \infty} \pi(\{y\})$ for all $x, y \in E$. (For periodic X , this is false.) Dominated convergence thus yields

$$\begin{aligned}\lim_{n \rightarrow \infty} \mathbf{P}_\pi [A^\varepsilon \cap \tau^{-n}(B)] &= \sum_{x, y \in E} \mathbf{E}_\pi [\mathbb{1}_{A^\varepsilon} \mathbb{1}_{\{X_N=x\}}] \pi(\{y\}) \mathbf{P}_y[B] \\ &= \mathbf{P}_\pi [A^\varepsilon] \mathbf{P}_\pi [B]\end{aligned}$$

Since $|\mathbf{P}_\pi [A^\varepsilon \cap \tau^{-n}(B)] - \mathbf{P}[A \cap \tau^{-n}(B)]| < \varepsilon$, the statement follows by letting $\varepsilon \rightarrow 0$.

Exercise 20.5.1 Show that "strongly mixing" implies "weakly mixing", which in turn implies "ergodic". Give an example of a measure-preserving dynamical system that is ergodic but not weakly mixing.

5.20.6 20.6 Entropy

The entropy $H(\mathbf{P})$ of a probability distribution $\mathbf{P}$ (see Definition 5.25) measures the amount of randomness in this distribution. In fact, the entropy of a delta distribution is zero and for a distribution on n points, the maximal entropy is achieved by the uniform distribution and equals $\log(n)$ (see Exercise 5.3.3). It is natural to use the entropy in order to quantify also the randomness of a dynamical system.

First we consider the situation of a simple shift: Let $\Omega = E^{\mathbb{N}_0}$, where E is a finite set equipped with the product σ -algebra $\mathcal{A} = (2^E)^{\otimes \mathbb{N}_0}$. Let τ be the shift on Ω and let $\mathbf{P}$ be an invariant probability measure. For $n \in \mathbb{N}$, denote by P_n the projection of $\mathbf{P}$ on $E^n = E^{\{0, \dots, n-1\}}$; that is,

$$P_n(\{(e_0, \dots, e_{n-1})\}) = \mathbf{P}[\{e_0\} \times \dots \times \{e_{n-1}\} \times E^{\{n, n+1, \dots\}}].$$

Denote by h_n the entropy of P_n . By Exercise 5.3.4, the entropy is subadditive:

$$h_{m+n} \leq h_m + h_n \quad \text{for } m, n \in \mathbb{N}.$$

Hence the following limit exists (see Exercise 20.6.2)

$$h := h(\mathbf{P}, \tau) := \lim_{n \rightarrow \infty} \frac{1}{n} h_n = \inf_{n \in \mathbb{N}} \frac{1}{n} h_n$$

Definition 20.30 (Entropy of the simple shift) $h(\mathbf{P}, \tau)$ is called the entropy of the dynamical system $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$.

Example 20.31 Assume that $\mathbf{P}$ is a product measure with marginals π on E . Then

$$h = H(\pi) = - \sum_{e \in E} \pi(\{e\}) \log(\pi(\{e\})).$$

Example 20.32 (Markov chain) Let $(X_n)_{n \in \mathbb{N}_0}$ be a Markov chain on E with transition matrix P and stationary distribution π . Let $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ be the corresponding dynamical system. For $x = (x_0, \dots, x_{n-1})$ and $0 \leq k < n-1$, let

$$p(k, x) = \pi(\{x_k\}) P(x_k, x_{k+1}) \dots P(x_{n-2}, x_{n-1}).$$

Then the entropy of P_n is (using stationarity of π in the third line)

$$\begin{aligned} H(P_n) &= - \sum_{x_0, \dots, x_{n-1} \in E} p(0, x) \log(p(0, x)) \\ &= - \sum_{x_0, \dots, x_{n-1} \in E} p(0, x) \left[\log(\pi(\{x_0\})) + \sum_{k=0}^{n-2} \log(P(x_k, x_{k+1})) \right] \\ &= H(\pi) - \sum_{k=0}^{n-2} \sum_{x_k, \dots, x_{n-1}} p(k, x) \log(P(x_k, x_{k+1})) \\ &= H(\pi) - (n-1) \sum_{x_0, x_1 \in E} \pi(\{x_0\}) P(x_0, x_1) \log(P(x_0, x_1)) \end{aligned}$$

We infer that the entropy of the dynamical system is

$$h(\mathbf{P}, \tau) = - \sum_{x, y \in E} \pi(\{x\}) P(x, y) \log(P(x, y)). \quad (20.9)$$

Example 20.33 (Integer rotation) Consider the rotation of Example 20.8. Let $n \in \mathbb{N} \setminus \{1\}$, $E = \mathbb{Z}/(n)$ and let $\mathbf{P}$ be the uniform distribution on Ω . Let $r \in \{1, \dots, n\}$ and

$$\tau : \Omega \rightarrow \Omega, \quad x \mapsto x + r \pmod{n}.$$

Clearly, $\tau^{(n)}$ is the identity map, hence $h_n = h_{2n} = \dots$ and thus $h(\mathbf{P}, \tau) = 0$.

We now come to the situation of the general dynamical system. Let $\mathcal{P}$ be a finite measurable partition of Ω ; that is, $\mathcal{P} = \{A_1, \dots, A_k\}$ for certain pairwise disjoint non-empty sets $A_1, \dots, A_k \in \mathcal{A}$ with $\Omega = A_1 \cup \dots \cup A_k$. Denote by $\mathcal{P}_n$ the partition that is generated by the sets $\bigcap_{l=0}^{n-1} \tau^{-l}(A_{i_l})$, $i_1, \dots, i_n \in \{1, \dots, k\}$. We define

$$h_n(\mathbf{P}, \tau; \mathcal{P}) = - \sum_{A \in \mathcal{P}_n} \mathbf{P}[A] \log(\mathbf{P}[A])$$

Similarly as in the simple shift case, we obtain the subadditivity of (h_n) and thus the existence of

$$h(\mathbf{P}, \tau; \mathcal{P}) := \lim_{n \rightarrow \infty} \frac{1}{n} h_n(\mathbf{P}, \tau; \mathcal{P}) = \inf_{n \in \mathbb{N}} \frac{1}{n} h_n(\mathbf{P}, \tau; \mathcal{P})$$

Definition 20.34 (Kolmogorov-Sinai entropy) The entropy of a (general) measurepreserving dynamical system $(\Omega, \mathcal{A}, \mathbf{P}, \tau)$ is

$$h(\mathbf{P}, \tau) = \sup_{\mathcal{P}} h(\mathbf{P}, \tau; \mathcal{P})$$

where the supremum is taken over all finite measurable partitions of Ω .

Theorem 20.35 (Kolmogorov-Sinai) Let $\mathcal{P}$ be a generator of $\mathcal{A}$; that is $\mathcal{A} = \sigma\left(\bigcup_{n \in \mathbb{N}_0} \tau^{-n}(\mathcal{P})\right)$. Then

$$h(\mathbf{P}, \tau) = h(\mathbf{P}, \tau; \mathcal{P})$$

Proof See, e.g., [88, Theorem 3.2.18], [167, Theorem 4.17] or [155].

The Kolmogorov-Sinai theorem shows that the entropy that was introduced in Definition 20.30 for simple shifts coincides with the entropy of Definition 20.34; simply take $\mathcal{P} = \{\{e\} \times E^{\mathbb{N}}, e \in E\}$ which generates the product σ -algebra on $\Omega = E^{\mathbb{N}_0}$.

Example 20.36 (Rotation) We come back to the rotation of Example 20.9. Let $\Omega = [0, 1)$, $\mathcal{A} = \mathcal{B}(\Omega)$, $\mathbf{P} = \lambda$ the Lebesgue measure, $r \in (0, 1)$ and $\tau_r(x) = x + r \pmod{1}$.

First assume that r is rational. Let $\mathcal{P}$ be an arbitrary finite measurable partition of Ω . Choose $n \in \mathbb{N}$ such that $rn \in \mathbb{N}_0$. As in Example 20.33 we obtain $h_n(\mathbf{P}, \tau_r; \mathcal{P}) = h_{kn}(\mathbf{P}, \tau_r; \mathcal{P})$ for all $k \in \mathbb{N}$, hence $h(\mathbf{P}, \tau_r, \mathcal{P}) = 0$. Concluding, we get $h(\mathbf{P}, \tau_r) = 0$.

Now assume that r is irrational. Choose the partition $\mathcal{P} = \{[0, 1/2), [1/2, 1)\}$. As r is irrational, it is easy to see that $\mathcal{A}$ is generated by $\bigcup_{n \in \mathbb{N}_0} \tau_r^{-n}(\mathcal{P})$. Hence $h(\mathbf{P}, \tau_r) = h(\mathbf{P}, \tau_r, \mathcal{P})$. In order to compute the latter quantity, we first determine the cardinality $\#\mathcal{P}_n$. To this end, consider the map

$$\begin{aligned} \phi_n : [0, 1) &\rightarrow \{0, 1\}^n \\ x &\mapsto (\mathbb{1}_{[1/2, 1)}(x), \mathbb{1}_{[1/2, 1)}(\tau_r(x)), \dots, \mathbb{1}_{[1/2, 1)}(\tau_r^{n-1}(x))) \end{aligned}$$

Clearly, we have $\#\phi_n([0, 1)) = \#\mathcal{P}_n$. As $x \in [0, 1)$ increases, each coordinate $\mathbb{1}_{[1/2, 1)}(\tau_r^k(x))$, $k = 1, \dots, n-1$, changes its value exactly twice. Only $\mathbb{1}_{[1/2, 1)}(x)$ changes the value exactly once. Summing up, we get $\#\phi_n([0, 1)) \leq 2n$. The maximal entropy of a probability measure on N points is achieved by the uniform distribution and is $\log(N)$. Consequently, $h_n(\mathbf{P}, \tau_r; \mathcal{P}) \leq \log(2n)$. We conclude that

$$h(\mathbf{P}, \tau_r) = h(\mathbf{P}, \tau_r; \mathcal{P}) = 0$$

Exercise 20.6.1 Let $\Omega = [0, 1)$ and $\tau : x \mapsto 2x \pmod{1}$. Let $\mathbf{P}$ be the Lebesgue measure on Ω . Determine $h(\mathbf{P}, \tau)$.

Exercise 20.6.2 Let $(a_n)_{n \in \mathbb{N}}$ be a sequence on nonnegative numbers. The sequence is called subadditive, if $a_{m+n} \leq a_m + a_n$ for all $m, n \in \mathbb{N}$. Show that the limit $\lim_{n \rightarrow \infty} a_n/n$ exists and that

$$\lim_{n \rightarrow \infty} \frac{1}{n} a_n = \inf_{n \in \mathbb{N}} \frac{1}{n} a_n$$

5.21 21 Brownian Motion

In Example 14.45, we constructed a (canonical) process $(X_t)_{t \in [0, \infty)}$ with independent stationary normally distributed increments. For example, such a process can be used to describe the motion of a particle immersed in water or the change of prices in the stock market. We are now interested in properties of this process X that cannot be described in terms of finite-dimensional distributions but reflect the whole path $t \mapsto X_t$. For example, we want to compute the distribution of the functional $F(X) := \sup_{t \in [0, 1]} X_t$. The first problem that has to be resolved is to show that $F(X)$ is a random variable.

In this chapter, we investigate continuity properties of paths of stochastic processes and show how they ensure measurability of some path functionals. Then we construct a version of X that has continuous paths, the so-called Wiener process or Brownian motion. Without exaggeration, it can be stated that Brownian motion is the central object of probability theory.

For further reading, we recommend, e.g., [86, 118, 144, 151].

5.21.1 21.1 Continuous Versions

A priori the paths of a canonical process are of course not continuous since every map $[0, \infty) \rightarrow \mathbb{R}$ is possible. Hence, it will be important to find out which paths are $\mathbf{P}$ -almost surely negligible.

Definition 21.1 Let X and Y be stochastic processes on $(\Omega, \mathcal{A}, \mathbf{P})$ with time set I and state space E . X and Y are called

(i) modifications or versions of each other if, for any $t \in I$, we have

$$X_t = Y_t \quad \mathbf{P}\text{-almost surely ,}$$

(ii) indistinguishable if there exists an $N \in \mathcal{A}$ with $\mathbf{P}[N] = 0$ such that

$$\{X_t \neq Y_t\} \subset N \quad \text{for all } t \in I$$

Clearly, indistinguishable processes are modifications of each other. Under certain assumptions on the continuity of the paths, however, the two notions coincide.

Definition 21.2 Let (E, d) and (E', d') be metric spaces and $\gamma \in (0, 1]$. A map $\varphi : E \rightarrow E'$ is called Hölder-continuous of order γ (briefly, Hölder- γ -continuous) at the point $r \in E$ if there exist $\varepsilon > 0$ and $C < \infty$ such that, for any $s \in E$ with $d(s, r) < \varepsilon$, we have

$$d'(\varphi(r), \varphi(s)) \leq C d(r, s)^\gamma. \quad (21.1)$$

φ is called locally Hölder-continuous of order γ if, for every $t \in E$, there exist $\varepsilon > 0$ and $C = C(t, \varepsilon) > 0$ such that, for all $s, r \in E$ with $d(s, t) < \varepsilon$ and $d(r, t) < \varepsilon$, the inequality (21.1) holds. Finally, φ is called Hölder-continuous of order γ if there exists a C such that (21.1) holds for all $s, r \in E$.

In the case $\gamma = 1$, Hölder continuity is Lipschitz continuity (see Definition 13.8). Furthermore, for $E = \mathbb{R}$ and $\gamma > 1$, every locally Hölder- γ -continuous function is constant. Evidently, a locally Hölder- γ -continuous map is Hölder- γ -continuous at every point. On the other hand, for a function φ that is Hölder- γ -continuous at a given point t , there need not exist an open neighborhood in which φ is continuous. In particular, φ need not be locally Hölder- γ -continuous.

We collect some simple properties of Hölder-continuous functions.

Lemma 21.3 Let $I \subset \mathbb{R}$ and let $f : I \rightarrow \mathbb{R}$ be locally Hölder-continuous of order $\gamma \in (0, 1]$. Then the following statements hold.

- (i) f is locally Hölder-continuous of order γ' for every $\gamma' \in (0, \gamma)$.
- (ii) If I is compact, then f is Hölder-continuous.
- (iii) Let I be a bounded interval of length $T > 0$. Assume that there exists an $\varepsilon > 0$ and an $C(\varepsilon) < \infty$ such that, for all $s, t \in I$ with $|t - s| \leq \varepsilon$, we have

$$|f(t) - f(s)| \leq C(\varepsilon) |t - s|^\gamma$$

Then f is Hölder-continuous of order γ with constant $C := C(\varepsilon) [T/\varepsilon]^{1-\gamma}$.

Proof (i) This is obvious since $|t - s|^\gamma \leq |t - s|^{\gamma'}$ for all $s, t \in I$ with $|t - s| \leq 1$.

(ii) For $t \in I$ and $\varepsilon > 0$, let $U_\varepsilon(t) := \{s \in I : |s - t| < \varepsilon\}$. For every $t \in I$, choose $\varepsilon(t) > 0$ and $C(t) < \infty$ such that

$$|f(r) - f(s)| \leq C(t) \cdot |r - s|^\gamma \quad \text{for all } r, s \in U_t := U_{\varepsilon(t)}(t).$$

There exists a finite subcovering $\mathfrak{U}' = \{U_{t_1}, \dots, U_{t_n}\}$ of the covering $\mathfrak{U} := \{U_t, t \in I\}$ of I . Let $\varrho > 0$ be a Lebesgue number of the covering $\mathfrak{U}'$; that is, $\varrho > 0$ is such that, for every $t \in I$, there exists a $U \in \mathfrak{U}$ such that $U_\varrho(t) \subset U$. Define

$$\bar{C} := \max \{C(t_1), \dots, C(t_n), 2\|f\|_\infty \varrho^{-\gamma}\}.$$

For $s, t \in I$ with $|t - s| < \varrho$, there is an $i \in \{1, \dots, n\}$ with $s, t \in U_{t_i}$. By assumption, we have $|f(t) - f(s)| \leq C(t_i) |t - s|^\gamma \leq \bar{C} |t - s|^\gamma$. Now let $s, t \in I$ with $|s - t| \geq \varrho$. Then

$$|f(t) - f(s)| \leq 2\|f\|_\infty \left(\frac{|t - s|}{\varrho} \right)^\gamma \leq \bar{C} |t - s|^\gamma.$$

Hence f is Hölder-continuous of order γ with constant $\bar{C}$.

(iii) Let $n = \lceil \frac{T}{\varepsilon} \rceil$. For $s, t \in I$, by assumption, $\frac{|t-s|}{n} \leq \varepsilon$ and thus

$$\begin{aligned} |f(t) - f(s)| &\leq \sum_{k=1}^n \left| f\left(s + (t-s)\frac{k}{n}\right) - f\left(s + (t-s)\frac{k-1}{n}\right) \right| \\ &\leq C(\varepsilon) n^{1-\gamma} |t - s|^\gamma = C |t - s|^\gamma \end{aligned}$$

□

Definition 21.4 (Path properties) Let $I \subset \mathbb{R}$ and let $X = (X_t, t \in I)$ be a stochastic process on some probability space $(\Omega, \mathcal{A}, \mathbf{P})$ with values in a metric space (E, d) . Let $\gamma \in (0, 1]$. For every $\omega \in \Omega$, we say that the map $I \rightarrow E, t \mapsto X_t(\omega)$ is a path of X .

We say that X has almost surely continuous paths, or briefly that X is a.s. continuous, if for almost all $\omega \in \Omega$, the path $t \mapsto X_t(\omega)$ is continuous. Similarly, we define locally Hölder- γ -continuous paths and so on.

Lemma 21.5 Let X and Y be modifications of each other. Assume that one of the following conditions holds.

- (i) I is countable.
- (ii) $I \subset \mathbb{R}$ is a (possibly unbounded) interval and X and Y are almost surely right continuous. Then X and Y are indistinguishable.

Proof Define $N_t := \{X_t \neq Y_t\}$ for $t \in I$ and $\bar{N} = \bigcup_{t \in I} N_t$. By assumption, $\mathbf{P}[N_t] = 0$ for every $t \in I$. We have to show that there exists an $N \in \mathcal{A}$ with $\bar{N} \subset N$ and $\mathbf{P}[N] = 0$.

- (i) If I is countable, then $N := \bar{N}$ is measurable and $\mathbf{P}[N] \leq \sum_{t \in I} \mathbf{P}[N_t] = 0$.
- (ii) Now let $I \subset \mathbb{R}$ be an interval and let X and Y be almost surely right continuous. Define

$$\bar{R} := \{X \text{ and } Y \text{ are right continuous}\}$$

and choose an $R \in \mathcal{A}$ with $R \subset \bar{R}$ and $\mathbf{P}[R] = 1$. Define

$$\tilde{I} := \begin{cases} \mathbb{Q} \cap I, & \text{if } I \text{ is open to the right,} \\ (\mathbb{Q} \cap I) \cup \max I, & \text{if } I \text{ is closed to the right,} \end{cases}$$

and $\tilde{N} := \bigcup_{r \in \tilde{I}} N_r$. By (i), we have $\mathbf{P}[\tilde{N}] = 0$. Furthermore, for every $t \in I$,

$$N_t \cap R \subset \bigcup_{r \geq t, r \in \tilde{I}} (N_r \cap R) \subset \tilde{N}.$$

Hence

$$\bar{N} \subset R^c \cup \bigcup_{t \in I} N_t \subset R^c \cup \tilde{N} =: N$$

and thus $\mathbf{P}[N] \leq \mathbf{P}[R^c] + \mathbf{P}[\tilde{N}] = 0$.

We come to the main theorem of this section.

Theorem 21.6 (Kolmogorov-Chentsov) Let $X = (X_t, t \in [0, \infty))$ be a realvalued process. Assume for every $T > 0$, there are numbers $\alpha, \beta, C > 0$ such that

$$\mathbf{E}[|X_t - X_s|^\alpha] \leq C|t - s|^{1+\beta} \quad \text{for all } s, t \in [0, T] \quad (21.2)$$

Then the following statements hold.

- (i) There is a modification $\tilde{X} = (\tilde{X}_t, t \in [0, \infty))$ of X whose paths are locally Hölder-continuous of every order $\gamma \in (0, \frac{\beta}{\alpha})$.
- (ii) Let $\gamma \in (0, \frac{\beta}{\alpha})$. For every $\varepsilon > 0$ and $T < \infty$, there exists a number $K < \infty$ that depends only on $\varepsilon, T, \alpha, \beta, C, \gamma$ such that

$$\mathbf{P}\left[\left|\tilde{X}_t - \tilde{X}_s\right| \leq K|t - s|^\gamma, s, t \in [0, T]\right] \geq 1 - \varepsilon. \quad (21.3)$$

Proof (i) It is enough to show that, for any $T > 0$, the process X on $[0, T]$ has a modification X^T that is locally Hölder-continuous of any order $\gamma \in (0, \beta/\alpha)$. For $S, T > 0$, by Lemma 21.5, two such modifications X^S and X^T are indistinguishable on $[0, S \wedge T]$; hence

$$\Omega_{S,T} := \{ \text{there is a } t \in [0, S \wedge T] \text{ with } X_t^T \neq X_t^S \}$$

is a null set and thus also $\Omega_\infty := \bigcup_{S,T \in \mathbb{N}} \Omega_{S,T}$ is a null set. Therefore, defining $\tilde{X}_t(\omega) := X_t^t(\omega)$, $t \geq 0$, for $\omega \in \Omega \setminus \Omega_\infty$, we get that $\tilde{X}$ is a locally Höldercontinuous modification of X on $[0, \infty)$.

Without loss of generality, assume $T = 1$. We show that X has a continuous modification on $[0, 1]$. By Chebyshev's inequality, for every $\varepsilon > 0$,

$$\mathbf{P}[|X_t - X_s| \geq \varepsilon] \leq C\varepsilon^{-\alpha}|t - s|^{1+\beta} \quad (21.4)$$

Hence

$$X_s \xrightarrow{s \rightarrow t} X_t \quad \text{in probability.} \quad (21.5)$$

The idea is first to construct $\tilde{X}$ on the dyadic rational numbers and then to extend it continuously to $[0, 1]$. To this end, we will need (21.5). In particular, for $\gamma > 0$, $n \in \mathbb{N}$ and $k \in \{1, \dots, 2^n\}$, we have

$$\mathbf{P}[|X_{k2^{-n}} - X_{(k-1)2^{-n}}| \geq 2^{-\gamma n}] \leq C2^{-n(1+\beta-\alpha\gamma)}.$$

Define

$$A_n = A_n(\gamma) := \{ \max \{ |X_{k2^{-n}} - X_{(k-1)2^{-n}}|, k \in \{1, \dots, 2^n\} \} \geq 2^{-\gamma n} \}$$

and

$$B_n := \bigcup_{m=n}^{\infty} A_m \quad \text{and} \quad N := \limsup_{n \rightarrow \infty} A_n = \bigcap_{n=1}^{\infty} B_n.$$

It follows that, for every $n \in \mathbb{N}$,

$$\mathbf{P}[A_n] \leq \sum_{k=1}^{2^n} \mathbf{P}[|X_{k2^{-n}} - X_{(k-1)2^{-n}}| \geq 2^{-\gamma n}] \leq C2^{-n(\beta-\alpha\gamma)}$$

Now fix $\gamma \in (0, \beta/\alpha)$ to obtain

$$\mathbf{P}[B_n] \leq \sum_{m=n}^{\infty} \mathbf{P}[A_m] \leq C \frac{2^{-(\beta-\alpha\gamma)n}}{1-2^{\alpha\gamma-\beta}} \xrightarrow{n \rightarrow \infty} 0 \quad (21.6)$$

hence $\mathbf{P}[N] = 0$. Now fix $\omega \in \Omega \setminus N$ and choose $n_0 = n_0(\omega)$ such that $\omega \notin \bigcup_{n=n_0}^{\infty} A_n$. Hence

$$|X_{k2^{-n}}(\omega) - X_{(k-1)2^{-n}}(\omega)| < 2^{-\gamma n} \quad \text{for } k \in \{1, \dots, 2^n\}, n \geq n_0. \quad (21.7)$$

Define the sets of finite dyadic rationals $D_m = \{k2^{-m}, k = 0, \dots, 2^m\}$, and let $D = \bigcup_{m \in \mathbb{N}} D_m$. Any $t \in D_m$ has a unique dyadic expansion

$$t = \sum_{i=0}^m b_i(t) 2^{-i} \quad \text{for some } b_i(t) \in \{0, 1\}, i = 0, \dots, m$$

Let $m \geq n \geq n_0$ and $s, t \in D_m, s \leq t$ with $|s - t| \leq 2^{-n}$. Let $u := \max(D_n \cap [0, s])$. Then

$$u \leq s < u + 2^{-n} \quad \text{and} \quad u \leq t < u + 2^{1-n}$$

and hence $b_i(t - u) = b_i(s - u) = 0$ for $i < n$. Define

$$t_l = u + \sum_{i=n}^l b_i(t - u) 2^{-i} \quad \text{for } l = n - 1, \dots, m.$$

Then, we have $t_{n-1} = u$ and $t_m = t$. Furthermore, $t_l \in D_l$ for $l = n, \dots, m$ and

$$t_l - t_{l-1} \leq 2^{-l} \quad \text{for } l = n, \dots, m$$

Hence, by (21.7),

$$|X_t(\omega) - X_u(\omega)| \leq \sum_{l=n}^m |X_{t_l}(\omega) - X_{t_{l-1}}(\omega)| \leq \sum_{l=n}^m 2^{-\gamma l} \leq \frac{2^{-\gamma n}}{1 - 2^{-\gamma}}.$$

Analogously, we obtain $|X_s(\omega) - X_u(\omega)| \leq 2^{-\gamma n} (1 - 2^{-\gamma})^{-1}$, and thus

$$|X_t(\omega) - X_s(\omega)| \leq 2 \frac{2^{-\gamma n}}{1 - 2^{-\gamma}} \quad (21.8)$$

Define $C_0 = 2^{1+\gamma} (1 - 2^{-\gamma})^{-1} < \infty$. Let $s, t \in D$ with $|s - t| \leq 2^{-n_0}$. By choosing the minimal $n \geq n_0$ such that $|t - s| \geq 2^{-n}$, we obtain by (21.8),

$$|X_t(\omega) - X_s(\omega)| \leq C_0 |t - s|^\gamma. \quad (21.9)$$

As in the proof of Lemma 21.3(iii), we infer (with $K := C_0 2^{(1-\gamma)n_0}$)

$$|X_t(\omega) - X_s(\omega)| \leq K |t - s|^\gamma \quad \text{for all } s, t \in D. \quad (21.10)$$

In other words, for dyadic rationals D , $X(\omega)$ is (globally) Hölder- γ -continuous. In particular, X is uniformly continuous on D ; hence it can be extended to $[0, 1]$. For $t \in D$, define $\tilde{X}_t := X_t$. For $t \in [0, 1] \setminus D$ and $\{s_n, n \in \mathbb{N}\} \subset D$ with $s_n \rightarrow t$, the sequence $(X_{s_n}(\omega))_{n \in \mathbb{N}}$ is a Cauchy sequence. Hence the limit

$$\tilde{X}_t(\omega) := \lim_{D \ni s \rightarrow t} X_s(\omega) \quad (21.11)$$

exists. Furthermore, the statement analogous to (21.10) holds not only for $s, t \in D$:

$$|\tilde{X}_t(\omega) - \tilde{X}_s(\omega)| \leq K |t - s|^\gamma \quad \text{for all } s, t \in [0, 1]. \quad (21.12)$$

Hence $\tilde{X}$ is locally Hölder-continuous of order γ . By (21.5) and (21.11), we have $\mathbf{P}[X_t \neq \tilde{X}_t] = 0$ for every $t \in [0, 1]$. Hence $\tilde{X}$ is a modification of X .
(ii) Let $\varepsilon > 0$ and choose $n \in \mathbb{N}$ large enough that (see (21.6))

$$\mathbf{P}[B_n] \leq C \frac{2^{-(\beta-\alpha\gamma)n}}{1 - 2^{\alpha\gamma-\beta}} < \varepsilon$$

For $\omega \notin B_n$, we conclude that (21.10) holds. However, this is exactly (21.3) with $T = 1$. For general T , the claim follows by linear scaling.

Remark 21.7 The statement of Theorem 21.6 remains true if X assumes values in some Polish space (E, ϱ) since in the proof we did not make use of the assumption that the range was in $\mathbb{R}$. However, if we change the time set, then the assumptions have to be strengthened: If $(X_t)_{t \in \mathbb{R}^d}$ is a process with values in E , and if, for certain $\alpha, \beta > 0$, all $T > 0$ and some $C < \infty$, we have

$$\mathbf{E}[\varrho(X_t, X_s)^\alpha] \leq C \|t - s\|_2^{d+\beta} \quad \text{for all } s, t \in [-T, T]^d, \quad (21.13)$$

then for every $\gamma \in (0, \beta/\alpha)$, there is a locally Hölder- γ -continuous version of X .

Exercise 21.1.1 Show the claim of Remark 21.7.

Exercise 21.1.2 Let $X = (X_t)_{t \geq 0}$ be a real-valued process with continuous paths. Show that, for all $0 \leq a < b$, the map $\omega \mapsto \int_a^b X_t(\omega) dt$ is measurable.

Exercise 21.1.3 (Optional sampling/stopping) Let $\mathbb{F}$ be a filtration and let $(X_t)_{t \geq 0}$ be an $\mathbb{F}$ -supermartingale with right continuous paths. Let σ and τ be bounded stopping times with $\sigma \leq \tau$. Define $\sigma^n := 2^{-n} \lceil 2^n \sigma \rceil$ and $\tau^n := 2^{-n} \lceil 2^n \tau \rceil$.

(i) Show that $\mathbf{E}[X_{\tau^n} | \mathcal{F}_{\sigma^n}] \xrightarrow{n \rightarrow \infty} \mathbf{E}[X_{\tau^n} | \mathcal{F}_\sigma]$ almost surely and in L^1 as well as $X_{\sigma^n} \xrightarrow{n \rightarrow \infty} X_\sigma$ almost surely and in L^1 .

(ii) Infer the optional sampling theorem for right continuous supermartingales by using the analogous statement for discrete time (Theorem 10.11); that is, $X_\sigma \geq \mathbf{E}[X_\tau | \mathcal{F}_\sigma]$.

(iii) Show that if Y is adapted, integrable and right continuous, then Y is a martingale if and only if $\mathbf{E}[Y_\tau] = \mathbf{E}[Y_0]$ for every bounded stopping time τ .

(iv) Assume that X is uniformly integrable and that $\sigma \leq \tau$ are finite (not necessarily bounded) stopping times. Show that $X_\sigma \geq \mathbf{E}[X_\tau | \mathcal{F}_\sigma]$.

(v) Now let τ be an arbitrary stopping time. Deduce the optional stopping theorem for right continuous supermartingales: $(X_{\tau \wedge t})_{t \geq 0}$ is a right continuous supermartingale.

Exercise 21.1.4 Let $X = (X_t)_{t \geq 0}$ be a stochastic process on $(\Omega, \mathcal{F}, \mathbf{P})$ with values in the Polish space E and with right continuous paths. Show the following.

(i) The map $(\omega, t) \mapsto X_t(\omega)$ is measurable with respect to $\mathcal{F} \otimes \mathcal{B}([0, \infty)) - \mathcal{B}(E)$.

(ii) If in addition X is adapted to the filtration $\mathbb{F}$, then for any $t \geq 0$, the map $\Omega \times [0, t] \rightarrow E, (\omega, s) \mapsto X_s(\omega)$ is $\mathcal{F}_t \otimes \mathcal{B}([0, t]) - \mathcal{B}(E)$ measurable.

(iii) If τ is an $\mathbb{F}$ -stopping time and X is adapted, then X_τ is an $\mathcal{F}_\tau$ -measurable random variable.

5.21.2 21.2 Construction and Path Properties

Definition 21.8 A real-valued stochastic process $B = (B_t, t \in [0, \infty))$ is called a Brownian motion if

- (i) $B_0 = 0$,
- (ii) B has independent, stationary increments (compare Definition 9.7),
- (iii) $B_t \sim \mathcal{N}_{0,t}$ for all $t > 0$, and
- (iv) $t \mapsto B_t$ is $\mathbf{P}$ -almost surely continuous.




Fig. 21.1 Computer simulation of a Brownian motion

See Fig. 21.1 for a computer simulation of a Brownian motion.

Theorem 21.9 There exists a probability space $(\Omega, \mathcal{A}, \mathbf{P})$ and a Brownian motion B on $(\Omega, \mathcal{A}, \mathbf{P})$. The paths of B are a.s. locally Hölder- γ -continuous for every $\gamma < \frac{1}{2}$.

Proof As in Example 14.45 or Corollary 16.10 there exists a stochastic process X that fulfills (i), (ii) and (iii). Evidently, $X_t - X_s \stackrel{\mathcal{D}}{=} \sqrt{t-s}X_1 \sim \mathcal{N}_{0,t-s}$ for all $t > s \geq 0$. Thus, for every $n \in \mathbb{N}$, writing $C_n := \mathbf{E}[X_1^{2n}] = \frac{(2n)!}{2^n n!} < \infty$, we have

$$\mathbf{E}[(X_t - X_s)^{2n}] = \mathbf{E}[(\sqrt{t-s}X_1)^{2n}] = C_n |t-s|^n.$$

Now let $n \geq 2$ and $\gamma \in (0, \frac{n-1}{2n})$. Theorem 21.6 yields the existence of a version B of X that has Hölder- γ -continuous paths. Since all continuous versions of a process are equivalent, B is locally Hölder- γ -continuous for every $\gamma \in (0, \frac{n-1}{2n})$ and every $n \geq 2$ and hence for every $\gamma \in (0, \frac{1}{2})$. $\square$

Recall that a stochastic process $(X_t)_{t \in I}$ is called a Gaussian process if, for every $n \in \mathbb{N}$ and for all $t_1, \dots, t_n \in I$, we have that

$$(X_{t_1}, \dots, X_{t_n}) \quad \text{is } n\text{-dimensional normally distributed.}$$

X is called centered if $\mathbf{E}[X_t] = 0$ for every $t \in I$. The map

$$\Gamma(s, t) := \text{Cov}[X_s, X_t] \quad \text{for } s, t \in I$$

is called the covariance function of X .

Remark 21.10 The covariance function determines the finite-dimensional distributions of a centered Gaussian process since a multidimensional normal distribution is determined by the vector of expectations and by the covariance matrix.

Theorem 21.11 Let $X = (X_t)_{t \in [0, \infty)}$ be a stochastic process. Then the following are equivalent:

- (i) X is a Brownian motion.
- (ii) X is a continuous centered Gaussian process with $\mathbf{Cov}[X_s, X_t] = s \wedge t$ for all $s, t \geq 0$.

Proof By Remark 21.10, X is characterized by (ii). Hence, it is enough to show that, for Brownian motion X , we have $\text{Cov}[X_s, X_t] = \min(s, t)$. This is indeed true since for $t > s$, the random variables X_s and $X_t - X_s$ are independent; hence

$$\text{Cov}[X_s, X_t] = \text{Cov}[X_s, X_t - X_s] + \text{Cov}[X_s, X_s] = \text{Var}[X_s] = s.$$

Corollary 21.12 (Scaling property of Brownian motion) If B is a Brownian motion and if $K \neq 0$, then $(K^{-1}B_{K^2t})_{t \geq 0}$ is also a Brownian motion.

Example 21.13 Another example of a continuous Gaussian process is the so-called Brownian bridge $X = (X_t)_{t \in [0,1]}$ that is defined by the covariance function $\Gamma(s, t) = s \wedge t - st$. We construct the Brownian bridge as follows.

Let $B = (B_t, t \in [0, 1])$ be a Brownian motion and let

$$X_t := B_t - tB_1.$$

Clearly, X is a centered Gaussian process with continuous paths. We compute the covariance function Γ of X ,

$$\begin{aligned} \Gamma(s, t) &= \text{Cov}[X_s, X_t] = \text{Cov}[B_s - sB_1, B_t - tB_1] \\ &= \text{Cov}[B_s, B_t] - s \text{Cov}[B_1, B_t] - t \text{Cov}[B_s, B_1] + st \text{Cov}[B_1, B_1] \\ &= \min(s, t) - st - st + st = \min(s, t) - st \end{aligned}$$

Theorem 21.14 Let $(B_t)_{t \geq 0}$ be a Brownian motion and

$$X_t = \begin{cases} tB_{1/t}, & \text{if } t > 0 \\ 0, & \text{if } t = 0 \end{cases}$$

Then X is a Brownian motion.

Proof Clearly, X is a Gaussian process. For $s, t > 0$, we have

$$\text{Cov}[X_s, X_t] = ts \cdot \text{Cov}[B_{1/s}, B_{1/t}] = ts \min(s^{-1}, t^{-1}) = \min(s, t).$$

Clearly, $t \mapsto X_t$ is continuous at every point $t > 0$. To show continuity at $t = 0$, consider

$$\begin{aligned} \limsup_{t \downarrow 0} X_t &= \limsup_{t \rightarrow \infty} \frac{1}{t} B_t \\ &\leq \limsup_{n \rightarrow \infty} \frac{1}{n} B_n + \limsup_{n \rightarrow \infty} \frac{1}{n} \sup \{B_t - B_n, t \in [n, n+1]\}. \end{aligned}$$

By the strong law of large numbers, we have $\lim_{n \rightarrow \infty} \frac{1}{n} B_n = 0$ a.s. Using a generalization of the reflection principle (Theorem 17.15; see also Theorem 21.19), for $x > 0$, we have (using the abbreviation $B_{[a,b]} := \{B_t : t \in [a, b]\}$)

$$\begin{aligned} \mathbf{P}[\sup B_{[n,n+1]} - B_n > x] &= \mathbf{P}[\sup B_{[0,1]} > x] = 2\mathbf{P}[B_1 > x] \\ &= \frac{2}{\sqrt{2\pi}} \int_x^\infty e^{-u^2/2} du \leq \frac{1}{x} e^{-x^2/2} \end{aligned}$$

In particular, $\sum_{n=1}^\infty \mathbf{P}[\sup B_{[n,n+1]} - B_n > n^\varepsilon] < \infty$ for every $\varepsilon > 0$. By the Borel-Cantelli lemma (Theorem 2.7), we infer

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sup \{B_t - B_n, t \in [n, n+1]\} = 0 \quad \text{almost surely.}$$

Hence X is also continuous at 0.

Theorem 21.15 (Blumenthal's 0-1 law, see [18]) Let B be a Brownian motion and let $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0} = \sigma(B)$ be the filtration generated by B . Further, let $\mathcal{F}_0^+ = \bigcap_{t > 0} \mathcal{F}_t$. Then $\mathcal{F}_0^+$ is a $\mathbf{P}$ -trivial σ -algebra.

Proof Define $Y^n = (B_{2^{-n}+t} - B_{2^{-n}})_{t \in [0, 2^{-n}]}$, $n \in \mathbb{N}$. Then $(Y^n)_{n \in \mathbb{N}}$ is an independent family of random variables (with values in $C([0, 2^{-n}])$). By Kolmogorov's 0-1 law (Theorem 2.37), the tail σ -algebra $\mathcal{T} = \bigcap_{n \in \mathbb{N}} \sigma(Y^m, m \geq n)$ is $\mathbf{P}$ -trivial. On the other hand, $\sigma(Y^m, m \geq n) = \mathcal{F}_{2^{-n+1}}$; hence

$$\mathcal{F}_0^+ = \bigcap_{t>0} \mathcal{F}_t = \bigcap_{n \in \mathbb{N}} \mathcal{F}_{2^{-n+1}} = \mathcal{T}$$

is $\mathbf{P}$ -trivial.

Example 21.16 Let B be a Brownian motion. For every $K > 0$, we have

$$\mathbf{P} \left[\inf \left\{ t > 0 : B_t \geq K\sqrt{t} \right\} = 0 \right] = 1 \quad (21.14)$$

To check this, define $A_s := \left\{ \inf \left\{ t > 0 : B_t \geq K\sqrt{t} \right\} \leq s \right\}$ and

$$A := \left\{ \inf \left\{ t > 0 : B_t \geq K\sqrt{t} \right\} = 0 \right\} = \bigcap_{s>0} A_s \in \mathcal{F}_0^+$$

Then $\mathbf{P}[A] \in \{0, 1\}$. By the scaling property of Brownian motion,

$$\mathbf{P}[A] = \inf_{s>0} \mathbf{P}[A_s] \geq \mathbf{P}[B_1 \geq K] > 0$$

and thus $\mathbf{P}[A] = 1$.

The preceding example shows that, for every $t \geq 0$, almost surely B is not Hölder- $\frac{1}{2}$ -continuous at t . Note that the order of quantifiers is subtle. We have not shown that almost surely B was not Hölder- $\frac{1}{2}$ -continuous at any $t \geq 0$ (however, see Remark 22.4). However, it is not too hard to show the following theorem, which for the case $\gamma = 1$ is due to Paley, Wiener and Zygmund [126]. The proof presented here goes back to an idea of Dvoretzky, Erdős and Kakutani (see [40]).

Theorem 21.17 (Paley-Wiener-Zygmund (1933)) For every $\gamma > \frac{1}{2}$, almost surely the paths of Brownian motion $(B_t)_{t \geq 0}$ are not Hölder-continuous of order γ at any point. In particular, the paths are almost surely nowhere differentiable.

Proof Let $\gamma > \frac{1}{2}$. It suffices to consider $B = (B_t)_{t \in [0,1]}$. Denote by $H_{\gamma,t}$ the set of maps $[0,1] \rightarrow \mathbb{R}$ that are Hölder- γ -continuous at t and define $H_\gamma := \bigcup_{t \in [0,1]} H_{\gamma,t}$. The aim is to show that almost surely $B \notin H_\gamma$.

If $t \in [0,1)$ and $w \in H_{\gamma,t}$, then for every $\delta > 0$ there exists a $c = c(\delta, w)$ with the property $|w_s - w_t| \leq c|s - t|^\gamma$ for every $s \in [0,1]$ with $|s - t| < \delta$. Choose a $k \in \mathbb{N}$ with $k > \frac{2}{2\gamma-1}$. Then, for $n \in \mathbb{N}$ with $n \geq n_0 := \Gamma(k+1)/\delta$, $i = \lfloor tn \rfloor + 1$ and $l \in \{0, \dots, k-1\}$, we get

$$\begin{aligned} |w_{(i+l+1)/n} - w_{(i+l)/n}| &\leq |w_{(i+l+1)/n} - w_t| + |w_{(i+l)/n} - w_t| \\ &\leq 2c(k+1)^\gamma n^{-\gamma}. \end{aligned}$$

Hence, for $N \geq 2c(k+1)^\gamma$, we have $w \in A_{N,n,i}$, where

$$A_{N,n,i} := \bigcap_{l=0}^{k-1} \left\{ w : |w_{(i+l+1)/n} - w_{(i+l)/n}| \leq Nn^{-\gamma} \right\}.$$

Define $A_{N,n} = \bigcup_{i=1}^n A_{N,n,i}$, $A_N = \liminf_{n \rightarrow \infty} A_{N,n}$ and $A = \bigcup_{N=1}^\infty A_N$. Clearly, $H_\gamma \subset A$. Owing to the independence of increments and since the density of the standard normal distribution is bounded by 1, we get

$$\begin{aligned} \mathbf{P}[B \in A_{N,n,i}] &= \mathbf{P}[|B_{1/n}| \leq Nn^{-\gamma}]^k = \mathbf{P}[|B_1| \leq Nn^{-\gamma+1/2}]^k \\ &\leq N^k n^{k(-\gamma+1/2)} \end{aligned}$$

By the choice of k and since the increments of B are stationary, we have

$$\begin{aligned} \mathbf{P}[B \in A_N] &= \lim_{n \rightarrow \infty} \mathbf{P} \left[\bigcap_{m \geq n} A_{N,m} \right] \leq \limsup_{n \rightarrow \infty} \mathbf{P}[A_{N,n}] \leq \limsup_{n \rightarrow \infty} \sum_{i=1}^n \mathbf{P}[A_{N,n,i}] \\ &\leq \limsup_{n \rightarrow \infty} n \mathbf{P}[B \in A_{N,n,1}] \leq N^k \limsup_{n \rightarrow \infty} n^{1+k(-\gamma+1/2)} = 0 \end{aligned}$$

Thus $\mathbf{P}[B \in A] = 0$. Therefore, we almost surely have $B \notin H_\gamma$.

Exercise 21.2.1 Let B be a Brownian motion and let λ be the Lebesgue measure on $[0, \infty)$.

(i) Compute the expectation and variance of $\int_0^1 B_s ds$. (For the measurability of the integral see Exercise 21.1.2.)

(ii) Show that almost surely $\lambda(\{t : B_t = 0\}) = 0$.

(iii) Compute the expectation and variance of

$$\int_0^1 \left(B_t - \int_0^1 B_s ds \right)^2 dt$$

Exercise 21.2.2 Let B be a Brownian motion. Show that $(B_t^2 - t)_{t \geq 0}$ is a martingale.

Exercise 21.2.3 Let B be a Brownian motion and $\sigma > 0$. Show that the process $\left(\exp \left(\sigma B_t - \frac{\sigma^2}{2} t \right) \right)_{t \geq 0}$ is a martingale.

Exercise 21.2.4 Let B be a Brownian motion, $a < 0 < b$. Define the stopping time $\tau_{a,b} = \inf \{t \geq 0 : B_t \in \{a, b\}\}$.

Show that almost surely $\tau_{a,b} < \infty$ and that $\mathbf{P}[B_{\tau_{a,b}} = b] = -\frac{a}{b-a}$. Furthermore, show (using Exercise 21.2.2) that $\mathbf{E}[\tau_{a,b}] = -ab$.

Exercise 21.2.5 Let B be a Brownian motion, $b > 0$ and $\tau_b = \inf \{t \geq 0 : B_t = b\}$. Show the following.

(i) $\mathbf{E}[e^{-\lambda \tau_b}] = e^{-b\sqrt{2\lambda}}$ for $\lambda \geq 0$.

Hint: Use Exercise 21.2.3 and the optional sampling theorem.

(ii) τ_b has a $\frac{1}{2}$ -stable distribution with Lévy measure

$$v(dx) = (b/(\sqrt{2\pi}))x^{-3/2}\mathbb{1}_{\{x>0\}}dx$$

(iii) The distribution of τ_b has density $f_b(x) = \frac{b}{\sqrt{2\pi}}e^{-b^2/(2x)}x^{-3/2}$.

Exercise 21.2.6 Let B be a Brownian motion, $a \in \mathbb{R}$, $b > 0$ and $\tau = \inf \{t \geq 0 : B_t = at + b\}$. For $\lambda \geq 0$, show that

$$\mathbf{E}[e^{-\lambda \tau}] = \exp \left(-ba - b\sqrt{a^2 + 2\lambda} \right).$$

Conclude that $\mathbf{P}[\tau < \infty] = 1 \wedge e^{-2ba}$.

5.21.3 21.3 Strong Markov Property

Denote by $\mathbf{P}_x$ the probability measure such that $B = (B_t)_{t \geq 0}$ is a Brownian motion started at $x \in \mathbb{R}$. To put it differently, under $\mathbf{P}_x$, the process $(B_t - x)_{t \geq 0}$ is a standard Brownian motion. While the (simple) Markov property of $(B, (\mathbf{P}_x)_{x \in \mathbb{R}})$ is evident, it takes some work to check the strong Markov property.

Theorem 21.18 (Strong Markov property) Brownian motion B with distributions $(\mathbf{P}_x)_{x \in \mathbb{R}}$ has the strong Markov property.

Proof Let $\mathbb{F} = \sigma(B)$ be the filtration generated by B and let $\tau < \infty$ be an $\mathbb{F}$ stopping time. We have to show that, for every bounded measurable $F : \mathbb{R}^{[0, \infty)} \rightarrow \mathbb{R}$, we have:

$$\mathbf{E}_x[F((B_{t+\tau})_{t \geq 0}) \mid \mathcal{F}_\tau] = \mathbf{E}_{B_\tau}[F(B)]. \quad (21.15)$$

It is enough to consider continuous bounded functions F that depend on only finitely many coordinates $t_1, \dots, t_N$ since these functions determine the distribution of $(B_{t+\tau})_{t \geq 0}$. Hence, let $f : \mathbb{R}^N \rightarrow \mathbb{R}$ be continuous and bounded and $F(B) = f(B_{t_1}, \dots, B_{t_N})$. Manifestly, the map $x \mapsto \mathbf{E}_x[F(B)] = \mathbf{E}_0[f(B_{t_1} + x, \dots, B_{t_N} + x)]$ is continuous and bounded. Now let $\tau^n := 2^{-n} \lfloor 2^n \tau + 1 \rfloor$ for $n \in \mathbb{N}$. Then τ^n is a stopping time and $\tau^n \downarrow \tau$; hence $B_{\tau^n} \xrightarrow{n \rightarrow \infty} B_\tau$ almost surely. Now every Markov process with countable time set (here all positive rational linear combinations of $1, t_1, \dots, t_N$) is a strong Markov process (by Theorem 17.14); hence we have

$$\begin{aligned} \mathbf{E}_x[F((B_{\tau^n+t})_{t \geq 0}) \mid \mathcal{F}_{\tau^n}] &= \mathbf{E}_x[f(B_{\tau^n+t_1}, \dots, B_{\tau^n+t_N}) \mid \mathcal{F}_{\tau^n}] \\ &= \mathbf{E}_{B_{\tau^n}}[f(B_{t_1}, \dots, B_{t_N})] \\ &\xrightarrow{n \rightarrow \infty} \mathbf{E}_{B_\tau}[f(B_{t_1}, \dots, B_{t_N})] = \mathbf{E}_{B_\tau}[F(B)]. \end{aligned} \quad (21.16)$$

As B is right continuous, we have $F((B_{\tau^n+t})_{t \geq 0}) \xrightarrow{n \rightarrow \infty} F((B_{\tau+t})_{t \geq 0})$ almost surely and in L^1 and thus

$$\begin{aligned} \mathbf{E}[|\mathbf{E}_x[F((B_{\tau^n+t})_{t \geq 0}) \mid \mathcal{F}_{\tau^n}] - \mathbf{E}_x[F((B_{\tau+t})_{t \geq 0}) \mid \mathcal{F}_{\tau^n}]|] \\ \leq \mathbf{E}_x[|F((B_{\tau^n+t})_{t \geq 0}) - F((B_{\tau+t})_{t \geq 0})|] \xrightarrow{n \rightarrow \infty} 0. \end{aligned} \quad (21.17)$$

Furthermore,

$$\mathcal{F}_{\tau^n} \downarrow \mathcal{F}_{\tau+} := \bigcap_{\sigma > \tau \text{ is a stopping time}} \mathcal{F}_\sigma \supset \mathcal{F}_\tau$$

By (21.16) and (21.17), using the convergence theorem for backwards martingales (Theorem 12.14), we get that in the sense of L^1 -limits

$$\begin{aligned} \mathbf{E}_{B_\tau}[F(B)] &= \lim_{n \rightarrow \infty} \mathbf{E}_x[F((B_{\tau^n+t})_{t \geq 0}) \mid \mathcal{F}_{\tau^n}] \\ &= \lim_{n \rightarrow \infty} \mathbf{E}_x[F((B_{\tau+t})_{t \geq 0}) \mid \mathcal{F}_{\tau^n}] = \mathbf{E}_x[F((B_{\tau+t})_{t \geq 0}) \mid \mathcal{F}_{\tau+}] \end{aligned}$$

The left-hand side is $\mathcal{F}_\tau$ -measurable. The tower property of conditional expectation thus yields (21.15).

Using the strong Markov property, we show the reflection principle for Brownian motion.

Theorem 21.19 (Reflection principle for Brownian motion) For every $a > 0$ and $T > 0$,

$$\mathbf{P}[\sup\{B_t : t \in [0, T]\} > a] = 2\mathbf{P}[B_T > a] \leq \frac{2\sqrt{T}}{\sqrt{2\pi}} \frac{1}{a} e^{-a^2/2T}$$

Proof By the scaling property of Brownian motion (Corollary 21.12), without loss of generality, we may assume $T = 1$. Let $\tau := \inf\{t \geq 0 : B_t \geq a\} \wedge 1$. By symmetry, we have $\mathbf{P}_a[B_{1-\tau} > a] = \frac{1}{2}$ if $\tau < 1$; hence

$$\begin{aligned} \mathbf{P}[B_1 > a] &= \mathbf{P}[B_1 > a \mid \tau < 1] \mathbf{P}[\tau < 1] \\ &= \mathbf{P}_a[B_{1-\tau} > a] \mathbf{P}[\tau < 1] = \frac{1}{2} \mathbf{P}[\tau < 1] \end{aligned}$$

For the inequality compute

$$\begin{aligned} \mathbf{P}[B_1 > a] &= \frac{1}{\sqrt{2\pi}} \int_a^\infty e^{-x^2/2} dx \\ &\leq \frac{1}{\sqrt{2\pi}} \frac{1}{a} \int_a^\infty x e^{-x^2/2} dx = \frac{1}{\sqrt{2\pi}} \frac{1}{a} e^{-a^2/2} \end{aligned}$$

As an application of the reflection principle we derive Paul Lévy's arcsine law [107, p. 216] for the last time a Brownian motion visits zero.

Theorem 21.20 (Lévy's arcsine law) Let $T > 0$ and $\zeta_T := \sup \{t \leq T : B_t = 0\}$. Then, for $t \in [0, T]$,

$$\mathbf{P}[\zeta_T \leq t] = \frac{2}{\pi} \arcsin(\sqrt{t/T})$$

Proof Without loss of generality, assume $T = 1$ and $\zeta = \zeta_1$. Let $\tilde{B}$ be a further, independent Brownian motion. By the reflection principle,

$$\begin{aligned} \mathbf{P}[\zeta \leq t] &= \mathbf{P}[B_s \neq 0 \text{ for all } s \in [t, 1]] \\ &= \int_{-\infty}^{\infty} \mathbf{P}[B_s \neq 0 \text{ for all } s \in [t, 1] \mid B_t = a] \mathbf{P}[B_t \in da] \\ &= \int_{-\infty}^{\infty} \mathbf{P}_{|a|}[\tilde{B}_s > 0 \text{ for all } s \in [0, 1-t]] \mathbf{P}[B_t \in da] \\ &= \int_{-\infty}^{\infty} \mathbf{P}_0[|\tilde{B}_{1-t}| \leq |a|] \mathbf{P}[B_t \in da] \\ &= \mathbf{P}[|\tilde{B}_{1-t}| \leq |B_t|] \end{aligned}$$

If X and Y are independent and $\mathcal{N}_{0,1}$ -distributed, then

$$(B_t, \tilde{B}_{1-t}) \stackrel{\mathcal{D}}{=} (\sqrt{t}X, \sqrt{1-t}Y)$$

Hence

$$\begin{aligned} \mathbf{P}[\zeta \leq t] &= \mathbf{P}[\sqrt{1-t}|Y| \leq \sqrt{t}|X|] \\ &= \mathbf{P}[Y^2 \leq t(X^2 + Y^2)] \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} dx \int_{-\infty}^{\infty} dy e^{-(x^2+y^2)/2} \mathbb{1}_{\{y^2 \leq t(x^2+y^2)\}} \end{aligned}$$

Passing to polar coordinates, we obtain

$$\mathbf{P}[\zeta \leq t] = \frac{1}{2\pi} \int_0^{\infty} r dr e^{-r^2/2} \int_0^{2\pi} d\varphi \mathbb{1}_{\{\sin(\varphi)^2 \leq t\}} = \frac{2}{\pi} \arcsin(\sqrt{t})$$

Exercise 21.3.1 (Hard problem!) Let $\mathbf{P}_x$ be the distribution of Brownian motion started at $x \in \mathbb{R}$. Let $a > 0$ and $\tau = \inf \{t \geq 0 : B_t \in \{0, a\}\}$. Use the reflection principle to show that, for every $x \in (0, a)$,

$$\mathbf{P}_x[\tau > T] = \sum_{n=-\infty}^{\infty} (-1)^n \mathbf{P}_x[B_T \in [na, (n+1)a]] \quad (21.18)$$

If f is the density of a probability distribution on $\mathbb{R}$ with characteristic function φ and $\sup_{x \in \mathbb{R}} x^2 f(x) < \infty$, then the Poisson summation formula holds,

$$\sum_{n=-\infty}^{\infty} f(s+n) = \sum_{k=-\infty}^{\infty} \varphi(k) e^{2\pi i s k} \quad \text{for all } s \in \mathbb{R} \quad (21.19)$$

Use (21.18) and (21.19) (compare also (21.37)) to conclude that

$$\mathbf{P}_x[\tau > T] = \frac{4}{\pi} \sum_{k=0}^{\infty} \frac{1}{2k+1} \exp\left(-\frac{(2k+1)^2 \pi^2 T}{2a^2}\right) \sin\left(\frac{(2k+1)\pi x}{a}\right). \quad (21.20)$$

5.21.4 21.4 Supplement: Feller Processes

In many situations, a continuous version of a process would be too much to expect, for instance, the Poisson process is generically discontinuous. However, often there is a version with right continuous paths that have left-sided limits. At this point, we only briefly make plausible the existence theorem for such regular versions of processes in the case of so-called Feller semigroups.

Definition 21.21 Let E be a Polish space. A map $f : [0, \infty) \rightarrow E$ is called *RCLL* (right continuous with left limits) or *càdlàg* (continue à droite, limites à gauche) if $f(t) = f(t+) := \lim_{s \downarrow t} f(s)$ for every $t \geq 0$ and if, for every $t > 0$, the left-sided limit $f(t-) := \lim_{s \uparrow t} f(s)$ exists and is finite.

Definition 21.22 A filtration $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ is called right continuous if $\mathbb{F} = \mathbb{F}^+$, where $\mathcal{F}_t^+ = \bigcap_{s > t} \mathcal{F}_s$. We say that a filtration $\mathbb{F}$ satisfies the usual conditions (from the French conditions habituelles) if $\mathbb{F}$ is right continuous and if $\mathcal{F}_0$ is $\mathbf{P}$ -complete.

Remark 21.23 If $\mathbb{F}$ is an arbitrary filtration and $\mathcal{F}_t^{+,*}$ is the completion of $\mathcal{F}_t^+$, then $\mathbb{F}^{+,*}$ satisfies the usual conditions.

Theorem 21.24 (Doob's regularization) Let $\mathbb{F}$ be a filtration that satisfies the usual conditions and let $X = (X_t)_{t \geq 0}$ be an $\mathbb{F}$ -supermartingale such that $t \mapsto \mathbf{E}[X_t]$ is right continuous. Then there exists a modification $\tilde{X}$ of X with RCLL paths.

Proof For $a, b \in \mathbb{Q}^+$, $a < b$ and $I \subset [0, \infty)$, let $U_I^{a,b}$ be the number of upcrossings of $(X_t)_{t \in I}$ over $[a, b]$. By the upcrossing inequality (Lemma 11.3), for every $N > 0$ and every finite set $I \subset [0, N]$, we have $\mathbf{E}[U_I^{a,b}] \leq (\mathbf{E}[|X_N|] + |a|)/(b - a)$. Define $U_N^{a,b} = U_{\mathbb{Q}^+ \cap [0, N]}^{a,b}$. Then $\mathbf{E}[U_N^{a,b}] \leq (\mathbf{E}[|X_N|] + |a|)/(b - a)$. By Exercise 11.1.1, for $\lambda > 0$, we have

$$\begin{aligned} \lambda \mathbf{P}[\sup\{|X_t| : t \in \mathbb{Q}^+ \cap [0, N]\} > \lambda] \\ &= \lambda \sup\{\mathbf{P}[\sup\{|X_t| : t \in I\} > \lambda] : I \subset \mathbb{Q}^+ \cap [0, N] \text{ finite}\} \\ &\leq 12\mathbf{E}[|X_0|] + 9\mathbf{E}[|X_N|]. \end{aligned}$$

Consider the event

$$A := \bigcap_{N \in \mathbb{N}} \left(\bigcap_{\substack{a, b \in \mathbb{Q}^+ \\ 0 \leq a < b \leq N}} \{U_N^{a,b} < \infty\} \cap \{\sup\{|X_t| : t \in \mathbb{Q}^+ \cap [0, N]\} < \infty\} \right)$$

We have $\mathbf{P}[A] = 1$; hence $A \in \mathcal{F}_t$ for every $t \geq 0$ since $\mathbb{F}$ satisfies the usual conditions. For $\omega \in A$, for every $t \geq 0$, the limit

$$\tilde{X}_t(\omega) := \lim_{\mathbb{Q}^+ \ni s \downarrow t, s > t} X_s(\omega)$$

exists and is RCLL. For $\omega \in A^c$, we define $\tilde{X}_t(\omega) = 0$. As $\mathbb{F}$ satisfies the usual conditions, $\tilde{X}$ is $\mathbb{F}$ -adapted. As X is a supermartingale, for every N , the family $(X_s)_{s \leq N}$ is uniformly integrable. Hence, by assumption,

$$\mathbf{E}[\tilde{X}_t] = \lim_{\mathbb{Q}^+ \ni s \downarrow t, s > t} \mathbf{E}[X_s] = \mathbf{E}[X_t]$$

However, since X is a supermartingale, for every $s > t$, we have

$$X_t \geq \mathbf{E}[X_s | \mathcal{F}_t] \xrightarrow{\mathbb{Q}^+ \ni s \downarrow t, s > t} \mathbf{E}[\tilde{X}_t | \mathcal{F}_t] = \tilde{X}_t \quad \text{in } L^1.$$

Therefore, $X_t = \tilde{X}_t$ almost surely and hence $\tilde{X}$ is a modification of X .

Corollary 21.25 Let $(v_t)_{t \geq 0}$ be a continuous convolution semigroup and assume that $\int |x|v_1(dx) < \infty$. Then there exists a Markov process X with RCLL paths and with independent stationary increments $\mathbf{P}_{X_t - X_s} = v_{t-s}$ for all $t > s$.

Let E be a locally compact Polish space and let $C_0(E)$ be the set of (bounded) continuous functions that vanish at infinity. If κ is a stochastic kernel from E to E and if f is measurable and bounded, then we define $\kappa f(x) = \int \kappa(x, dy)f(y)$.

Definition 21.26 A Markov semigroup $(\kappa_t)_{t \geq 0}$ on E is called a Feller semigroup if

$$f(x) = \lim_{t \rightarrow 0} \kappa_t f(x) \quad \text{for all } x \in E, f \in C_0(E)$$

and $\kappa_t f \in C_0(E)$ for every $f \in C_0(E)$.

Let X be a Markov process with transition kernels $(\kappa_t)_{t \geq 0}$ and with respect to a filtration $\mathbb{F}$ that satisfies the usual conditions.

Let $g \in C_0(E)$, $g \geq 0$. Let $h = \int_0^\infty e^{-t} \kappa_t g dt$. Then

$$e^{-s} \kappa_s h = e^{-s} \int_0^\infty e^{-t} \kappa_s \kappa_t g dt = \int_s^\infty e^{-t} \kappa_t g dt \leq h$$

Hence $X^g := (e^{-t} h(X_t))_{t \geq 0}$ is an $\mathbb{F}$ -supermartingale.

The Feller property and Theorem 21.24 ensure the existence of an RCLL version $\tilde{X}^g$ of X^g . It takes a little more work to show that there exists a countable set $G \subset C_0(E)$ and a process $\tilde{X}$ that is uniquely determined by $\tilde{X}^g, g \in G$, and is an RCLL version of X . See, e.g., [146, Chapter III.7ff].

Let us take a moment's thought and look back at how we derived the strong Markov property of Brownian motion in Section 21.3. Indeed, there we needed only right continuity of the paths and a certain continuity of the distribution as a function of the starting point, which is exactly the Feller property. With a little more work, one can show the following theorem (see, e.g., [146, Chapter III.8ff] or [144, Chapter III, Theorem 2.7]).

Theorem 21.27 Let $(\kappa_t)_{t \geq 0}$ be a Feller semigroup on the locally compact Polish space E . Then there exists a strong Markov process $(X_t)_{t \geq 0}$ with RCLL paths and transition kernels $(\kappa_t)_{t \geq 0}$.

Such a process X is called a Feller process.

Exercise 21.4.1 (Doob's inequality) Let $X = (X_t)_{t \geq 0}$ be a martingale or a nonnegative submartingale with RCLL paths. For $T \geq 0$, let $|X|_T^* = \sup_{t \in [0, T]} |X_t|$. Show Doob's inequalities:

- (i) For any $p \geq 1$ and $\lambda > 0$, we have $\lambda^p \mathbf{P}[|X|_T^* \geq \lambda] \leq \mathbf{E}[|X_T|^p]$.
- (ii) For any $p > 1$, we have $\mathbf{E}[|X_T|^p] \leq \mathbf{E}[(|X|_T^*)^p] \leq \left(\frac{p}{p-1}\right)^p \mathbf{E}[|X_T|^p]$.

Construct a counterexample that shows that right continuity of the paths of X is essential.

Exercise 21.4.2 (Martingale convergence theorems) Let X be a stochastic process with RCLL paths. Use Doob's inequality (Exercise 21.4.1) to show that the martingale convergence theorems (a.s. convergence (Theorem 11.4), a.s. and L^1 convergence for uniformly integrable martingales (Theorem 11.7) and the L^p martingale convergence theorem (Theorem 11.10)) hold for X .

Exercise 21.4.3 Let $p \geq 1$ and let $X^1, X^2, X^3, \dots$ be L^p -integrable martingales. Assume that, for every $t \geq 0$, there exists an $\tilde{X}_t \in \mathcal{L}^p(\mathbf{P})$ such that $X_t^n \xrightarrow{n \rightarrow \infty} \tilde{X}_t$ in L^p .

- (i) Show that $(\tilde{X}_t)_{t \geq 0}$ is a martingale.
- (ii) Use Doob's inequality to show the following. If $p > 1$ and if $X^1, X^2, \dots$ are a.s. continuous, then there is a continuous martingale X with the following properties: X is a modification of $\tilde{X}$ and $X_t^n \xrightarrow{n \rightarrow \infty} X_t$ in L^p for every $t \geq 0$.

Exercise 21.4.4 Let X be a stochastic process with values in a Polish space E and with RCLL paths. Let $\mathbb{F} = \sigma(X)$ be the filtration generated by X and define $\mathbb{F}^+ := (\mathcal{F}_t^+)_{t \geq 0}$ by $\mathcal{F}_t^+ = \bigcap_{s > t} \mathcal{F}_s$. Let $U \subset E$ be open and let $C \subset E$ be closed. For every set $A \subset E$, define $\tau_A := \inf \{t > 0 : X_t \in A\}$. Show the following.

- (i) τ_C is an $\mathbb{F}$ -stopping time (and an $\mathbb{F}^+$ -stopping time).
- (ii) τ_U is an $\mathbb{F}^+$ -stopping time but in general (even for continuous X) is not an $\mathbb{F}$ -stopping time.

Exercise 21.4.5 Show the statement of Remark 21.23. Conclude that if $\mathbb{F}$ is a filtration and if B is a Brownian motion that is an $\mathbb{F}$ -martingale, then B is also an $\mathbb{F}^{+,*}$ -martingale.

5.21.5 21.5 Construction via l^2 -Approximation

We give an alternative construction of Brownian motion by functional analytic means as an L^2 -approximation. For simplicity, as the time interval we take $[0, 1]$ instead of $[0, \infty)$.

Let $H = L^2([0, 1])$ be the Hilbert space of square integrable (with respect to Lebesgue measure λ) functions $[0, 1] \rightarrow \mathbb{R}$ with inner product

$$\langle f, g \rangle = \int_{[0,1]} f(x)g(x)\lambda(dx)$$

and with norm $\|f\| = \sqrt{\langle f, f \rangle}$ (compare Section 7.3). Two functions $f, g \in H$ are considered equal if $f = g$ λ -a.e. Let $(b_n)_{n \in \mathbb{N}}$ be an orthonormal basis (ONB) of H ; that is, $\langle b_m, b_n \rangle = \mathbb{1}_{\{m=n\}}$ and

$$\lim_{n \rightarrow \infty} \left\| f - \sum_{m=1}^n \langle f, b_m \rangle b_m \right\| = 0 \quad \text{for all } f \in H.$$

In particular, for every $f \in H$, Parseval's equation

$$\|f\|^2 = \sum_{m=1}^{\infty} \langle f, b_m \rangle^2 \tag{21.21}$$

holds and for $f, g \in H$

$$\langle f, g \rangle = \sum_{m=1}^{\infty} \langle f, b_m \rangle \langle g, b_m \rangle \tag{21.22}$$

Now consider an i.i.d. sequence $(\xi_n)_{n \in \mathbb{N}}$ of $\mathcal{N}_{0,1}$ -random variables on some probability space $(\Omega, \mathcal{A}, \mathbf{P})$. For $n \in \mathbb{N}$ and $t \in [0, 1]$, define

$$X_t^n = \int \mathbb{1}_{[0,t]}(s) \left(\sum_{m=1}^n \xi_m b_m(s) \right) \lambda(ds) = \sum_{m=1}^n \xi_m \langle \mathbb{1}_{[0,t]}, b_m \rangle.$$

Clearly, for $n \geq m$,

$$\begin{aligned} \mathbf{E} [(X_t^m - X_t^n)^2] &= \mathbf{E} \left[\left(\sum_{k=m+1}^n \xi_k \langle \mathbb{1}_{[0,t]}, b_k \rangle \right) \left(\sum_{l=m+1}^n \xi_l \langle \mathbb{1}_{[0,t]}, b_l \rangle \right) \right] \\ &= \sum_{k=m+1}^n \langle \mathbb{1}_{[0,t]}, b_k \rangle^2 \leq \sum_{k=m+1}^{\infty} \langle \mathbb{1}_{[0,t]}, b_k \rangle^2. \end{aligned}$$

Since $\sum_{k=1}^{\infty} \langle \mathbb{1}_{[0,t]}, b_k \rangle^2 = \|\mathbb{1}_{[0,t]}\|^2 = t < \infty$, we have $X_t^n \in L^2(\mathbf{P})$ and

$$\lim_{m \rightarrow \infty} \sup_{n \geq m} \mathbf{E} [(X_t^m - X_t^n)^2] = 0$$

Hence $(X_t^n)_{n \in \mathbb{N}}$ is a Cauchy sequence in $L^2(\mathbf{P})$ and thus (since $L^2(\mathbf{P})$ is complete, see Theorem 7.3) has an L^2 -limit X_t . Thus, for $N \in \mathbb{N}$ and $0 \leq t_1, \dots, t_N \leq 1$,

$$\lim_{n \rightarrow \infty} \mathbf{E} \left[\sum_{i=1}^N (X_{t_i}^n - X_{t_i})^2 \right] = 0.$$

In particular, $(X_{t_1}^n, \dots, X_{t_N}^n) \xrightarrow{n \rightarrow \infty} (X_{t_1}, \dots, X_{t_N})$ in $\mathbf{P}$ -probability.

Manifestly, $(X_{t_1}^n, \dots, X_{t_N}^n)$ is normally distributed and centered. For $s, t \in [0, 1]$, we have

$$\begin{aligned} \text{Cov}[X_s^n, X_t^n] &= \mathbf{E} \left[\left(\sum_{k=1}^n \xi_k \langle \mathbb{1}_{[0,s]}, b_k \rangle \right) \left(\sum_{l=1}^n \xi_l \langle \mathbb{1}_{[0,t]}, b_l \rangle \right) \right] \\ &= \sum_{k,l=1}^n \mathbf{E} [\xi_k \xi_l] \langle \mathbb{1}_{[0,s]}, b_k \rangle \langle \mathbb{1}_{[0,t]}, b_l \rangle \\ &= \sum_{k=1}^n \langle \mathbb{1}_{[0,s]}, b_k \rangle \langle \mathbb{1}_{[0,t]}, b_k \rangle \\ &\xrightarrow{n \rightarrow \infty} \langle \mathbb{1}_{[0,s]}, \mathbb{1}_{[0,t]} \rangle = \min(s, t) \end{aligned}$$

Hence $(X_t)_{t \in [0,1]}$ is a centered Gaussian process with

$$\text{Cov}[X_s, X_t] = \min(s, t). \quad (21.23)$$

Lévy Construction of Brownian Motion

Up to continuity of paths, X is thus a Brownian motion. A continuous version of X can be obtained via the Kolmogorov-Chentsov theorem (Theorem 21.6). However,




Fig. 21.2 The processes $\{X^n\}_{n=0,1,2,3,10}$ of the Lévy construction of Brownian motion

by a clever choice of the ONB $(b_n)_{n \in \mathbb{N}}$, we can construct X directly as a continuous process. The Haar functions $b_{n,k}$ are one such choice: Let $b_{0,1} \equiv 1$ and for $n \in \mathbb{N}$ and $k = 1, \dots, 2^{n-1}$, let

$$b_{n,k}(t) = \begin{cases} 2^{(n-1)/2}, & \text{if } \frac{2k-2}{2^n} \leq t < \frac{2k-1}{2^n}, \\ -2^{(n-1)/2}, & \text{if } \frac{2k-1}{2^n} \leq t < \frac{2k}{2^n}, \\ 0, & \text{else.} \end{cases}$$

Then $(b_{n,k})$ is an orthonormal system: $\langle b_{m,k}, b_{n,l} \rangle = \mathbb{1}_{\{(m,k)=(n,l)\}}$. It is easy to check that $(b_{n,k})$ is a basis (exercise!). Define the Schauder functions by

$$B_{n,k}(t) = \int_{[0,t]} b_{n,k}(s) \lambda(ds) = \langle \mathbb{1}_{[0,t]}, b_{n,k} \rangle.$$

Let $\xi_{0,1}, (\xi_{n,k})_{n \in \mathbb{N}, k=1, \dots, 2^{n-1}}$ be independent and $\mathcal{N}_{0,1}$ -distributed. Let

$$X^n := \xi_{0,1} B_{0,1} + \sum_{m=1}^n \sum_{k=1}^{2^{m-1}} \xi_{m,k} B_{m,k},$$

and define $\tilde{X}_t$ as the $L^2(\mathbf{P})$ -limit $\tilde{X}_t = L^2 - \lim_{n \rightarrow \infty} X_t^n$. See Fig. 21.2 for a computer simulation of $X^n, n = 0, 1, 2, 3, 10$.

Theorem 21.28 (Brownian motion, L^2 -approximation) There is a continuous version X of $\tilde{X}$. X is a Brownian motion and we have

$$\lim_{n \rightarrow \infty} \|X^n - X\|_\infty = 0 \quad \mathbf{P}\text{-almost surely} \quad (21.24)$$

Proof By (21.24), we have $X_t = \tilde{X}_t$ a.s. for all $t \in [0, 1]$. As uniform limits of continuous functions are continuous, (21.24) implies that X is continuous. Hence, by (21.23) (and Theorem 21.11), X is a Brownian motion. Therefore, it is enough to prove the existence of an X such that (21.24) holds.

Since $(C([0, 1]), \|\cdot\|_\infty)$ is complete, it suffices to show that $\mathbf{P}$ -almost surely (X^n) is a Cauchy sequence in $(C([0, 1]), \|\cdot\|_\infty)$. Note that $\|B_{n,k}\|_\infty \leq 2^{-(n+1)/2}$ if $n \in \mathbb{N}$ and $B_{n,k} B_{n,l} = 0$ if $k \neq l$. Hence

$$\|X^n - X^{n-1}\|_\infty \leq 2^{-(n+1)/2} \max \{|\xi_{n,k}|, k = 1, \dots, 2^{n-1}\}.$$

Therefore,

$$\begin{aligned} \mathbf{P} [\|X^n - X^{n-1}\|_\infty > 2^{-n/4}] &\leq \sum_{k=1}^{2^{n-1}} \mathbf{P} [|\xi_{n,k}| > 2^{(n+2)/4}] \\ &= 2^{n-1} \frac{2}{\sqrt{2\pi}} \int_{2^{(n+2)/4}}^{\infty} e^{-x^2/2} dx \\ &\leq 2^n \exp(-2^{n/2}) \end{aligned}$$

Evidently, $\sum_{n=1}^{\infty} \mathbf{P} [\|X^n - X^{n-1}\|_\infty > 2^{-n/4}] < \infty$; hence, by the Borel-Cantelli lemma,

$$\mathbf{P} [\|X^n - X^{n-1}\|_\infty > 2^{-n/4} \quad \text{only finitely often}] = 1.$$

We conclude that $\lim_{n \rightarrow \infty} \sup_{m \geq n} \|X^m - X^n\|_\infty = 0$ $\mathbf{P}$ -almost surely.

Brownian Motion and White Noise

The construction of Brownian motion via Haar functions has the advantage that continuity of the paths is straightforward. For some applications, however, a decomposition in trigonometric functions is preferable. Here as the orthonormal basis of $L^2([0, 1])$ we use $b_0 = 1$ and

$$b_n(x) = \sqrt{2} \cos(n\pi x) \quad \text{for } n \in \mathbb{N}.$$

For $t \in [0, 1]$ and $n \in \mathbb{N}_0$, define

$$B_n(t) = \int_0^t b_n(s) \lambda(ds)$$




Fig. 21.3 The processes X^n , $n=0,1,4,64,8192$ from the Fourier Construction of Brownian motion

that is, $B_0(t) = t$ and

$$B_n(t) = \frac{\sqrt{2}}{n\pi} \sin(n\pi t) \quad \text{for } n \in \mathbb{N}.$$

Let $\xi_n, n \in \mathbb{N}_0$, be independent standard normally distributed random variables. Define $A_0 = \xi_0$ and

$$A_n := \frac{\sqrt{2}}{\pi n} \xi_n \quad \text{for } n \in \mathbb{N}.$$

Finally, let

$$X^n := \sum_{k=0}^n \xi_k B_k;$$

that is,

$$X^n(t) = \xi_0 t + \sum_{k=1}^n A_k \sin(k\pi t).$$

See Fig. 21.3 for a computer simulation of $X^n, n = 0, 1, 4, 64, 8192$.

As shown above, the sequence (X^n) converges in $L^2([0, 1])$ towards a process X , which (up to continuity of paths) has all properties of Brownian motion:

$$X_t = \xi_0 t + \sum_{n=1}^{\infty} \frac{\sqrt{2}}{n\pi} \xi_n \sin(n\pi t).$$

This representation of the Brownian motions goes back to Paley and Wiener who also show that along a suitable subsequence the series converges uniformly almost surely and hence the limit X is indeed continuous, see [125, Theorem XLIII, p. 148]. The representation is also

sometimes called Karhunen-Loève expansion. More precisely, up to the first summand, it is the Karhunen-Loève expansion of the Brownian bridge $(X_t - tX_1)_{t \in [0,1]}$ (see, e.g., [1, Chapter 3.3]).

Taking the formal derivative

$$\dot{X}_t := \frac{d}{dt}X_t = \xi_0 + \sqrt{2} \sum_{n=1}^{\infty} \xi_n \cos(n\pi t)$$

we get independent identically distributed Fourier coefficients for all frequencies. Hence, the formal object $\dot{X}$ is often referred to as white noise as opposed to colored noise where the coefficients for the different frequencies have different distributions.

The Fourier basis is not too well suited to showing continuity of paths. For example, the sufficient criterion of absolute summability of coefficients (A_n) fails (see Exercise 21.5.5).

Example 21.29 (Stochastic integral à la Paley-Wiener) Assume that $(\xi_n)_{n \in \mathbb{N}}$ is an i.i.d. sequence of $\mathcal{N}_{0,1}$ -distributed random variables. Let $(b_n)_{n \in \mathbb{N}}$ be an orthonormal basis of $L^2([0, 1])$ such that $W_t := \lim_{n \rightarrow \infty} \sum_{k=1}^n \xi_k \langle \mathbb{1}_{[0,t]}, b_k \rangle$, $t \in [0, 1]$, is a Brownian motion. For $f \in L^2([0, 1])$, define

$$I(f) := \sum_{n=1}^{\infty} \xi_n \langle f, b_n \rangle$$

By Parseval's equation and the Bienaymé formula, we have

$$\|f\|_2^2 = \sum_{n=1}^{\infty} \langle f, b_n \rangle^2 = \text{Var}[I(f)] = \mathbf{E} [I(f)^2]$$

Hence

$$I : L^2([0, 1]) \rightarrow L^2(\mathbf{P}), \quad f \mapsto I(f) \quad \text{is an isometry.} \quad (21.25)$$

We call

$$\int_0^t f(s) dW_s := I(f \mathbb{1}_{[0,t]}), \quad t \in [0, 1], f \in L^2([0, 1])$$

the stochastic integral of f with respect to W . In the special case of the Fourier basis $b_0(x) = 1$ and $b_n(x) = \sqrt{2} \cos(n\pi x)$, $n \in \mathbb{N}$, this construction goes back to Paley and Wiener [125, Theorem XLV, p. 154].

The process $X_t := \int_0^t f(s) dW_s$, $t \in [0, 1]$, is centered Gaussian with covariance function

$$\text{Cov}[X_s, X_t] = \int_0^{s \wedge t} f^2(u) du$$

In fact, it is obvious that X is centered and Gaussian (since it is a limit of the Gaussian processes of partial sums) and has the given covariance function. Furthermore, the existence of a continuous version can be obtained as for Brownian motion by employing the fourth moments of the increments, which for normal random variables can be computed from the variances (compare Theorem 21.9). In the following we will assume for the stochastic integral that such a continuous version is chosen.

In the special case, $f = \sum_{i=1}^n \alpha_i \mathbb{1}_{(t_{i-1}, t_i]}$ for some $n \in \mathbb{N}$ and $0 = t_0 < t_1 < \dots < t_n$ and $\alpha_1, \dots, \alpha_n \in \mathbb{R}$, we obtain

$$\int_0^1 f(s) dW_s = \sum_{i=1}^n \alpha_i (W_{t_i} - W_{t_{i-1}})$$

Exercise 21.5.1 Use the representation of Brownian motion $(W_t)_{t \in [0,1]}$ as a random linear combination of the Schauder functions $(B_{n,k})$ to show that the Brownian bridge $Y = (Y_t)_{t \in [0,1]} = (W_t - tW_1)_{t \in [0,1]}$ is a continuous, Gaussian process with covariance function $\text{Cov}[Y_t, Y_s] = (s \wedge t) - st$. Further, show that

$$\mathbf{P}_Y = \lim_{\varepsilon \downarrow 0} \mathbf{P}[W \in \cdot \mid W_1 \in (-\varepsilon, \varepsilon)].$$

Exercise 21.5.2 (Compare Example 8.32) Fix $T \in (0, 1)$. Use an orthonormal basis $b_{0,1}, (c_{n,k}), (d_{n,k})$ of suitably modified Haar functions (such that the $c_{n,k}$ have support $[0, T]$ and the $d_{n,k}$ have support $[T, 1]$) to show that a regular conditional distribution of W_T given W_1 is defined by

$$\mathbf{P}[W_T \in \cdot \mid W_1 = x] = \mathcal{N}_{T,x,T}$$

Exercise 21.5.3 Define $Y := (Y_t)_{t \in [0,1]}$ by $Y_1 = 0$ and

$$Y_t = (1-t) \int_0^t (1-s)^{-1} dW_s \quad \text{for } t \in [0, 1]$$

Show that Y is a Brownian bridge.

Hint: Show that Y is a continuous Gaussian process with the correct covariance function. In particular, it has to be shown that $\lim_{t \uparrow 1} Y_t = 0$ almost surely.

Exercise 21.5.4 Let $d \in \mathbb{N}$. Use a suitable orthonormal basis on $[0, 1]^d$ to show:

(i) There is a Gaussian process $(W_t)_{t \in [0,1]^d}$ with covariance function

$$\text{Cov}[W_t, W_s] = \prod_{i=1}^d (t_i \wedge s_i)$$

(ii) There is a modification of W such that $t \mapsto W_t$ is almost surely continuous (see Remark 21.7).

A process W with properties (i) and (ii) is called a Brownian sheet.

Exercise 21.5.5 Consider the coefficients $(A_n)_{n \in \mathbb{N}_0}$ of the Fourier basis of the construction of Brownian motion. Show the following statements:

- (i) $\sum_{n=0}^{\infty} A_n^2 < \infty$ almost surely.
- (ii) $\sum_{n=0}^{\infty} |A_n| = \infty$ almost surely.
- (iii) $\sum_{k=0}^n A_k, n \in \mathbb{N}$ converges almost surely.

Hint: Kolmogorov's three-series theorem (Theorem 15.50).

Exercise 21.5.6 Let $t \in (0, 1)$ and $f_0(x) := t$ as well as

$$f_n(x) := \frac{2 \sin(n\pi t)}{n\pi} \cos(n\pi x) \quad \text{for } n \in \mathbb{N}, x \in [0, 1]$$

Show that $\sum_{n=0}^{\infty} f_n(x) = \mathbb{1}_{[0,t]}(x)$ for $x \in (0, 1) \setminus \{t\}$.

5.21.6 21.6 The Space $c([0, \infty))$

Are functionals that depend on the whole path of a Brownian motion measurable? For example, is $\sup\{X_t, t \in [0, 1]\}$ measurable? For general stochastic processes, this is false since the supremum depends on more than countably many coordinates. However, for processes with continuous paths, this is true, as we will show in this section in a somewhat more general framework.

We may consider Brownian motion as the canonical process on the space $\Omega := C([0, \infty))$ of continuous paths.

We start by collecting some properties of the space $\Omega = C([0, \infty)) \subset \mathbb{R}^{[0, \infty)}$. Define the evaluation map

$$X_t : \Omega \rightarrow \mathbb{R}, \quad \omega \mapsto \omega(t), \quad (21.26)$$

that is, the restriction of the canonical projection $\mathbb{R}^{[0, \infty)} \rightarrow \mathbb{R}$ to Ω .

For $f, g \in C([0, \infty))$ and $n \in \mathbb{N}$, let $d_n(f, g) := \left\| (f - g)|_{[0, n]} \right\|_\infty \wedge 1$ and

$$d(f, g) = \sum_{n=1}^{\infty} 2^{-n} d_n(f, g). \quad (21.27)$$

Theorem 21.30 d is a complete metric on $\Omega := C([0, \infty))$ that induces the topology of uniform convergence on compact sets. The space (Ω, d) is separable and hence Polish.

Proof Clearly, every d_n is a complete metric on $(C([0, n]), \|\cdot\|_\infty)$. Thus, for every Cauchy sequence (f_N) in (Ω, d) and every $n \in \mathbb{N}$, there exists a $g_n \in \Omega$ with $d_n(f_N, g_n) \xrightarrow{N \rightarrow \infty} 0$. Evidently, $g_n(x) = g_m(x)$ for every $x \leq m \wedge n$; hence there exists a $g \in \Omega$ with $g(x) = g_n(x)$ for every $x \leq n$ for every $n \in \mathbb{N}$. Hence, clearly, $d(f_N, g) \xrightarrow{N \rightarrow \infty} 0$, and thus d is complete.

The set of polynomials with rational coefficients is countable and by the Weierstraß theorem, it is dense in any $(C([0, n]), \|\cdot\|_\infty)$; hence it is dense in (Ω, d) .

Theorem 21.31 With respect to the Borel σ -algebra $\mathcal{B}(\Omega, d)$, the canonical projections $X_t, t \in [0, \infty)$ are measurable. On the other hand, the X_t generate $\mathcal{B}(\Omega, d)$. Hence

$$(\mathcal{B}(\mathbb{R}))^{\otimes [0, \infty)}|_\Omega = \sigma(X_t, t \in [0, \infty)) = \mathcal{B}(\Omega, d)$$

Proof The first equation holds by definition. For the second one, we must show the mutual inclusions.

" $\subset$ " Clearly, every $X_t : \Omega \rightarrow \mathbb{R}$ is continuous and hence $(\mathcal{B}(\Omega, d) - \mathcal{B}(\mathbb{R}))$ measurable. Thus $\sigma(X_t, t \in [0, \infty)) \subset \mathcal{B}(\Omega, d)$.

" $\supset$ " We have to show that open subsets of (Ω, d) are in $\mathcal{A} := (\mathcal{B}(\mathbb{R}))^{\otimes [0, \infty)}$. Since (Ω, d) is separable (Theorem 21.30), every open set is a countable union of ε -balls. Hence it suffices to show that for fixed $\omega_0 \in \Omega$, the map $\omega \mapsto d(\omega_0, \omega)$ is $\mathcal{A}$ -measurable. To this end it is enough to show that for any $n \in \mathbb{N}$, the map $\omega \mapsto Y_n(\omega) := d_n(\omega_0, \omega)$ (see (21.27)) is $\mathcal{A}$ -measurable. However, the map

$$\omega \mapsto Z_t(\omega) := |X_t(\omega) - X_t(\omega_0)| \wedge 1$$

is $\mathcal{A}$ -measurable. Since each ω is continuous, Y_n is a countable supremum

$$Y_n = \sup_{t \in [0, n] \cap \mathbb{Q}} Z_t$$

and is hence $\mathcal{A}$ -measurable.

In the following, let $\mathcal{A} := \sigma(X_t, t \in [0, \infty))$.

Corollary 21.32 The map $F_1 : \Omega \rightarrow [0, \infty), \omega \mapsto \sup\{\omega(t) : t \in [0, 1]\}$ is $\mathcal{A}$ measurable.

Proof F_1 is continuous with respect to d and hence $\mathcal{B}(\Omega, d)$ -measurable.

If B is a Brownian motion (on some probability space $(\tilde{\Omega}, \tilde{\mathcal{A}}, \tilde{\mathbf{P}})$), then there exists an $\bar{\Omega} \in \tilde{\mathcal{A}}$ with $\tilde{\mathbf{P}}[\bar{\Omega}] = 1$ and $B(\omega) \in C([0, \infty))$ for every $\omega \in \bar{\Omega}$. Let $\bar{\mathcal{A}} = \tilde{\mathcal{A}}|_{\bar{\Omega}}$ and $\bar{\mathbf{P}} = \tilde{\mathbf{P}}|_{\bar{\mathcal{A}}}$. Then $B : \bar{\Omega} \rightarrow C([0, \infty))$ is measurable with respect to $(\bar{\mathcal{A}}, \mathcal{A})$. With respect to the image measure $\mathbf{P} = \bar{\mathbf{P}} \circ B^{-1}$ on $\Omega = C([0, \infty))$, the canonical process $X = (X_t, t \in [0, \infty))$ on $C([0, \infty))$ is a Brownian motion.

Definition 21.33 Let $\mathbf{P}$ be the probability measure on $\Omega = C([0, \infty))$ with respect to which the canonical process X is a Brownian motion. Then $\mathbf{P}$ is called the Wiener measure. The

triple $(\Omega, \mathcal{A}, \mathbf{P})$ is called the Wiener space, and X is called the canonical Brownian motion or the Wiener process.

Remark 21.34 Sometimes we want a Brownian motion to start not at $X_0 = 0$ but at an arbitrary point x . Denote by $\mathbf{P}_x$ that measure on $C([0, \infty))$ for which $\tilde{X} = (X_t - x, t \in [0, \infty))$ is a Brownian motion (with $\tilde{X}_0 = 0$).

Exercise 21.6.1 Show that the map

$$F_\infty : \Omega \rightarrow [0, \infty], \quad \omega \mapsto \sup\{\omega(t) : t \in [0, \infty)\}$$

is $\mathcal{A}$ -measurable.

5.21.7 21.7 Convergence of Probability Measures on $C([0, \infty))$

Let X and $(X^n)_{n \in \mathbb{N}}$ be random variables with values in $C([0, \infty))$ (i.e., continuous stochastic processes) with distributions $\mathbf{P}_X$ and $(\mathbf{P}_{X^n})_{n \in \mathbb{N}}$.

Definition 21.35 We say that the finite-dimensional distributions of (X^n) converge to those of X if, for every $k \in \mathbb{N}$ and $t_1, \dots, t_k \in [0, \infty)$, we have

$$(X_{t_1}^n, \dots, X_{t_k}^n) \xrightarrow{n \rightarrow \infty} (X_{t_1}, \dots, X_{t_k}).$$

In this case, we write

$$X^n \xrightarrow[\text{fdd}]{n \rightarrow \infty} X \quad \text{or} \quad \mathbf{P}_{X^n} \xrightarrow[\text{fdd}]{n \rightarrow \infty} \mathbf{P}_X.$$

Lemma 21.36 $P_n \xrightarrow[\text{fdd}]{n \rightarrow \infty} P$ and $P_n \xrightarrow[\text{fdd}]{n \rightarrow \infty} Q$ imply $P = Q$.

Proof By Theorem 14.12(iii), the finite-dimensional distributions determine P uniquely.

Theorem 21.37 Weak convergence in $\mathcal{M}_1(\Omega, d)$ implies fdd-convergence:

$$P_n \xrightarrow{n \rightarrow \infty} P \implies P_n \xrightarrow[\text{fdd}]{n \rightarrow \infty} P.$$

Proof Let $k \in \mathbb{N}$ and $t_1, \dots, t_k \in [0, \infty)$. The map

$$\varphi : C([0, \infty)) \rightarrow \mathbb{R}^k, \quad \omega \mapsto (\omega(t_1), \dots, \omega(t_k))$$

is continuous. By the continuous mapping theorem (Theorem 13.25 on p. 257), we have $P_n \circ \varphi^{-1} \xrightarrow{n \rightarrow \infty} P \circ \varphi^{-1}$; hence $P_n \xrightarrow[\text{fdd}]{n \rightarrow \infty} P$.

The converse statement in the preceding theorem does not hold. However, we still have the following.

Theorem 21.38 Let $(P_n)_{n \in \mathbb{N}}$ and P be probability measures on $C([0, \infty))$. Then the following are equivalent:

- (i) $P_n \xrightarrow{n \rightarrow \infty} P$ and $(P_n)_{n \in \mathbb{N}}$ is tight.
- (ii) $P_n \xrightarrow{\text{fdd}} P$ weakly.

Proof "(ii) $\implies$ (i)" This is a direct consequence of Prohorov's theorem (Theorem 13.29 with $E = C([0, \infty))$).

"(i) $\implies$ (ii)" By Prohorov's theorem, $(P_n)_{n \in \mathbb{N}}$ is relatively sequentially compact. Let Q be a limit point for $(P_{n_k})_{k \in \mathbb{N}}$ along some subsequence (n_k) . Then $P_{n_k} \xrightarrow{\text{fdd}} Q$, $k \rightarrow \infty$. By Lemma 21.36, we have $P = Q$.

Next we derive a useful criterion for tightness of sets $\{P_n\} \subset \mathcal{M}_1(C([0, \infty)))$. We start by recalling the Arzelà-Ascoli characterization of relatively compact sets in $C([0, \infty))$ (see, e.g., [37, Theorem 2.4.7] or [173, Theorem III.3]).

For $N, \delta > 0$ and $\omega \in C([0, \infty))$, let

$$V^N(\omega, \delta) := \sup\{|\omega(t) - \omega(s)| : |t - s| \leq \delta, s, t \leq N\}.$$

Theorem 21.39 (Arzelà-Ascoli) A set $A \subset C([0, \infty))$ is relatively compact if and only if the following two conditions hold.

- (i) $\{\omega(0), \omega \in A\} \subset \mathbb{R}$ is bounded.
- (ii) For every N , we have $\lim_{\delta \downarrow 0} \sup_{\omega \in A} V^N(\omega, \delta) = 0$.

Theorem 21.40 A family $(P_i, i \in I)$ of probability measures on $C([0, \infty))$ is weakly relatively compact if and only if the following two conditions hold.

- (i) $(P_i \circ X_0^{-1}, i \in I)$ is tight; that is, for every $\varepsilon > 0$, there is a $K > 0$ such that

$$P_i(\{\omega : |\omega(0)| > K\}) \leq \varepsilon \quad \text{for all } i \in I. \quad (21.28)$$

- (ii) For all $\eta, \varepsilon > 0$ and $N \in \mathbb{N}$, there is a $\delta > 0$ such that

$$P_i(\{\omega : V^N(\omega, \delta) > \eta\}) \leq \varepsilon \quad \text{for all } i \in I. \quad (21.29)$$

Proof " " By Prohorov's theorem (Theorem 13.29), weak relative compactness of $(P_i, i \in I)$ implies tightness of this family. Thus, for every $\varepsilon > 0$, there exists a compact set $A \subset C([0, \infty))$ with $P_i(A) > 1 - \varepsilon$ for every $i \in I$. Using the Arzelà-Ascoli characterization of the compactness of A , we infer (i) and (ii).

"." Now assume that (i) and (ii) hold. Then, for $\varepsilon > 0$ and $k, N \in \mathbb{N}$, choose numbers K_ε and $\delta_{N,k,\varepsilon}$ such that

$$\sup_{i \in I} P_i(\{\omega : |\omega(0)| > K_\varepsilon\}) \leq \frac{\varepsilon}{2}$$

and

$$\sup_{i \in I} P_i\left(\left\{\omega : V^N(\omega, \delta_{N,k,\varepsilon}) > \frac{1}{k}\right\}\right) \leq 2^{-N-k-1}\varepsilon.$$

Define

$$C_{N,\varepsilon} = \left\{\omega : |\omega(0)| \leq K_\varepsilon, V^N(\omega, \delta_{N,k,\varepsilon}) \leq \frac{1}{k} \text{ for all } k \in \mathbb{N}\right\}.$$

By the Arzelà-Ascoli theorem, $C_\varepsilon := \bigcap_{N \in \mathbb{N}} C_{N,\varepsilon}$ is relatively compact in $C([0, \infty))$ and we have

$$P_i(C_\varepsilon^c) \leq \frac{\varepsilon}{2} + \sum_{k,N=1}^{\infty} P_i(\{\omega : V^N(\omega, \delta_{N,k,\varepsilon}) > 1/k\}) \leq \varepsilon \quad \text{for all } i \in I$$

Hence the claim follows.

Corollary 21.41 Let $(X_i, i \in I)$ and $(Y_i, i \in I)$ be families of random variables in $C([0, \infty))$. Assume that $(\mathbf{P}_{X_i}, i \in I)$ and $(\mathbf{P}_{Y_i}, i \in I)$ are tight. Then $(\mathbf{P}_{X_i+Y_i}, i \in I)$ is tight.

Proof Apply the triangle inequality in order to check (i) and (ii) in the preceding theorem.

The following is an important tool to check weak relative compactness.

Theorem 21.42 (Kolmogorov's criterion for weak relative compactness) Let $(X^i, i \in I)$ be a sequence of continuous stochastic processes. Assume that the following conditions are satisfied.

- (i) The family $(\mathbf{P}[X_0^i \in \cdot], i \in I)$ of initial distributions is tight.
- (ii) For any $N > 0$ there are numbers $C, \alpha, \beta > 0$ such that, for all $s, t \in [0, N]$ and every $i \in I$, we have

$$\mathbf{E}[|X_s^i - X_t^i|^\alpha] \leq C|s - t|^{\beta+1}.$$

Then the family $(\mathbf{P}_{X^i}, i \in I) = (\mathcal{L}[X^i], i \in I)$ of distributions of X^i is weakly relatively compact in $\mathcal{M}_1(C([0, \infty)))$.

Proof We check the conditions of Theorem 21.40. The first condition of Theorem 21.40 is exactly (i).

Let $N > 0$. By the Kolmogorov-Chentsov theorem (Theorem 21.6(ii)), for $\varepsilon > 0$ and $\gamma \in (0, \beta/\alpha)$, there exists a K such that, for every $i \in I$, we have

$$\mathbf{P} [|X_t^i - X_s^i| \leq K|t - s|^\gamma \text{ for all } s, t \in [0, N]] \geq 1 - \varepsilon.$$

This clearly implies (21.29) with $\delta = (\eta/K)^{1/\gamma}$. $\square$

5.21.8 21.8 Donsker's Theorem

Let $Y_1, Y_2, \dots$ be i.i.d. random variables with $\mathbf{E}[Y_1] = 0$ and $\text{Var}[Y_1] = \sigma^2 > 0$. For $t > 0$, let $S_t^n = \sum_{i=1}^{\lfloor nt \rfloor} Y_i$ and $\tilde{S}_t^n = \frac{1}{\sqrt{\sigma^2 n}} S_t^n$. By the central limit theorem, for $t > s \geq 0$, we have $\mathcal{L}[\tilde{S}_t^n - \tilde{S}_s^n] \xrightarrow{n \rightarrow \infty} \mathcal{N}_{0, t-s}$.

Let $B = (B_t, t \geq 0)$ be a Brownian motion. Then

$$\mathcal{L}[\tilde{S}_t^n - \tilde{S}_s^n] \xrightarrow{n \rightarrow \infty} \mathcal{L}[B_t - B_s] \quad \text{for any } t > s \geq 0.$$

For $N \in \mathbb{N}$ and $0 = t_0 < t_1 < \dots < t_N$, the random variables $\tilde{S}_{t_i}^n - \tilde{S}_{t_{i-1}}^n, i = 1, \dots, N$, are independent, and hence, we have

$$\mathcal{L}[(\tilde{S}_{t_1}^n - \tilde{S}_{t_0}^n, \dots, \tilde{S}_{t_N}^n - \tilde{S}_{t_{N-1}}^n)] \xrightarrow{n \rightarrow \infty} \mathcal{L}[(B_{t_1} - B_{t_0}, \dots, B_{t_N} - B_{t_{N-1}})].$$

We infer that

$$\mathcal{L}[(\tilde{S}_{t_1}^n, \dots, \tilde{S}_{t_N}^n)] \xrightarrow{n \rightarrow \infty} \mathcal{L}[(B_{t_1}, \dots, B_{t_N})]. \quad (21.30)$$

We now define $\bar{S}^n$ as $\tilde{S}^n$ but linearly interpolated:

$$\bar{S}_t^n = \frac{1}{\sqrt{\sigma^2 n}} \sum_{i=1}^{\lfloor nt \rfloor} Y_i + \frac{tn - \lfloor tn \rfloor}{\sqrt{\sigma^2 n}} Y_{\lfloor tn \rfloor + 1}. \quad (21.31)$$

Then, for $\varepsilon > 0$,

$$\begin{aligned} \mathbf{P} [|\tilde{S}_t^n - \bar{S}_t^n| > \varepsilon] &\leq \varepsilon^{-2} \mathbf{E} \left[(\tilde{S}_t^n - \bar{S}_t^n)^2 \right] \\ &\leq \frac{1}{\varepsilon^2 n} \frac{1}{\sigma^2} \mathbf{E} [Y_1^2] = \frac{1}{\varepsilon^2 n} \xrightarrow{n \rightarrow \infty} 0. \end{aligned}$$

By Slutsky's theorem (Theorem 13.18), we thus have convergence of the finitedimensional distributions to the Wiener measure $\mathbf{P}_W$:

$$\mathbf{P}_{\bar{S}^n} \xrightarrow[\text{fdd}]{n \rightarrow \infty} \mathbf{P}_W. \quad (21.32)$$

The aim of this section is to strengthen this convergence statement to weak convergence of probability measures on $C([0, \infty))$. The main theorem of this section is the functional central limit theorem, which goes back to Donsker [35]. Theorems

of this type are also called invariance principles since the limiting distribution is the same for all distributions Y_i with expectation 0 and the same variance.

Theorem 21.43 (Donsker's invariance principle) In the sense of weak convergence on $C([0, \infty))$, the distributions of $\bar{S}^n$ converge to the Wiener measure,

$$\mathcal{L}[\bar{S}^n] \xrightarrow{n \rightarrow \infty} \mathbf{P}_W. \quad (21.33)$$

Proof Owing to (21.32) and Theorem 21.38, it is enough to show that $(\mathcal{L}[\bar{S}^n], n \in \mathbb{N})$ is tight. To this end, we want to apply Kolmogorov's moment criterion. However, as in the proof of existence of Brownian motion, second moments are not enough; rather we need fourth moments in order that we can choose $\beta > 0$. Hence the strategy is to truncate the Y_i to obtain fourth moments.

For $K > 0$, define

$$\begin{aligned} Y_i^K &:= Y_i \mathbb{1}_{\{|Y_i| \leq K/2\}} - \mathbf{E}[Y_i \mathbb{1}_{\{|Y_i| \leq K/2\}}] \quad \text{and} \\ Z_i^K &:= Y_i - Y_i^K \quad \text{for } i \in \mathbb{N}. \end{aligned}$$

Then $\mathbf{E}[Y_i^K] = \mathbf{E}[Z_i^K] = 0$ and $\text{Var}[Z_i^K] \xrightarrow{K \rightarrow \infty} 0$ as well as $\text{Var}[Y_i^K] \leq \sigma^2, i \in \mathbb{N}$. Clearly, $|Y_i^K| \leq K$ for every i . Define

$$T_n^K := \sum_{i=1}^n Y_i^K \quad \text{and} \quad U_n^K := \sum_{i=1}^n Z_i^K \quad \text{for } n \in \mathbb{N}.$$

Let $\bar{T}_t^{K,n}$ and $\bar{U}_t^{K,n}$ be the linearly interpolated versions of

$$\tilde{T}_t^{K,n} := \frac{1}{\sqrt{\sigma^2 n}} T_{[nt]}^K \quad \text{and} \quad \tilde{U}_t^{K,n} := \frac{1}{\sqrt{\sigma^2 n}} U_{[nt]}^K \quad \text{for } t \geq 0.$$

Evidently, $\bar{S}^n = \bar{T}^{K,n} + \bar{U}^{K,n}$. By Corollary 21.41, it is enough to show that, for a sequence $(K_n)_{n \in \mathbb{N}}$ (chosen later), the families $(\mathcal{L}[\bar{U}^{K_n,n}], n \in \mathbb{N})$ and $(\mathcal{L}[\bar{T}^{K_n,n}], n \in \mathbb{N})$ are tight.

We consider first the remainder term. As U^K is a martingale, Doob's inequality (Theorem 11.2) yields

$$\mathbf{P} \left[\sup_{l=1, \dots, n} |U_l^K| > \varepsilon \sqrt{n} \right] \leq \varepsilon^{-2} \text{Var}[Z_1^K] \quad \text{for every } \varepsilon > 0.$$

Now, if $K_n \uparrow \infty, n \rightarrow \infty$, then for every $N > 0$, we have

$$\mathbf{P} \left[\sup_{t \in [0, N]} |\bar{U}_t^{K_n,n}| > \varepsilon \right] \leq \frac{N}{\varepsilon^2 \sigma^2} \text{Var}[Z_1^{K_n}] \xrightarrow{n \rightarrow \infty} 0$$

hence $\bar{U}^{K_n,n} \xrightarrow{n \rightarrow \infty} 0$ in $C([0, \infty))$. In particular, $(\mathcal{L}[\bar{U}^{K_n,n}], n \in \mathbb{N})$ is tight.

Next, for $N > 0$ and $s, t \in [0, N]$, we compute the fourth moments of the differences $\bar{T}_{t+s}^{K_n,n} - \bar{T}_s^{K_n,n}$ for the main term. In the following, let $K_n = n^{1/4}$. Fix $n \in \mathbb{N}$. We distinguish two cases:

Case 1: $t < n^{-1}$. Let $k := \lfloor (t+s)n \rfloor$. If $sn \geq k$, then

$$\bar{T}_{t+s}^{K_n,n} - \bar{T}_s^{K_n,n} = \frac{tn}{\sqrt{n\sigma^2}} Y_{k+1}^{K_n}.$$

If $sn < k$, then

$$\bar{T}_{t+s}^{K_n,n} - \bar{T}_s^{K_n,n} = \frac{1}{\sqrt{n\sigma^2}} ((t+s)n - k) Y_{k+1}^{K_n} + (k - sn) Y_k^{K_n}.$$

In either case, we have

$$\left| \bar{T}_{t+s}^{K_n, n} - \bar{T}_s^{K_n, n} \right| \leq \frac{t\sqrt{n}}{\sigma} (|Y_k^{K_n}| + |Y_{k+1}^{K_n}|),$$

hence

$$\begin{aligned} \mathbf{E} \left[\left(\bar{T}_{t+s}^{K_n, n} - \bar{T}_s^{K_n, n} \right)^4 \right] &\leq \frac{n^2 t^4}{\sigma^4} (2K_n)^2 \mathbf{E} \left[(|Y_1^{K_n}| + |Y_2^{K_n}|)^2 \right] \\ &\leq \frac{16n^{5/2} t^4}{\sigma^4} \text{Var} [Y_1^{K_n}] \leq \frac{16}{\sigma^2} t^{3/2} \end{aligned} \quad (21.34)$$

Case 2: $t \geq n^{-1}$. Using the binomial theorem, we get (note that the mixed terms with odd moments vanish since $\mathbf{E} [Y_1^{K_n}] = 0$)

$$\begin{aligned} \mathbf{E} \left[(T_n^{K_n})^4 \right] &= n \mathbf{E} \left[(Y_1^{K_n})^4 \right] + 3n(n-1) \mathbf{E} \left[(Y_1^{K_n})^2 \right]^2 \\ &\leq n K_n^2 \sigma^2 + 3n(n-1) \sigma^4. \end{aligned} \quad (21.35)$$

Note that, for independent real random variables X, Y with $\mathbf{E}[X] = \mathbf{E}[Y] = 0$ and $\mathbf{E}[X^4], \mathbf{E}[Y^4] < \infty$ and for $a \in [-1, 1]$, we have

$$\begin{aligned} \mathbf{E} [(aX + Y)^4] &= a^4 \mathbf{E} [X^4] + 6a^2 \mathbf{E} [X^2] \mathbf{E} [Y^2] + \mathbf{E} [Y^4] \\ &\leq \mathbf{E} [X^4] + 6 \mathbf{E} [X^2] \mathbf{E} [Y^2] + \mathbf{E} [Y^4] = \mathbf{E} [(X + Y)^4]. \end{aligned}$$

We apply this twice (with $a = \lceil (t+s)n \rceil - (t+s)n$ and $a = sn - \lfloor sn \rfloor$) and obtain (using the estimate $\lceil (t+s)n \rceil - \lfloor sn \rfloor \leq tn + 2 \leq 3tn$) from (21.35) (since $t \leq N$)

$$\begin{aligned} \mathbf{E} \left[\left(\bar{T}_{t+s}^{K_n, n} - \bar{T}_s^{K_n, n} \right)^4 \right] &\leq n^{-2} \sigma^{-4} \mathbf{E} \left[\left(T_{\lceil (t+s)n \rceil}^{K_n} - T_{\lfloor sn \rfloor}^{K_n} \right)^4 \right] \\ &= n^{-2} \sigma^{-4} \mathbf{E} \left[\left(T_{\lceil (t+s)n \rceil - \lfloor sn \rfloor}^{K_n} \right)^4 \right] \\ &\leq \frac{3tn K_n^2}{n^2 \sigma^2} + 18t^2 = \frac{3}{\sigma^2} tn^{-1/2} + 18t^2 \\ &\leq \frac{3}{\sigma^2} t^{3/2} + 18t^2 \leq \left(\frac{3}{\sigma^2} + 18\sqrt{N} \right) t^{3/2} \end{aligned} \quad (21.36)$$

By (21.34) and (21.36), for every $N > 0$, there exists a $C = C(N, \sigma^2)$ such that, for every $n \in \mathbb{N}$ and all $s, t \in [0, N]$, we have

$$\mathbf{E} \left[\left(\bar{T}_{t+s}^{K_n, n} - \bar{T}_s^{K_n, n} \right)^4 \right] \leq C t^{3/2}$$

Hence, by Kolmogorov's moment criterion (Theorem 21.42 with $\alpha = 4$ and $\beta = 1/2$), $(\mathcal{L} [\bar{T}^{K_n, n}], n \in \mathbb{N})$ is tight in $\mathcal{M}_1(C([0, \infty)))$.

Exercise 21.8.1 Let $X_1, X_2, \dots$ be i.i.d. random variables with continuous distribution function F . Let $G_n : [0, 1] \rightarrow \mathbb{R}, t \mapsto n^{-1/2} \sum_{i=1}^n (\mathbb{1}_{[0, t]}(F(X_i)) - t)$ and $M_n := \|G_n\|_\infty$. Further, let $M = \sup_{t \in [0, 1]} |B_t|$, where B is a Brownian bridge.

(i) Show that $\mathbf{E} [G_n(t)] = 0$ and $\mathbf{Cov} [G_n(s), G_n(t)] = s \wedge t - st$ for $s, t \in [0, 1]$.

(ii) Show that $\mathbf{E} [(G_n(t) - G_n(s))^4] \leq C((t-s)^2 + |t-s|/n)$ for some $C > 0$.

(iii) Conclude that a suitable continuous version of G_n converges weakly to B . For example, choose

$$H_n(t) = n^{-1/2} \sum_{i=1}^n (h_n(F(X_i) - t) - g_n(t)),$$

where h_n is a suitable smoothed version of $\mathbb{1}_{(-\infty, 0]}$, for example, $h_n(s) = 1 - (s/\varepsilon_n \vee 0) \wedge 1$ for some sequence $\varepsilon_n \downarrow 0$, and $g_n(t) := \int_0^1 h_n(t - u) du$.

(iv) Finally, show that $M_n \xrightarrow{n \rightarrow \infty} M$.

Remark: The distribution of M can be expressed by the Kolmogorov-Smirnov formula ([101] and [156]; see, e.g., [133])

$$\mathbf{P}[M > x] = 2 \sum_{n=1}^{\infty} (-1)^{n-1} e^{-2n^2 x^2} \quad (21.37)$$

Compare (21.20). Using the statistic M_n , one can test if random variables of a known distribution are independent. Let $X_1, X_2, \dots$ and $\tilde{X}_1, \tilde{X}_2, \dots$ be independent random variables with unknown continuous distribution functions F and $\tilde{F}$ and with empirical distribution functions F_n and $\tilde{F}_n$. Further, let

$$D_n := \sup_{t \in \mathbb{R}} |F_n(t) - \tilde{F}_n(t)|$$

Under the assumption that $F = \tilde{F}$ holds, $\sqrt{n/2} D_n$ converges in distribution to M . This fact is the basis for nonparametric tests on the equality of distributions.

5.21.9 21.9 Pathwise Convergence of Branching Processes*

In this section, we investigate the convergence of rescaled Galton-Watson processes (branching processes). As for sums of independent random variables, we first show convergence for a fixed time point to the distribution of a certain limiting process. The next step is to show convergence of finite-dimensional distributions. Finally, using Kolmogorov's moment criterion for tightness, we show convergence in the path space $C([0, \infty))$.

Consider a Galton-Watson process $(Z_n)_{n \in \mathbb{N}_0}$ with geometric offspring distribution

$$p(k) = 2^{-k-1} \quad \text{for } k \in \mathbb{N}_0$$

That is, let $X_{n,i}, n, i \in \mathbb{N}_0$ be i.i.d. random variables on $\mathbb{N}_0$ with $\mathbf{P}[X_{n,i} = k] = p(k), k \in \mathbb{N}_0$, and based on the initial state Z_0 define inductively

$$Z_{n+1} = \sum_{i=1}^{Z_n} X_{n,i}$$

Thus Z is a Markov chain with transition probabilities $p(i, j) = p^{*i}(j)$, where p^{*i} is the i th convolution power of p . In other words, if $Z, Z^1, \dots, Z^i$ are independent copies of our Galton-Watson process, with $Z_0 = i$ and $Z_0^1 = \dots = Z_0^i = 1$, then

$$Z \stackrel{\mathcal{D}}{=} Z^1 + \dots + Z^i \quad (21.38)$$

We consider now the probability generating function of $X_{1,1}, \psi^{(1)}(s) := \psi(s) := \mathbf{E}[s^{X_{1,1}}], s \in [0, 1]$. Denote by $\psi^{(n)} := \psi^{(n-1)} \circ \psi$ its n th iterate for $n \in \mathbb{N}$. Then, by Lemma 3.10,

$$\mathbf{E}_i[s^{Z_n}] = \mathbf{E}_1[s^{Z_n}]^i = (\psi^{(n)}(s))^i.$$

For the geometric offspring distribution, $\psi^{(n)}$ can be computed explicitly.

Lemma 21.44 For the branching process with critical geometric offspring distribution, the n th iterate of the probability generating function is

$$\psi^{(n)}(s) = \frac{n - (n-1)s}{n+1 - ns}.$$

Proof Compute

$$\psi(s) = \sum_{k=0}^{\infty} 2^{-k-1} s^k = \frac{1}{-s+2}$$

In order to compute the iterated function, first consider general linear rational functions of the form $f(x) = \frac{ax+b}{cx+d}$. For such f , define the matrix $M_f = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$. For two linear rational functions f and g , we have $M_{f \circ g} = M_f \cdot M_g$. The powers of M are easy to compute:

$$M_\psi = \begin{pmatrix} 0 & 1 \\ -1 & 2 \end{pmatrix}, \quad M_\psi^2 = \begin{pmatrix} -1 & 2 \\ -2 & 3 \end{pmatrix}, \quad M_\psi^3 = \begin{pmatrix} -2 & 3 \\ -3 & 4 \end{pmatrix},$$

and inductively

$$M_\psi^n = \begin{pmatrix} -(n-1) & n \\ -n & n+1 \end{pmatrix}.$$

If we let $s = e^{-\lambda}$, then we get the Laplace transform of Z_n ,

$$\mathbf{E}_i [e^{-\lambda Z_n}] = \psi^{(n)}(e^{-\lambda})^i.$$

By Example 6.29, we can compute the moments of Z_n by differentiating the Laplace transform. That is, we obtain the following lemma.

Lemma 21.45 The moments of Z_n are

$$\mathbf{E}_i [Z_n^k] = (-1)^k \frac{d^k}{d\lambda^k} \left(\psi^{(n)}(e^{-\lambda})^i \right) \Big|_{\lambda=0}. \quad (21.39)$$

In particular, the first six moments are

$$\begin{aligned} \mathbf{E}_i [Z_n] &= i \\ \mathbf{E}_i [Z_n^2] &= 2in + i^2 \\ \mathbf{E}_i [Z_n^3] &= 6in^2 + 6i^2n + i^3 \\ \mathbf{E}_i [Z_n^4] &= 24in^3 + 36i^2n^2 + (12i^3 + 2i)n + i^4 \\ \mathbf{E}_i [Z_n^5] &= 120in^4 + 240i^2n^3 + (120i^3 + 30i)n^2 + (20i^4 + 10i^2)n + i^5, \\ \mathbf{E}_i [Z_n^6] &= 720in^5 + 1800i^2n^4 + (1200i^3 + 360i)n^3 \\ &\quad + (300i^4 + 240i^2)n^2 + (30i^5 + 30i^3 + 2i)n + i^6. \end{aligned} \quad (21.40)$$

Hence, Z is a martingale, and the first six centered moments are

$$\begin{aligned}
\mathbf{E}_i [(Z_n - i)^2] &= 2in, \\
\mathbf{E}_i [(Z_n - i)^3] &= 6in^2, \\
\mathbf{E}_i [(Z_n - i)^4] &= 24in^3 + 12i^2n^2 + 2in, \\
\mathbf{E}_i [(Z_n - i)^5] &= 120in^4 + 120i^2n^3 + 30in^2, \\
\mathbf{E}_i [(Z_n - i)^6] &= 720in^5 + 1080i^2n^4 + (120i^3 + 360i)n^3 \\
&\quad + 60i^2n^2 + 2in.
\end{aligned} \tag{21.41}$$

Proof The exact formulas for the first six moments are obtained by tenaciously computing the right-hand side of (21.39).

Now consider the following rescaling: Fix $x \geq 0$, start with $Z_0 = \lfloor nx \rfloor$ individuals and consider $\tilde{Z}_t^n := \frac{Z_{\lfloor tn \rfloor}}{n}$ for $t \geq 0$. We abbreviate

$$\mathcal{L}_x [\tilde{Z}^n] := \mathcal{L}_{\lfloor nx \rfloor} \left[(n^{-1} Z_{\lfloor nt \rfloor})_{t \geq 0} \right]. \tag{21.42}$$

Evidently, $\mathbf{E}_x [\tilde{Z}_t^n] = \frac{\lfloor nx \rfloor}{n} \leq x$ for every n ; hence $(\mathcal{L}_x [\tilde{Z}_t^n], n \in \mathbb{N})$ is tight. By considering Laplace transforms, we obtain that, for every $\lambda \geq 0$, the sequence of distributions converges:

$$\begin{aligned}
\lim_{n \rightarrow \infty} \mathbf{E}_x [e^{-\lambda \tilde{Z}_t^n}] &= \lim_{n \rightarrow \infty} (\psi^{(\lfloor tn \rfloor)} (e^{-\lambda/n}))^{nx} \\
&= \lim_{n \rightarrow \infty} \left(\frac{nt - (nt - 1)e^{-\lambda/n}}{nt + 1 - nte^{-\lambda/n}} \right)^{nx} \\
&= \lim_{n \rightarrow \infty} \left(1 - \frac{1 - e^{-\lambda/n}}{n(1 - e^{-\lambda/n})t + 1} \right)^{nx} \\
&= \exp \left(- \lim_{n \rightarrow \infty} \frac{xn(1 - e^{-\lambda/n})}{n(1 - e^{-\lambda/n})t + 1} \right) \\
&= \exp \left(- \frac{\lambda}{\lambda + 1/t} (x/t) \right) := \psi_t(\lambda)^x
\end{aligned} \tag{21.43}$$

However, the function ψ_t^x is the Laplace transform of the compound Poisson distribution $\text{CPoi}_{(x/t)} \exp_{1/t}$ (see Definition 16.3).

Consider the stochastic kernel $\kappa_t(x, dy) := \text{CPoi}_{(x/t)} \exp_{1/t}(dy)$. This is the kernel on $[0, \infty)$ whose Laplace transform is given by

$$\int_0^\infty \kappa_t(x, dy) e^{-\lambda y} = \psi_t(\lambda)^x \tag{21.44}$$

Lemma 21.46 $(\kappa_t)_{t \geq 0}$ is a Markov semigroup and there exists a Markov process $(Y_t)_{t \geq 0}$ with transition kernels $\mathbf{P}_x [Y_t \in dy] = \kappa_t(x, dy)$.

Proof It suffices to check that the Chapman-Kolmogorov equation $\kappa_t \cdot \kappa_s = \kappa_{s+t}$ holds. We compute the Laplace transform for these kernels. For $\lambda \geq 0$, applying (21.44) twice yields

$$\begin{aligned}
\iint \kappa_t(x, dy) \kappa_s(y, dz) e^{-\lambda z} &= \int \kappa_t(x, dy) \exp\left(-\frac{\lambda y}{\lambda s + 1}\right) \\
&= \exp\left(-\frac{\frac{\lambda}{\lambda s + 1} x}{\frac{\lambda}{\lambda s + 1} t + 1}\right) \\
&= \exp\left(-\frac{\lambda x}{\lambda(t + s) + 1}\right) \\
&= \int \kappa_{t+s}(x, dz) e^{-\lambda z}
\end{aligned}$$

Next we show that Y has a continuous version. To this end, we compute some of its moments and then use the Kolmogorov-Chentsov theorem (Theorem 21.6).

Lemma 21.47 The first k moments of Y_t can be computed by differentiating the Laplace transform,

$$\mathbf{E}_x [Y_t^k] = (-1)^k \frac{d^k}{d\lambda^k} (\psi(\lambda)^x) \big|_{\lambda=0}$$

where

$$\psi_t(\lambda) = \exp\left(-\frac{\lambda}{\lambda t + 1}\right).$$

In particular, we have

$$\begin{aligned}
\mathbf{E}_x [Y_t] &= x \\
\mathbf{E}_x [Y_t^2] &= 2xt + x^2 \\
\mathbf{E}_x [Y_t^3] &= 6xt^2 + 6x^2t + x^3 \\
\mathbf{E}_x [Y_t^4] &= 24xt^3 + 36x^2t^2 + 12x^3t + x^4 \\
\mathbf{E}_x [Y_t^5] &= 120xt^4 + 240x^2t^3 + 120x^3t^2 + 20x^4t + x^5 \\
\mathbf{E}_x [Y_t^6] &= 720xt^5 + 1800x^2t^4 + 1200x^3t^3 + 300x^4t^2 + 30x^5t + x^6
\end{aligned} \tag{21.45}$$

Hence Y is a martingale, and the first centered moments are

$$\begin{aligned}
\mathbf{E}_x [(Y_t - x)^2] &= 2xt, \\
\mathbf{E}_x [(Y_t - x)^3] &= 6xt^2, \\
\mathbf{E}_x [(Y_t - x)^4] &= 24xt^3 + 12x^2t^2, \\
\mathbf{E}_x [(Y_t - x)^5] &= 120xt^4 + 120x^2t^3, \\
\mathbf{E}_x [(Y_t - x)^6] &= 720xt^5 + 1080x^2t^4 + 120x^3t^3.
\end{aligned} \tag{21.46}$$

Theorem 21.48 There is a continuous version of the Markov process Y with transition kernels $(\kappa_t)_{t \geq 0}$ given by (21.44). This version is called Feller's (continuous) branching diffusion.

See Fig. 26.4 for a computer simulation of Feller's branching diffusion.

Proof For fixed $N > 0$ and $s, t \in [0, N]$, we have

$$\begin{aligned}
\mathbf{E}_x [(Y_{t+s} - Y_s)^4] &= \mathbf{E}_x [\mathbf{E}_{Y_s} [(Y_t - Y_0)^4]] \\
&= \mathbf{E}_x [24Y_s t^3 + 12Y_s^2 t^2] \\
&= 24xt^3 + 12(2sx + x^2) t^2 \\
&\leq (48Nx + 12x^2) t^2.
\end{aligned}$$

Thus Y satisfies the condition of Theorem 21.6 (Kolmogorov-Chentsov) with $\alpha = 4$ and $\beta = 1$.

Remark 21.49

- (i) By using higher moments, it can be shown that the paths of Y are Höldercontinuous of any order $\gamma \in (0, \frac{1}{2})$.
- (ii) It can be shown that Y is the (unique strong) solution of the stochastic (Itô-) differential equation (see Examples 26.11 and 26.31)

$$dY_t = \sqrt{2Y_t} dW_t, \quad (21.47)$$

where W is a Brownian motion.

Theorem 21.50 We have $\mathcal{L}_x [\tilde{Z}^n] \xrightarrow[\text{fdd}]{n \rightarrow \infty} \mathcal{L}_x[Y]$.

Proof As in (21.43) for $0 \leq t_1 \leq t_2$, $\lambda_1, \lambda_2 \geq 0$ and $x \geq 0$, we get

$$\begin{aligned} & \lim_{n \rightarrow \infty} \mathbf{E}_x \left[e^{-(\lambda_1 \tilde{Z}_{t_1}^n + \lambda_2 \tilde{Z}_{t_2}^n)} \right] \\ &= \lim_{n \rightarrow \infty} \mathbf{E}_x \left[\mathbf{E}_x \left[e^{-\lambda_2 \tilde{Z}_{t_2}^n | \tilde{Z}_{t_1}^n} \right] e^{-\lambda_1 \tilde{Z}_{t_1}^n} \right] \\ &= \lim_{n \rightarrow \infty} \mathbf{E}_x \left[\exp \left(-\frac{\lambda_2}{\lambda_2(t_2 - t_1) + 1} \tilde{Z}_{t_1}^n \right) e^{-\lambda_1 \tilde{Z}_{t_1}^n} \right] \\ &= \exp \left(-\frac{\left(\frac{\lambda_2}{\lambda_2(t_2 - t_1) + 1} + \lambda_1 \right) x}{\left(\frac{\lambda_2}{\lambda_2(t_2 - t_1) + 1} + \lambda_1 \right) t_1 + 1} \right) \\ &= \mathbf{E}_x [\exp(-(\lambda_1 Y_{t_1} + \lambda_2 Y_{t_2}))] \end{aligned}$$

Hence, we obtain

$$\mathcal{L}_x [\lambda_1 \tilde{Z}_{t_1}^n + \lambda_2 \tilde{Z}_{t_2}^n] \xrightarrow{n \rightarrow \infty} \mathcal{L}_x [\lambda_1 Y_{t_1} + \lambda_2 Y_{t_2}].$$

Using the Cramér-Wold device (Theorem 15.56), this implies

$$\mathcal{L}_x \left[\left(\tilde{Z}_{t_1}^n, \tilde{Z}_{t_2}^n \right) \right] \xrightarrow{n \rightarrow \infty} \mathcal{L}_x [(Y_{t_1}, Y_{t_2})].$$

Iterating the argument, for every $k \in \mathbb{N}$ and $0 \leq t_1 \leq t_2 \leq \dots \leq t_k$, we get

$$\mathcal{L}_x \left[\left(\tilde{Z}_{t_i}^n \right)_{i=1, \dots, k} \right] \xrightarrow{n \rightarrow \infty} \mathcal{L}_x [(Y_{t_i})_{i=1, \dots, k}]$$

However, this was the claim.

The final step is to show convergence in path space. To this end, we have to modify the rescaled processes so that they become continuous. Assume that $(Z_i^n)_{i \in \mathbb{N}_0}$, $n \in \mathbb{N}$ is a sequence of Galton-Watson processes with $Z_0^n = \lfloor nx \rfloor$. Define the linearly interpolated processes

$$\bar{Z}_t^n := (t - n^{-1} \lfloor tn \rfloor) (Z_{\lfloor tn \rfloor + 1}^n - Z_{\lfloor tn \rfloor}^n) + \frac{1}{n} Z_{\lfloor tn \rfloor}^n.$$

Theorem 21.51 (Lindvall (1972), see [109]) As $n \rightarrow \infty$, in the sense of weak convergence in $\mathcal{M}_1(C([0, \infty)))$, the rescaled Galton-Watson processes $\bar{Z}^n$ converge to Feller's diffusion Y :

$$\mathcal{L}_x [\bar{Z}^n] \xrightarrow{n \rightarrow \infty} \mathcal{L}_x[Y]$$

Proof We have shown already the convergence of the finite-dimensional distributions. By Theorem 21.38, it is thus enough to show tightness of $(\mathcal{L}_x [\bar{Z}^n], n \in \mathbb{N})$ in $\mathcal{M}_1(C([0, \infty)))$. To this end, we apply Kolmogorov's moment criterion (Theorem 21.42 with $\alpha = 4$ and $\beta = 1$

). Hence, for fixed $N > 0$, we compute the fourth moments $\mathbf{E}_x \left[(\bar{Z}_{t+s}^n - \bar{Z}_s^n)^4 \right]$ for $s, t \in [0, N]$. We distinguish two cases:

Case 1: $t < \frac{1}{n}$. Let $k = \lfloor (t+s)n \rfloor$. First assume that $\lfloor sn \rfloor = k$. Then (by Lemma 21.45)

$$\begin{aligned} \mathbf{E}_x \left[(\bar{Z}_{t+s}^n - \bar{Z}_s^n)^4 \right] &= n^{-4} (tn)^4 \mathbf{E}_{\lfloor nx \rfloor} \left[(Z_{k+1}^n - Z_k^n)^4 \right] \\ &= t^4 \mathbf{E}_{\lfloor nx \rfloor} \left[24Z_k^n + 12(Z_k^n)^2 + 2Z_k^n \right] \\ &= t^4 (26\lfloor nx \rfloor + 24\lfloor nx \rfloor k + \lfloor nx \rfloor^2) \\ &\leq 26xt^3 + 24xst^2 + x^2t^2 \\ &\leq (50Nx + x^2) t^2. \end{aligned}$$

In the case $\lfloor sn \rfloor = k - 1$, we get a similar estimate. Therefore, there is a constant $C = C(N, x)$ such that

$$\mathbf{E}_x \left[(\bar{Z}_{s+t}^n - \bar{Z}_s^n)^4 \right] \leq Ct^2 \quad \text{for all } s, t \in [0, N] \text{ with } t < \frac{1}{n}. \quad (21.48)$$

Case 2: $t \geq \frac{1}{n}$. Define $k := \lceil (t+s)n \rceil - \lfloor sn \rfloor \leq tn + 1 \leq 2tn$. Then (by Lemma 21.45)

$$\begin{aligned} \mathbf{E}_x \left[(\bar{Z}_{t+s}^n - \bar{Z}_s^n)^4 \right] &\leq n^{-4} \mathbf{E}_{\lfloor nx \rfloor} \left[(Z_{\lfloor (t+s)n \rfloor}^n - Z_{\lfloor sn \rfloor}^n)^4 \right] \\ &= n^{-4} \mathbf{E}_{\lfloor nx \rfloor} \left[\mathbf{E}_{Z_{\lfloor sn \rfloor}^n} \left[(Z_k^n - Z_0^n)^4 \right] \right] \\ &= n^{-4} \mathbf{E}_{\lfloor nx \rfloor} \left[24Z_{\lfloor sn \rfloor}^n k^3 + 12(Z_{\lfloor sn \rfloor}^n)^2 k^2 + 2Z_{\lfloor sn \rfloor}^n k \right] \\ &\leq n^{-4} (24xn(2tn)^3 + (24xnsn + 12x^2n^2) (2tn)^2 + 4xtn^2) \\ &\leq 192xt^3 + (96xs + 48x^2) t^2 + 4xn^{-1}t^2 \\ &\leq (292Nx + 48x^2) t^2. \end{aligned} \quad (21.49)$$

Combining the estimates (21.48) and (21.49), the assumptions of Kolmogorov's moment criterion for tightness (Theorem 21.42) are fulfilled with $\alpha = 4$ and $\beta = 1$. Hence the sequence $(\mathcal{L}_x[\bar{Z}^n], n \in \mathbb{N})$ is tight.

5.21.10 21.10 Square Variation and Local Martingales

By the Paley-Wiener-Zygmund theorem (Theorem 21.17), the paths $t \mapsto W_t$ of Brownian motion are almost surely nowhere differentiable and hence have locally infinite variation. In particular, the stochastic integral $\int_0^1 f(s) dW_s$ that we introduced in Example 21.29 cannot be understood as a Lebesgue-Stieltjes integral. However, as a preparation for the construction of integrals of this type for larger classes of integrands and integrators (in Chapter 25), here we investigate the path properties of Brownian motion and, somewhat more generally, of continuous local martingales in more detail.

Definition 21.52 Let $G : [0, \infty) \rightarrow \mathbb{R}$ be a continuous function. For any $t \geq 0$, define the variation up to t by

$$V_t^1(G) := \sup \left\{ \sum_{i=0}^{n-1} |G_{t_{i+1}} - G_{t_i}| : 0 = t_0 \leq t_1 \leq \dots \leq t_n = t, n \in \mathbb{N} \right\}.$$

We say that G has locally finite variation if $V_t^1(G) < \infty$ for all $t \geq 0$. We write $\mathcal{C}_v$ for the vector space of continuous functions G with continuous variation $t \mapsto V_t^1(G)$.

Remark 21.53 Clearly, $V^1(F + G) \leq V^1(F) + V^1(G)$ and $V^1(\alpha G) = |\alpha|V^1(G)$ for all continuous $F, G : [0, \infty) \rightarrow \mathbb{R}$ and for all $\alpha \in \mathbb{R}$. Hence $\mathcal{C}_v$ is indeed a vector space. $\square$

Remark 21.54

(i) If G is of the form $G_t = \int_0^t f(s)ds$ for some locally integrable function f , then we have $G \in \mathcal{C}_v$ with $V_t^1(G) = \int_0^t |f(s)|ds$.

(ii) If $G = G^+ - G^-$ is the difference of two continuous monotone increasing functions G^+ and G^- , then

$$V_t^1(G) - V_s^1(G) \leq (G_t^+ - G_s^+) + (G_t^- - G_s^-) \text{ for all } t > s, \quad (21.50)$$

hence we have $G \in \mathcal{C}_v$. In (21.50), equality holds if G^- and G^+ "do not grow on the same sets"; that is, more formally, if G^- and G^+ are the distribution functions of mutually singular measures μ^- and μ^+ . The measures μ^- and μ^+ are then the Jordan decomposition of the signed measure $\mu = \mu^+ - \mu^-$ whose distribution function is G . Then the Lebesgue-Stieltjes integral is defined by

$$\int_0^t F(s)dG_s := \int_{[0,t]} Fd\mu^+ - \int_{[0,t]} Fd\mu^- \quad (21.51)$$

(iii) If $G \in \mathcal{C}_v$, then clearly

$$G_t^+ := \frac{1}{2} (V_t^1(G) + G_t) \quad \text{and} \quad G_t^- := \frac{1}{2} (V_t^1(G) - G_t)$$

is a decomposition of G as in (ii). $\square$

The fact that the paths of Brownian motion are nowhere differentiable can be used to infer that the paths have infinite variation. However, there is also a simple direct argument.

Theorem 21.55 Let W be a Brownian motion. Then $V_t^1(W) = \infty$ almost surely for every $t > 0$.

Proof It is enough to consider $t = 1$ and to show

$$Y_n := \sum_{i=1}^{2^n} |W_{i2^{-n}} - W_{(i-1)2^{-n}}| \xrightarrow{n \rightarrow \infty} \infty \quad \text{a.s.} \quad (21.52)$$

We have $\mathbf{E}[Y_n] = 2^{n/2} \mathbf{E}[|W_1|] = 2^{n/2} \sqrt{2/\pi}$ and $\text{Var}[Y_n] = 1 - 2/\pi$. By Chebyshev's inequality,

$$\sum_{n=1}^{\infty} \mathbf{P} \left[Y_n \leq \frac{1}{2} 2^{n/2} \sqrt{2/\pi} \right] \leq \sum_{n=1}^{\infty} \frac{2\pi - 4}{2^n} = 2\pi - 4 < \infty$$

Using the Borel-Cantelli lemma, this implies (21.52).

Evidently, the variation is too crude a measure to quantify essential path properties of Brownian motion. Hence, instead of the increments (in the definition of the variation), we will sum up the (smaller) squared increments. For the definition of this square variation, more care is needed than in Definition 21.52 for the variation.

Definition 21.56 A sequence $\mathcal{P} = (\mathcal{P}^n)_{n \in \mathbb{N}}$ of countable subsets of $[0, \infty)$,

$$\mathcal{P}^n := \{t_0, t_1, t_2, \dots\} \quad \text{with } 0 = t_0 < t_1 < t_2 < \dots,$$

is called an admissible partition sequence if

- (i) $\mathcal{P}^1 \subset \mathcal{P}^2 \subset \dots$,
- (ii) $\sup \mathcal{P}^n = \infty$ for every $n \in \mathbb{N}$, and
- (iii) the mesh size

$$|\mathcal{P}^n| := \sup_{t \in \mathcal{P}^n} \min_{s \in \mathcal{P}^n, s \neq t} |s - t|$$

tends to 0 as $n \rightarrow \infty$.

If $0 \leq S < T$, then define

$$\mathcal{P}_{S,T}^n := \mathcal{P}^n \cap [S, T) \quad \text{and} \quad \mathcal{P}_T^n := \mathcal{P}^n \cap [0, T).$$

If $t = t_k \in \mathcal{P}_T^n$, then let $t' := t_{k+1} \wedge T = \min \{s \in \mathcal{P}_T^n \cup \{T\} : s > t\}$.

Example 21.57 $\mathcal{P}^n = \{k2^{-n} : k = 0, 1, 2, \dots\}$.

Definition 21.58 For continuous $F, G : [0, \infty) \rightarrow \mathbb{R}$ and for $p \geq 1$, define the p -variation of G (along $\mathcal{P}$) by

$$V_T^p(G) := V_T^{\mathcal{P},p}(G) := \lim_{n \rightarrow \infty} \sum_{t \in \mathcal{P}_T^n} |G_{t'} - G_t|^p \quad \text{for } T \geq 0$$

if the limit exists. In particular, $\langle G \rangle := V^2(G)$ is called the square variation of G . If $T \mapsto V_T^2(G)$ is continuous, then we write $G \in \mathcal{C}_{\text{qv}} := \mathcal{C}_{\text{qv}}^{\mathcal{P}}$.

If, for every $T \geq 0$, the limit

$$V_T^{\mathcal{P},2}(F, G) := \lim_{n \rightarrow \infty} \sum_{t \in \mathcal{P}_T^n} (F_{t'} - F_t)(G_{t'} - G_t)$$

exists, then we call $\langle F, G \rangle := V^2(F, G) := V^{\mathcal{P},2}(F, G)$ the quadratic covariation of F and G (along $\mathcal{P}$).

Remark 21.59 If $p' > p$ and $V_T^p(G) < \infty$, then $V_T^{p'}(G) = 0$. In particular, we have $\langle G \rangle \equiv 0$ if G has locally finite variation.

Remark 21.60 By the triangle inequality, we have

$$\sum_{t \in \mathcal{P}_T^{n+1}} |G_{t'} - G_t| \geq \sum_{t \in \mathcal{P}_T^n} |G_{t'} - G_t| \quad \text{for all } n \in \mathbb{N}, T \geq 0$$

Hence in the case $p = 1$, the limit always exists and coincides with $V^1(G)$ from Definition 21.52 (and is hence independent of the particular choice of $\mathcal{P}$). A similar inequality does not hold for V^2 and thus the limit need not exist or may depend on the choice of $\mathcal{P}$. In the following, we will, however, show that, for a large class of continuous stochastic processes, V^2 exists almost surely along a suitable subsequence of partitions and is almost surely unique.

Remark 21.61

(i) If $\langle F + G \rangle_T$ and $\langle F - G \rangle_T$ exist, then the covariation $\langle F, G \rangle_T$ exists and the polarization formula holds:

$$\langle F, G \rangle_T = \frac{1}{4} (\langle F + G \rangle_T - \langle F - G \rangle_T).$$

(ii) If $\langle F \rangle_T, \langle G \rangle_T$ and $\langle F, G \rangle_T$ exist, then by the Cauchy-Schwarz inequality, we have for the approximating sums

$$V_T^1(\langle F, G \rangle) \leq \sqrt{\langle F \rangle_T \langle G \rangle_T}$$

Remark 21.62 If $f \in C^1(\mathbb{R})$ and $G \in \mathcal{C}_{qv}$, then (exercise!) in the sense of the Lebesgue-Stieltjes integral

$$\langle f(G) \rangle_T = \int_0^T (f'(G_s))^2 d\langle G \rangle_s$$

Corollary 21.63 If F has locally finite square variation and if $\langle G \rangle \equiv 0$ (hence, in particular, if G has locally finite variation), then $\langle F, G \rangle \equiv 0$ and $\langle F + G \rangle = \langle F \rangle$.

Theorem 21.64 For Brownian motion W and for every admissible sequence of partitions, we have

$$\langle W \rangle_T = T \quad \text{for all } T \geq 0 \quad \text{a.s.}$$

Proof We prove this only for the case where

$$\sum_{n=1}^{\infty} |\mathcal{P}^n| < \infty \quad (21.53)$$

For the general case, we only sketch the argument.

Accordingly, assume (21.53). If $\langle W \rangle$ exists, then $T \mapsto \langle W \rangle_T$ is monotone increasing. Hence, it is enough to show that $\langle W \rangle_T$ exists for every $T \in \mathbb{Q}^+ = \mathbb{Q} \cap [0, \infty)$ and that almost surely $\langle W \rangle_T = T$. Since $(\tilde{W}_t)_{t \geq 0} = (T^{-1/2}W_{tT})_{t \geq 0}$ is a Brownian motion, and since $\langle \tilde{W} \rangle_1 = T^{-1}\langle W \rangle_T$, it is enough to consider the case $T = 1$.

Define

$$Y_n := \sum_{t \in \mathcal{P}_1^n} (W_{t'} - W_t)^2 \quad \text{for all } n \in \mathbb{N}$$

Then $\mathbf{E}[Y_n] = \sum_{t \in \mathcal{P}_1^n} (t' - t) = 1$ and

$$\text{Var}[Y_n] = \sum_{t \in \mathcal{P}_1^n} \text{Var}[(W_{t'} - W_t)^2] = 2 \sum_{t \in \mathcal{P}_1^n} (t' - t)^2 \leq 2|\mathcal{P}^n|.$$

By assumption (21.53), we thus have $\sum_{n=1}^{\infty} \text{Var}[Y_n] \leq 2 \sum_{n=1}^{\infty} |\mathcal{P}^n| < \infty$; hence $Y_n \xrightarrow{n \rightarrow \infty} 1$ almost surely.

If we drop the assumption (21.53), then we still have $\text{Var}[Y_n] \xrightarrow{n \rightarrow \infty} 0$; hence $Y_n \xrightarrow{n \rightarrow \infty} 1$ in probability. However, it is not too hard to show that $(Y_n)_{n \in \mathbb{N}}$ is a backwards martingale (see, e.g., [140, Theorem I.28]) and thus converges almost surely to 1.

Corollary 21.65 If W and $\tilde{W}$ are independent Brownian motions, then we have $\langle W, \tilde{W} \rangle_T = 0$.

Proof The continuous processes $((W + \tilde{W})/\sqrt{2})$ and $((W - \tilde{W})/\sqrt{2})$ have independent normally distributed increments. Hence they are Brownian motions. By Remark 21.61(i), we have

$$\begin{aligned} 4\langle W, \tilde{W} \rangle_T &= \langle W + \tilde{W} \rangle_T - \langle W - \tilde{W} \rangle_T \\ &= 2\langle (W + \tilde{W})/\sqrt{2} \rangle_T - 2\langle (W - \tilde{W})/\sqrt{2} \rangle_T \\ &= 2T - 2T = 0 \end{aligned}$$

Clearly, $(W_t \tilde{W}_t)_{t \geq 0}$ is a continuous martingale. Now, by Exercise 21.4.2, the process $(W_t^2 - t)_{t \geq 0}$ is also a continuous martingale. Thus, as shown above, the processes $W^2 - \langle \bar{W} \rangle$ and $W\tilde{W} - \langle W, \tilde{W} \rangle$ are martingales. We will see (Theorem 21.70) that the square variation $\langle M(\omega) \rangle$ of a square integrable continuous martingale M always exists (for almost all ω) and that the process $\langle M \rangle$ is uniquely determined by the property that $M^2 - \langle M \rangle$ is a martingale.

In order to obtain a similar statement for continuous martingales that are not square integrable, we make the following definition.

Definition 21.66 (Local martingale) Let $\mathbb{F}$ be a filtration on $(\Omega, \mathcal{F}, \mathbf{P})$ and let τ be an $\mathbb{F}$ -stopping time. An adapted real-valued stochastic process $M = (M_t)_{t \geq 0}$ is called a local martingale up to time τ if there exists a sequence $(\tau_n)_{n \in \mathbb{N}}$ of stopping times such that $\tau_n \uparrow \tau$ almost surely and such that, for every $n \in \mathbb{N}$, the stopped process $M^{\tau_n} = (M_{\tau_n \wedge t})_{t \geq 0}$ is a uniformly integrable martingale. Such a sequence $(\tau_n)_{n \in \mathbb{N}}$ is called a localising sequence for M . M is called a local martingale if M is a local martingale up to time $\tau \equiv \infty$. Denote by $\mathcal{M}_{\text{loc}}$ the space of continuous local martingales.

Remark 21.67 Let M be a continuous adapted process and let τ be a stopping time. Then the following are equivalent:

- (i) M is a local martingale up to time τ .
- (ii) There is a sequence $(\tau_n)_{n \in \mathbb{N}}$ of stopping times with $\tau_n \uparrow \tau$ almost surely and such that every M^{τ_n} is a martingale.
- (iii) There is a sequence $(\tau_n)_{n \in \mathbb{N}}$ of stopping times with $\tau_n \uparrow \tau$ almost surely and such that every M^{τ_n} is a bounded martingale.

Indeed, (iii) $\implies$ (i) $\implies$ (ii) is trivial. Hence assume that (ii) holds, and define τ'_n by

$$\tau'_n := \inf \{t \geq 0 : |M_t| \geq n\} \quad \text{for all } n \in \mathbb{N}.$$

Since M is continuous, we have $\tau'_n \uparrow \infty$. Hence $(\sigma_n)_{n \in \mathbb{N}} := (\tau_n \wedge \tau'_n)_{n \in \mathbb{N}}$ is a localising sequence for M such that every M^{σ_n} is a bounded martingale.

Remark 21.68 A bounded local martingale M is a martingale. Indeed, assume that $|M_t| \leq C < \infty$ almost surely for all $t \geq 0$ and that $(\tau_n)_{n \in \mathbb{N}}$ is a localising sequence for M . Let $t > s \geq 0$ and $A \in \mathcal{F}_s$. Then $A \cap \{\tau_n \leq s\} \in \mathcal{F}_{\tau_n \wedge s}$ and hence

$$\mathbf{E} [M_{\tau_n \wedge t} \mathbb{1}_{A \cap \{\tau_n \leq s\}}] = \mathbf{E} [M_{\tau_n \wedge s} \mathbb{1}_{A \cap \{\tau_n \leq s\}}].$$

Since $\tau_n \uparrow \infty$, the dominated convergence theorem (Corollary 6.26) yields

$$\mathbf{E} [M_t \mathbb{1}_A] = \mathbf{E} [M_s \mathbb{1}_A].$$

Hence $\mathbf{E} [M_t | \mathcal{F}_s] = M_s$ and thus M is a martingale.

Example 21.69

- (i) Every martingale is a local martingale.
- (ii) In Remark 21.68, we saw that bounded local martingales are martingales. On the other hand, even a uniformly integrable local martingale need not be a martingale: Let $W = (W^1, W^2, W^3)$ be a three-dimensional Brownian motion (that is, W^1, W^2 and W^3 are independent Brownian motions) that starts at $W_0 = x \in \mathbb{R}^3 \setminus \{0\}$. Let

$$u(y) = \|y\|^{-1} \quad \text{for } y \in \mathbb{R}^3 \setminus \{0\}.$$

It is easy to check that u is harmonic; that is, $\Delta u(y) = 0$ for all $y \neq 0$. We will see later (Corollary 25.34) that this implies that $M := (u(W_t))_{t \geq 0}$ is a local martingale. Define a localising sequence for M by

$$\tau_n := \inf \{t > 0 : M_t \geq n\} = \inf \{t > 0 : \|W_t\| \leq 1/n\}, \quad n \in \mathbb{N}.$$

An explicit computation with the three-dimensional normal distribution shows $\mathbf{E} [M_t] \leq t^{-1/2} \xrightarrow{t \rightarrow \infty} 0$; hence M is integrable but is not a martingale. Since $M_t \xrightarrow{t \rightarrow \infty} 0$ in L^1 , M is uniformly integrable.

Theorem 21.70 Let M be a continuous local martingale.

(i) There exists a unique continuous, monotone increasing, adapted process $\langle M \rangle = (\langle M \rangle_t)_{t \geq 0}$ with $\langle M \rangle_0 = 0$ such that

$$(M_t^2 - \langle M \rangle_t)_{t \geq 0} \quad \text{is a continuous local martingale.}$$

(ii) If M is a continuous square integrable martingale, then $M^2 - \langle M \rangle$ is a martingale.

(iii) For every admissible sequence of partitions $\mathcal{P} = (\mathcal{P}^n)_{n \in \mathbb{N}}$, we have

$$U_T^n := \sum_{t \in \mathcal{P}_T^n} (M_{t'} - M_t)^2 \xrightarrow{n \rightarrow \infty} \langle M \rangle_T \quad \text{in probability for all } T \geq 0.$$

The process $\langle M \rangle$ is called the square variation process of M .

Remark 21.71 By possibly passing in (iii) to a subsequence $\mathcal{P}'$ (that might depend on T), we may assume that $U_T^n \xrightarrow{n \rightarrow \infty} \langle M \rangle_T$ almost surely. Using the diagonal sequence argument, we obtain (as in the proof of Helly's theorem) a sequence of partitions such that $U_T^n \xrightarrow{n \rightarrow \infty} \langle M \rangle_T$ almost surely for all $T \in \mathbb{Q}^+$. Since both $T \mapsto U_T^n$ and $T \mapsto \langle M \rangle_T$ are monotone and continuous, we get $U_T^n \xrightarrow{n \rightarrow \infty} \langle M \rangle_T$ for all $T \geq 0$ almost surely. Hence, for this sequence of partitions, the pathwise square variation almost surely equals the square variation process:

$$\langle M(\omega) \rangle = V^2(M(\omega)) = \langle M \rangle(\omega)$$

Proof of Theorem 21.70 Step 1. First let $|M_t| \leq C$ almost surely for all $t \geq 0$ for some $C < \infty$. Then, in particular, M is a martingale (by Remark 21.68). Write $U_T^n = M_T^2 - M_0^2 - N_T^n$, where

$$N_T^n = 2 \sum_{t \in \mathcal{P}_T^n} M_t (M_{t'} - M_t), \quad T \geq 0,$$

is a continuous martingale. Assume that we can show that, for every $T \geq 0$, $(U_T^n)_{n \in \mathbb{N}}$ is a Cauchy sequence in $\mathcal{L}^2(\mathbf{P})$. Then also $(N_T^n)_{n \in \mathbb{N}}$ is a Cauchy sequence, and we can define $\tilde{N}_T$ as the L^2 -limit of $(N_T^n)_{n \in \mathbb{N}}$. By Exercise 21.4.3, $\tilde{N}$ has a continuous modification N , and we have $N_T^n \xrightarrow{n \rightarrow \infty} N_T$ in L^2 for all $T \geq 0$. Thus there exists a continuous process $\langle M \rangle$ with

$$U_T^n \xrightarrow{n \rightarrow \infty} \langle M \rangle_T \quad \text{in } L^2 \quad \text{for all } T \geq 0, \quad (21.54)$$

and $N = M^2 - M_0^2 - \langle M \rangle$ is a continuous martingale.

It remains to show that, for all $T \geq 0$,

$$(U_T^n)_{n \in \mathbb{N}} \quad \text{is a Cauchy sequence in } L^2. \quad (21.55)$$

For $m \in \mathbb{N}$, let

$$Z_m := \max \{ (M_t - M_s)^2 : s \in \mathcal{P}_T^m, t \in \mathcal{P}_{s,s'}^n, n \geq m \}$$

Since M is almost surely uniformly continuous on $[0, T]$, we have $Z_m \xrightarrow{m \rightarrow \infty} 0$ almost surely. As $Z_m \leq 4C^2$, we infer

$$\mathbf{E} [Z_m^2] \xrightarrow{m \rightarrow \infty} 0. \quad (21.56)$$

For $n \in \mathbb{N}$ and numbers $a_0, \dots, a_n$, we have

$$(a_n - a_0)^2 - \sum_{k=0}^{n-1} (a_{k+1} - a_k)^2 = 2 \sum_{k=0}^{n-1} (a_k - a_0) (a_{k+1} - a_k).$$

In the following computation, we apply this to each summand in the outer sum to obtain for $m \in \mathbb{N}$ and $n \geq m$

$$\begin{aligned} U_T^m - U_T^n &= \sum_{s \in \mathcal{P}_T^m} \left((M_{s'} - M_s)^2 - \sum_{t \in \mathcal{P}_{s,s'}^n} (M_{t'} - M_t)^2 \right) \\ &= 2 \sum_{s \in \mathcal{P}_T^m} \sum_{t \in \mathcal{P}_{s,s'}^n} (M_t - M_s) (M_{t'} - M_t) \end{aligned} \quad (21.57)$$

Since M is a martingale, for $s_1, s_2 \in \mathcal{P}_T^m$ and $t_1 \in \mathcal{P}_{s_1,s'_1}^n, t_2 \in \mathcal{P}_{s_2,s'_2}^n$ with $t_1 < t_2$, we have

$$\begin{aligned} &\mathbf{E} [(M_{t_1} - M_{s_1}) (M_{t'_1} - M_{t_1}) (M_{t_2} - M_{s_2}) (M_{t'_2} - M_{t_2})] \\ &= \mathbf{E} [(M_{t_1} - M_{s_1}) (M_{t'_1} - M_{t_1}) (M_{t_2} - M_{s_2}) \mathbf{E} [M_{t'_2} - M_{t_2} \mid \mathcal{F}_{t_2}]] = 0 \end{aligned}$$

If we use (21.57) to compute the expectation of $(U_T^m - U_T^n)^2$, then the mixed terms vanish and we get (using the Cauchy-Schwarz inequality in the third line)

$$\begin{aligned} \mathbf{E} [(U_T^n - U_T^m)^2] &= 4\mathbf{E} \left[\sum_{s \in \mathcal{P}_T^m} \sum_{t \in \mathcal{P}_{s,s'}^n} (M_t - M_s)^2 (M_{t'} - M_t)^2 \right] \\ &\leq 4\mathbf{E} \left[Z_m \sum_{t \in \mathcal{P}_T^n} (M_{t'} - M_t)^2 \right] \\ &\leq 4\mathbf{E} [Z_m^2]^{1/2} \mathbf{E} \left[\left(\sum_{t \in \mathcal{P}_T^n} (M_{t'} - M_t)^2 \right)^2 \right]^{1/2} \end{aligned} \quad (21.58)$$

For the second factor,

$$\begin{aligned} &\mathbf{E} \left[\left(\sum_{t \in \mathcal{P}_T^n} (M_{t'} - M_t)^2 \right)^2 \right] \\ &= \mathbf{E} \left[\sum_{t \in \mathcal{P}_T^n} (M_{t'} - M_t)^4 \right] \\ &\quad + 2\mathbf{E} \left[\sum_{s \in \mathcal{P}_T^n} (M_{s'} - M_s)^2 \sum_{t \in \mathcal{P}_{s',T}^n} (M_{t'} - M_t)^2 \right] \end{aligned} \quad (21.59)$$

The first summand in (21.59) is bounded by

$$4C^2 \mathbf{E} \left[\sum_{t \in \mathcal{P}_T^n} (M_{t'} - M_t)^2 \right] = 4C^2 \mathbf{E} [(M_T - M_0)^2] \leq 16C^4$$

The second summand in (21.59) equals

$$\begin{aligned}
& 2\mathbf{E} \left[\sum_{s \in \mathcal{P}_T^n} (M_{s'} - M_s)^2 \mathbf{E} \left[\sum_{t \in \mathcal{P}_{s',T}^n} (M_{t'} - M_t)^2 \mid \mathcal{F}_{s'} \right] \right] \\
&= 2\mathbf{E} \left[\sum_{s \in \mathcal{P}_T^n} (M_{s'} - M_s)^2 \mathbf{E} [(M_T - M_{s'})^2 \mid \mathcal{F}_{s'}] \right] \\
&\leq 8C^2 \mathbf{E} [(M_T - M_0)^2] \leq 32C^4
\end{aligned}$$

Together with (21.58) and (21.56), we obtain

$$\sup_{n \geq m} \mathbf{E} [(U_T^n - U_T^m)^2] \leq 16\sqrt{3}C^2 \mathbf{E} [Z_m^2]^{1/2} \xrightarrow{m \rightarrow \infty} 0$$

This shows (21.55).

Step 2. Now let $M \in \mathcal{M}_{\text{loc},c}$ and let $(\tau_N)_{N \in \mathbb{N}}$ be a localising sequence such that every M^{τ_N} is a bounded martingale (see Remark 21.67). By Step 1, for $T \geq 0$ and $N \in \mathbb{N}$, we have

$$U_T^{N,n} := \sum_{t \in \mathcal{P}_T^n} (M_t^{\tau_N} - M_t^{\tau_N})^2 \xrightarrow{n \rightarrow \infty} \langle M^{\tau_N} \rangle_T \text{ in } L^2$$

Since $U_T^{N,n} = U_T^{N+1,n}$ if $T \leq \tau_N$, there is a continuous process U with $U_T^{N,n} \xrightarrow{n \rightarrow \infty} U_T$ in probability if $T \leq \tau_N$. Thus $\langle M^{\tau_N} \rangle_T = \langle M \rangle_T := U_T$ if $T \leq \tau_N$. Since $\tau_N \uparrow \infty$ almost surely, for all $T \geq 0$,

$$U_T^n \xrightarrow{n \rightarrow \infty} \langle M \rangle_T \text{ in probability.}$$

As $((M^{\tau_N})^2 - \langle M^{\tau_N} \rangle_T)_{T \geq 0}$ is a continuous martingale and since $\langle M^{\tau_N} \rangle = \langle M \rangle^{\tau_N}$, we have $M^2 - \langle M \rangle \in \mathcal{M}_{\text{loc},c}$.

Step 3. It remains to show (ii). Let M be a continuous square integrable martingale and let $(\tau_n)_{n \in \mathbb{N}}$ be a localising sequence for the local martingale $M^2 - \langle M \rangle$. Let $T > 0$ and let $\tau \leq T$ be a stopping time. Then $M_{\tau_n \wedge \tau}^2 \leq \mathbf{E} [M_T^2 \mid \mathcal{F}_{\tau_n \wedge \tau}]$ since M^2 is a nonnegative submartingale. Hence $(M_{\tau_n \wedge \tau}^2)_{n \in \mathbb{N}}$ is uniformly integrable and thus (using the monotone convergence theorem in the last step)

$$\begin{aligned}
\mathbf{E} [M_\tau^2] &= \lim_{n \rightarrow \infty} \mathbf{E} [M_{\tau_n \wedge \tau}^2] = \lim_{n \rightarrow \infty} \mathbf{E} [\langle M \rangle_{\tau_n \wedge \tau}] + \mathbf{E} [M_0^2] \\
&= \mathbf{E} [\langle M \rangle_\tau] + \mathbf{E} [M_0^2].
\end{aligned}$$

Thus, by the optional sampling theorem, $M^2 - \langle M \rangle$ is a martingale.

Step 4 (Uniqueness). Let A and A' be continuous, monotone increasing, adapted processes with $A_0 = A'_0$ such that $M^2 - A$ and $M^2 - A'$ are local martingales. Then also $N = A - A'$ is a local martingale, and for almost all ω , the path $N(\omega)$ has locally finite variation. Thus $\langle N \rangle \equiv 0$ and hence $N^2 - \langle N \rangle = N^2$ is a continuous

local martingale with $N_0 = 0$. Let $(\tau_n)_{n \in \mathbb{N}}$ be a localising sequence for N^2 . Then $\mathbf{E} [N_{\tau_n \wedge t}^2] = 0$ for any $n \in \mathbb{N}$ and $t \geq 0$; hence $N_{\tau_n \wedge t}^2 = 0$ almost surely and thus $N_t^2 = \lim_{n \rightarrow \infty} N_{\tau_n \wedge t}^2 = 0$ almost surely. We conclude $A = A'$.

Corollary 21.72 Let M be a continuous local martingale with $\langle M \rangle \equiv 0$. Then $M_t = M_0$ for all $t \geq 0$ almost surely. In particular, this holds if the paths of M have locally finite variation.

Corollary 21.73 Let $M, N \in \mathcal{M}_{\text{loc},c}$. Then there exists a unique continuous adapted process $\langle M, N \rangle$ with almost surely locally finite variation and $\langle M, N \rangle_0 = 0$ such that

$$MN - \langle M, N \rangle \text{ is a continuous local martingale.}$$

$\langle M, N \rangle$ is called the quadratic covariation process of M and N . For every admissible sequence of partitions $\mathcal{P}$ and for every $T \geq 0$, we have

$$\langle M, N \rangle_T = \lim_{n \rightarrow \infty} \sum_{t \in \mathcal{P}_T^n} (M_{t'} - M_t) (N_{t'} - N_t) \quad \text{in probability.} \quad (21.60)$$

Proof Existence. Manifestly, $M + N, M - N \in \mathcal{M}_{\text{loc},c}$. Define

$$\langle M, N \rangle := \frac{1}{4} (\langle M + N \rangle - \langle M - N \rangle).$$

As the difference of two monotone increasing functions, $\langle M, N \rangle$ has locally finite variation. Using Theorem 21.70(iii), we get (21.60). Furthermore,

$$MN - \langle M, N \rangle = \frac{1}{4} ((M + N)^2 - \langle M + N \rangle) - \frac{1}{4} ((M - N)^2 - \langle M - N \rangle)$$

is a local martingale.

Uniqueness. Let A and A' with $A_0 = A'_0 = 0$ be continuous, adapted and with locally finite variation such that $MN - A$ and $MN - A'$ are in $\mathcal{M}_{\text{loc},c}$. Then $A - A' \in \mathcal{M}_{\text{loc},c}$ have locally finite variation; hence $A - A' = 0$.

Corollary 21.74 If $M \in \mathcal{M}_{\text{loc},c}$ and A are continuous and adapted with $\langle A \rangle \equiv 0$, then $\langle M + A \rangle = \langle M \rangle$.

If M is a continuous local martingale up to the stopping time τ , then $M^\tau \in \mathcal{M}_{\text{loc},c}$, and we write $\langle M \rangle_t := \langle M^\tau \rangle_t$ for $t < \tau$.

Theorem 21.75 Let τ be a stopping time, M be a continuous local martingale up to τ and $\tau_0 < \tau$ a stopping time with $\mathbf{E}[\langle M \rangle_{\tau_0}] < \infty$. Then $\mathbf{E}[M_{\tau_0}] = \mathbf{E}[M_0]$, and M^{τ_0} is an L^2 -bounded martingale if $\mathbf{E}[M_0^2] < \infty$.

Proof Let $\tau_n \uparrow \tau$ be a localising sequence of stopping times for M such that every M^{τ_n} is even a bounded martingale (see Remark 21.67). Then $M^{\tau_0 \wedge \tau_n}$ is also a bounded martingale, and for every $t \geq 0$, we have

$$\mathbf{E}[M_{\tau_0 \wedge \tau_n \wedge t}^2] = \mathbf{E}[M_0^2] + \mathbf{E}[\langle M \rangle_{\tau_0 \wedge \tau_n \wedge t}] \leq \mathbf{E}[M_0^2] + \mathbf{E}[\langle M \rangle_{\tau_0}] < \infty. \quad (21.61)$$

Hence $((M_{\tau_0 \wedge \tau_n \wedge t}), n \in \mathbb{N}, t \geq 0)$ is bounded in L^2 and is thus uniformly integrable. Therefore, by the optional sampling theorem for uniformly integrable martingales,

$$\mathbf{E}[M_{\tau_0}] = \lim_{n \rightarrow \infty} \mathbf{E}[M_{\tau_0 \wedge \tau_n}] = \mathbf{E}[M_0]$$

and, for $t > s$,

$$\begin{aligned} \mathbf{E}[M_t^{\tau_0} \mid \mathcal{F}_s] &= \mathbf{E}\left[\lim_{n \rightarrow \infty} M_t^{\tau_0 \wedge \tau_n} \mid \mathcal{F}_s\right] \\ &= \lim_{n \rightarrow \infty} \mathbf{E}[M_t^{\tau_0 \wedge \tau_n} \mid \mathcal{F}_s] \\ &= \lim_{n \rightarrow \infty} M_s^{\tau_0 \wedge \tau_n} = M_s^{\tau_0} \end{aligned}$$

Hence M^{τ_0} is a martingale.

Corollary 21.76 If $M \in \mathcal{M}_{\text{loc},c}$ with $\mathbf{E}[M_0^2] < \infty$ and $\mathbf{E}[\langle M \rangle_t] < \infty$ for every $t \geq 0$, then M is a square integrable martingale.

Exercise 21.10.1 Show that the random variables $(Y_n)_{n \in \mathbb{N}}$ from the proof of Theorem 21.64 form a backwards martingale.

Exercise 21.10.2 Let $f : [0, \infty) \rightarrow \mathbb{R}$ be continuous and let $X \in \mathcal{C}_{\text{qv}}^{\mathcal{P}}$ for the admissible sequence of partitions $\mathcal{P}$. Show that

$$\int_0^T f(s) d\langle X \rangle_s = \lim_{n \rightarrow \infty} \sum_{t \in \mathcal{P}_T^n} f(t) (X_{t'} - X_t)^2 \quad \text{for all } T \geq 0$$

Exercise 21.10.3 Show by a counterexample that if M is a continuous local martingale with $M_0 = 0$ and if τ is a stopping time with $\mathbf{E}[\langle M \rangle_\tau] = \infty$, then this does not necessarily imply $\mathbf{E}[M_\tau^2] = \infty$.

5.22 22 Law of the Iterated Logarithm

For sums of independent random variables we already know two limit theorems: the law of large numbers and the central limit theorem. The law of large numbers describes for large $n \in \mathbb{N}$, the typical behavior, or average value behavior, of sums of n random variables. On the other hand, the central limit theorem quantifies the typical fluctuations about this average value.

In Chapter 23, we will study atypically large deviations from the average value in greater detail. The aim of this chapter is to quantify the typical fluctuations of the whole process as $n \rightarrow \infty$. The main message is: While for fixed time the partial sum S_n deviates by approximately $\sqrt{n}$ from its expected value (central limit theorem), the maximal fluctuation up to time n is of order $\sqrt{n \log \log n}$ (Hartman-Wintner theorem, Theorem 22.11).

We start with the simpler task of computing the fluctuations for Brownian motion (Theorem 22.1). After that, we will see how sums of independent centered random variables (with finite variance) can be embedded in a Brownian motion (Skorohod's theorem, Theorem 22.5). This embedding will be used to prove the Hartman-Wintner theorem.

In this chapter, we follow essentially the exposition of [39, Section 8.8].

5.22.1 22.1 Iterated Logarithm for the Brownian Motion

Let $(B_t)_{t \geq 0}$ be a Brownian motion. In Example 21.16, as an application of Blumenthal's $0-1$ law, we saw that $\limsup_{t \downarrow 0} B_t / \sqrt{t} = \infty$ a.s. Since by Theorem 21.14, $(tB_{1/t})_{t \geq 0}$ also is a Brownian motion, we get

$$\limsup_{t \rightarrow \infty} \frac{B_t}{\sqrt{t}} = \infty \quad \text{a.s.}$$

The aim of this section is to replace $\sqrt{t}$ by a function such that the limes superior is finite and nontrivial.

Theorem 22.1 (Law of the iterated logarithm for Brownian motion)

$$\limsup_{t \rightarrow \infty} \frac{B_t}{\sqrt{2t \log \log(t)}} = 1 \quad \text{a.s.} \quad (22.1)$$

Before proving the theorem, we state an elementary lemma.

Lemma 22.2 Let $X \sim \mathcal{N}_{0,1}$ be standard normally distributed. Then, for any $x > 0$,

$$\frac{1}{\sqrt{2\pi}} \frac{1}{x + \frac{1}{x}} e^{-x^2/2} \leq \mathbf{P}[X \geq x] \leq \frac{1}{\sqrt{2\pi}} \frac{1}{x} e^{-x^2/2} \quad (22.2)$$

Proof Let $\varphi(t) = \frac{1}{\sqrt{2\pi}} e^{-t^2/2}$ be the density of the standard normal distribution. Partial integration yields the second inequality in (22.2),

$$\mathbf{P}[X \geq x] = \int_x^\infty \frac{1}{t} (t\varphi(t)) dt = -\frac{1}{t} \varphi(t) \Big|_x^\infty - \int_x^\infty \frac{1}{t^2} \varphi(t) dt \leq \frac{1}{x} \varphi(x)$$

Similarly, we get

$$\mathbf{P}[X \geq x] \geq \frac{1}{x}\varphi(x) - \frac{1}{x^2} \int_x^\infty \varphi(t)dt = \frac{1}{x}\varphi(x) - \frac{1}{x^2}\mathbf{P}[X \geq x]$$

This implies the first inequality in (22.2).

Proof of Theorem 22.1 Step 1. " $\leq$ " Let $\alpha > 1$, and define $t_n = \alpha^n$ for $n \in \mathbb{N}$. Later, we let $\alpha \downarrow 1$. Define $f(t) = 2\alpha^2 \log \log t$. Then by the reflection principle (Theorem 21.19) and using the abbreviation $B_{[a,b]} := \{B_t : t \in [a,b]\}$, we obtain

$$\begin{aligned} & \mathbf{P}[\sup B_{[t_n, t_{n+1}]} > \sqrt{t_n f(t_n)}] \\ & \leq \mathbf{P}\left[t_{n+1}^{-1/2} \sup B_{[0, t_{n+1}]} > \sqrt{f(t_n)/\alpha}\right] \\ & = \mathbf{P}\left[\sup B_{[0, 1]} > \sqrt{f(t_n)/\alpha}\right] \\ & \leq \sqrt{\frac{\alpha}{f(t_n)}} e^{-f(t_n)/2\alpha} \\ & = (\log \alpha)^{-\alpha} \sqrt{\frac{\alpha}{f(t_n)}} n^{-\alpha} \\ & \leq n^{-\alpha} \quad \text{for large enough } n. \end{aligned} \tag{22.3}$$

In the next to last step, we used

$$\frac{f(t_n)}{2\alpha} = \alpha(\log(n \log \alpha)) = \alpha \log n + \alpha \log \log \alpha$$

Since $\alpha > 1$, the right-hand side of (22.3) is summable in n :

$$\sum_{n=1}^{\infty} \mathbf{P}\left[\sup B_{[t_n, t_{n+1}]} > \sqrt{t_n f(t_n)}\right] < \infty$$

The Borel-Cantelli lemma (Theorem 2.7) then yields (note that $t \mapsto \sqrt{t f(t)}$ is monotone increasing)

$$\limsup_{t \rightarrow \infty} \frac{B_t}{\sqrt{t f(t)}} \leq 1 \quad \text{a.s.}$$

Now let $\alpha \downarrow 1$ to obtain

$$\limsup_{t \rightarrow \infty} \frac{B_t}{\sqrt{2t \log \log t}} \leq 1 \quad \text{a.s.} \tag{22.4}$$

Step 2. " $\geq$ " Here we show the other inequality in (22.1). To this end, we let $\alpha \rightarrow \infty$. Let $\beta := \frac{\alpha}{\alpha-1} > 1$ and $g(t) = \frac{2}{\beta^2} \log \log t$. Choose n_0 large enough that $\beta g(t_n) \geq 1$ for all $n \geq n_0$. Then, by Brownian scaling (note that $t_n - t_{n-1} = \frac{1}{\beta} t_n$) and (22.2) (since $(x + \frac{1}{x})^{-1} \geq \frac{1}{2} \frac{1}{x}$ for $x = (\beta g(t_n))^{1/2} \geq 1$),

$$\begin{aligned} \mathbf{P}\left[B_{t_n} - B_{t_{n-1}} > \sqrt{t_n g(t_n)}\right] &= \mathbf{P}\left[B_1 > \sqrt{\beta g(t_n)}\right] \\ &\geq \frac{1}{\sqrt{2\pi}} \frac{1}{2} \frac{1}{\sqrt{\beta g(t_n)}} e^{-\beta g(t_n)/2} \\ &= \frac{1}{\sqrt{2\pi}} \frac{1}{2} (\log \alpha)^{-1/\beta} \frac{1}{\sqrt{\beta g(t_n)}} n^{-1/\beta}. \end{aligned}$$

If $\varepsilon \in (0, 1 - 1/\beta)$, then, for sufficiently large $n \in \mathbb{N}$, the right-hand side of the above equation is $\geq n^{-\varepsilon} n^{-1/\beta} \geq n^{-1}$. Hence

$$\sum_{n=2}^{\infty} \mathbf{P} \left[B_{t_n} - B_{t_{n-1}} > \sqrt{t_n g(t_n)} \right] = \infty$$

The events are independent and hence the Borel-Cantelli lemma yields

$$\mathbf{P} \left[B_{t_n} - B_{t_{n-1}} > \sqrt{t_n g(t_n)} \text{ for infinitely many } n \right] = 1 \quad (22.5)$$

Since $\frac{t_n \log \log t_n}{t_{n-1} \log \log t_{n-1}} \xrightarrow{n \rightarrow \infty} \alpha$, (22.4) and symmetry of Brownian motion imply that, for $\varepsilon > 0$,

$$B_{t_{n-1}} > -(1 + \varepsilon) \alpha^{-1/2} \sqrt{2t_n \log \log t_n} \quad \text{for almost all } n \in \mathbb{N} \quad \text{a.s.} \quad (22.6)$$

From (22.5) and (22.6), it follows that

$$\limsup_{n \rightarrow \infty} \frac{B_{t_n}}{\sqrt{2t_n \log \log t_n}} \geq \frac{1}{\beta} - (1 + \varepsilon) \alpha^{-1/2} = \frac{\alpha - 1}{\alpha} - (1 + \varepsilon) \alpha^{-1/2} \quad \text{a.s.}$$

Now, letting $\alpha \rightarrow \infty$ gives $\limsup_{t \rightarrow \infty} \frac{B_t}{\sqrt{2t \log \log t}} \geq 1$ a.s. Together with (22.4), this implies the claim of the theorem.

Corollary 22.3 For every $s \geq 0$, a.s. we have

$$\limsup_{t \downarrow 0} \frac{B_{s+t} - B_s}{\sqrt{2t \log \log(1/t)}} = 1$$

Proof Without loss of generality, assume $s = 0$. Apply Theorem 22.1 to the Brownian motion $(tB_{1/t})$ (see Theorem 21.14).

Remark 22.4 The statement of Corollary 22.3 is about the typical points s of Brownian motion B . However, there might be points in which Brownian motion moves faster than $\sqrt{2t \log \log(1/t)}$. The precise statement is due to Paul Lévy [106]: Denote by $h(\delta) := \sqrt{2\delta \log(1/\delta)}$ Lévy's modulus of continuity. Then

$$\mathbf{P} \left[\lim_{\delta \downarrow 0} \sup_{\substack{s, t \in [0, 1] \\ 0 \leq t-s \leq \delta}} |B_t - B_s| / h(\delta) = 1 \right] = 1. \quad (22.7)$$

(See, e.g., [144, Theorem I.2.5] for a proof.) This implies in particular that almost surely B is not locally Hölder- $\frac{1}{2}$ -continuous.

5.22.2 22.2 Skorohod's Embedding Theorem

In order to carry over the result of the previous section to sums of square integrable centered random variables, we use an embedding of such random variables in a Brownian motion that is due to Skorohod. This technique also provides an alternative proof of Donsker's invariance principle (Theorem 21.43).

Theorem 22.5 (Skorohod's embedding theorem) Let X be a real random variable with $\mathbf{E}[X] = 0$ and $\text{Var}[X] < \infty$. Then on a suitable probability space we can construct a random variable Ξ , a Brownian motion B that is independent of Ξ and an $\mathbb{F}$ -stopping time τ such that

$$B_\tau \stackrel{\mathcal{D}}{=} X \quad \text{and} \quad \mathbf{E}[\tau] = \text{Var}[X]$$

Here the filtration $\mathbb{F}$ is given by $\mathcal{F}_t = \sigma(\Xi, (B_s)_{s \leq t})$.

Remark 22.6 In the above theorem we can do without the additional random variable; that is, we can choose $\mathbb{F} = \sigma(B)$. The proof is rather involved, though (see p. 515).

Corollary 22.7 Let $X_1, X_2, \dots$ be i.i.d. real random variables with $\mathbf{E}[X_1] = 0$ and $\text{Var}[X_1] < \infty$. Further, let $S_n = X_1 + \dots + X_n, n \in \mathbb{N}$. Then on a suitable probability space there exists a filtration $\mathbb{F}$, a Brownian motion B and $\mathbb{F}$ -stopping times $0 = \tau_0 \leq \tau_1 \leq \tau_2 \leq \dots$ such that $(\tau_n - \tau_{n-1})_{n \in \mathbb{N}}$ is i.i.d., $\mathbf{E}[\tau_1] = \text{Var}[X_1]$ and $(B_{\tau_n})_{n \in \mathbb{N}} \stackrel{D}{=} (S_n)_{n \in \mathbb{N}}$.

Proof of Corollary 22.7 We only sketch the proof. The details are left to the reader. Choose independent triples $(B^{(n)}, \Xi^{(n)}, \tau^{(n)}), n \in \mathbb{N}$, as in Theorem 22.5. Let $\tau_n = \tau^{(1)} + \dots + \tau^{(n)}$. For $t \leq \tau_1$ let $B_t := B_t^{(1)}$, and define recursively

$$B_t = B_{\tau_n} + B_{t-\tau_n}^{(n+1)}, \quad \text{if } \tau_n < t \leq \tau_{n+1}$$

Using repeatedly the strong Markov property of Brownian motion, we see that B is a Brownian motion. Now let $\mathcal{F}_t = \sigma((\Xi_n)_{n \in \mathbb{N}}, (B_s)_{s \leq t})$.

We prepare for the proof of Theorem 22.5 with a lemma. In order to allow measures as integrands, we use the following notation: If $\mu \in \mathcal{M}(E)$ is a measure and $f \in \mathcal{L}^1(\mu)$ is nonnegative, then define $\int \mu(dx) f(x) \delta_x := f\mu$, where $f\mu$ is the measure with density f with respect to μ . This is consistent since for measurable $A \subset E$, we then have

$$\left(\int \mu(dx) f(x) \delta_x \right) (A) = \int \mu(dx) f(x) \delta_x(A) = \int \mu(dx) f(x) \mathbb{1}_A(x) = f\mu(A)$$

Lemma 22.8 Let $\mu \in \mathcal{M}_1(\mathbb{R})$ with $\int x\mu(dx) = 0$ and $\sigma^2 := \int x^2\mu(dx) < \infty$. Then there exists a probability measure $\theta \in \mathcal{M}_1((-\infty, 0) \times [0, \infty))$ with

$$\mu = \int \theta(d(u, v)) \left(\frac{v}{v-u} \delta_u + \frac{-u}{v-u} \delta_v \right) \quad (22.8)$$

Furthermore, $\sigma^2 = - \int uv\theta(d(u, v))$.

Proof Define $m := \int_{[0, \infty)} v\mu(dv) = - \int_{(-\infty, 0)} u\mu(du)$. If $m = 0$, then $\theta = \delta_{(-1, 0)}$ is a possible choice. Assume now $m > 0$ and define θ by

$$\theta(d(u, v)) := m^{-1}(v-u)\mu(du)\mu(dv) \quad \text{for } u < 0 \text{ and } v \geq 0$$

Then

$$\begin{aligned} \int \theta(d(u, v)) &= m^{-1} \int_{(-\infty, 0)} \mu(du) \int_{[0, \infty)} \mu(dv) (v-u) \\ &= m^{-1} \int_{(-\infty, 0)} \mu(du) [m - u\mu([0, \infty))] \\ &= m^{-1} (m\mu((-\infty, 0)) + m\mu([0, \infty))) = 1 \end{aligned}$$

Hence, θ is in fact a probability measure. Furthermore,

$$\begin{aligned} \int \theta(d(u, v)) \left(\frac{v}{v-u} \delta_u + \frac{-u}{v-u} \delta_v \right) &= m^{-1} \int_{(-\infty, 0)} \mu(du) \int_{[0, \infty)} \mu(dv) (v\delta_u - u\delta_v) \\ &= \int_{(-\infty, 0)} \mu(du) \delta_u + \int_{[0, \infty)} \mu(dv) \delta_v = \mu \end{aligned}$$

By (22.8), we infer

$$\sigma^2 = \int \mu(dx) x^2 = \int \theta(d(u, v)) \left(\frac{v}{v-u} u^2 + \frac{-u}{v-u} v^2 \right) = - \int \theta(d(u, v)) uv$$

Proof Theorem 22.5 First assume that X takes only the two values $u < 0$ and $v \geq 0$: $\mathbf{P}[X = u] = \frac{v}{v-u} = 1 - \mathbf{P}[X = v]$. Let

$$\tau_{u,v} = \inf \{t > 0 : B_t \in \{u, v\}\}.$$

By Exercise 21.2.4, we have $\mathbf{E}[B_{\tau_{u,v}}] = 0$; hence $B_{\tau_{u,v}} \stackrel{\mathcal{D}}{=} X$ and $\mathbf{E}[\tau_{u,v}] = -uv$. Now let X be arbitrary with $\mathbf{E}[X] = 0$ and $\sigma^2 := \mathbf{E}[X^2] < \infty$. Define $\mu = \mathbf{P}_X$ and $\theta = \theta_\mu$ as in Lemma 22.8. Further, let $\Xi = (\Xi_u, \Xi_v)$ be a random variable with values in $(-\infty, 0) \times [0, \infty)$ and with distribution θ .

Let $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ where $\mathcal{F}_t := \sigma(\Xi, B_s : s \in [0, t])$. Define $\tau := \tau_{\Xi_u, \Xi_v}$. By continuity of B , we get

$$\{\tau \leq t\} = \bigcap_{\substack{u, v \in \mathbb{Q} \\ u < 0 \leq v}} (\{\Xi \in (-\infty, u] \times [v, \infty)\} \cap \{\tau_{u,v} \leq t\}) \in \mathcal{F}_t$$

Hence τ is an $\mathbb{F}$ -stopping time (but not a $\sigma(B)$ -stopping time). For $x < 0$,

$$\begin{aligned} \mathbf{P}[X \leq x] &= \int_{(-\infty, x] \times [0, \infty)} \theta(d(u, v)) \frac{v}{v-u} \\ &= \int_{(-\infty, x] \times [0, \infty)} \theta(d(u, v)) \mathbf{P}[B_{\tau_{u,v}} = u] \\ &= \mathbf{P}[B_\tau \leq x] \end{aligned}$$

For $x \geq 0$, we similarly get $\mathbf{P}[X > x] = \mathbf{P}[B_\tau > x]$. Summing up, we have $B_\tau \stackrel{\mathcal{D}}{=} X$. Furthermore,

$$\mathbf{E}[\tau] = -\mathbf{E}[\Xi_u \Xi_v] = - \int \theta(d(u, v)) uv = \sigma^2$$

Supplement: Proof of Remark 22.6

Here we prove that in Skorohod's embedding theorem we can really do without randomized stopping times; that is, we can choose a stopping time with respect to the filtration generated by the Brownian motion B . In other words, the stopping time can be chosen without using additional random variables, such as the Ξ in the proof given above.

An elegant proof that is based on stochastic analysis methods can be found in Azéma and Yor; see [6, 7]. See also [118] for a more elementary version of that proof. Here, however, we follow an elementary route whose basic idea goes back to Dubins.

For $u < 0 < v$, let $\tau_{u,v} = \inf \{t > 0 : B_t \in \{u, v\}\}$. Hence, if X is a centered random variable that takes only the values u and v , then, as shown in the proof of Theorem 22.5, $B_{\tau_{u,v}} \stackrel{\mathcal{D}}{=} X$ and $\mathbf{E}[\tau_{u,v}] = \mathbf{E}[X^2]$.

In a first step, we generalize this statement to binary splitting martingales. (Recall from Definition 9.42 that a binary splitting process at each time step has a choice of just two different values, which may however depend on the history of the process.) In a second step, we show that square integrable centered random variables can be expressed as limits of such martingales.

Theorem 22.9 Let $(X_n)_{n \in \mathbb{N}_0}$ be a binary splitting martingale with $X_0 = 0$. Let B be a Brownian motion and let $\mathbb{F} = \sigma(B)$ be its canonical filtration. Then there exist $\mathbb{F}$ -stopping times $0 = \tau_0 \leq \tau_1 \leq \dots$ such that

$$(X_n)_{n \in \mathbb{N}_0} \stackrel{\mathcal{D}}{=} (B_{\tau_n})_{n \in \mathbb{N}_0}$$

and such that $\mathbf{E}[\tau_n] = \mathbf{E}[X_n^2]$ holds for all $n \in \mathbb{N}_0$.

If $(X_n)_{n \in \mathbb{N}_0}$ is bounded in L^2 and thus converges almost surely and in L^2 to some square integrable X_∞ , then $\tau := \sup_{n \in \mathbb{N}} \tau_n < \infty$ a.s., $\mathbf{E}[\tau] = \text{Var}[X_\infty]$ and $X_\infty \stackrel{\mathcal{D}}{=} B_\tau$.

Proof For $n \in \mathbb{N}$, let $f_n : \mathbb{R}^{n-1} \times \{-1, +1\} \rightarrow \mathbb{R}$ and let D_n be a $\{-1, +1\}$ -valued random variable such that $X_n = f_n(X_1, \dots, X_{n-1}, D_n)$ holds (compare Definition 9.42). Without loss of generality, we may assume that f_n is monotone increasing in D_n . Let $\tau_0 := 0$ and inductively define

$$\tau_n := \inf \{t > \tau_{n-1} : B_t \in \{f_n(B_{\tau_1}, \dots, B_{\tau_{n-1}}, -1), f_n(B_{\tau_1}, \dots, B_{\tau_{n-1}}, +1)\}\}.$$

Let $\tilde{X}_n := B_{\tau_n}$ and

$$\tilde{D}_n := \begin{cases} 1, & \text{if } \tilde{X}_n \geq \tilde{X}_{n-1} \\ -1, & \text{else} \end{cases}$$

By Exercise 21.2.4 and using the strong Markov property (at τ_{n-1}), we get

$$\begin{aligned} \mathbf{P}[\tilde{D}_n = 1 \mid \tilde{X}_1, \dots, \tilde{X}_{n-1}] \\ = \frac{\tilde{X}_{n-1} - f_n(\tilde{X}_1, \dots, \tilde{X}_{n-1}, -1)}{f_n(\tilde{X}_1, \dots, \tilde{X}_{n-1}, +1) - f_n(\tilde{X}_1, \dots, \tilde{X}_{n-1}, -1)} \end{aligned}$$

and $\mathbf{E}[\tau_n - \tau_{n-1}] = \mathbf{E}\left[\left(\tilde{X}_n - \tilde{X}_{n-1}\right)^2\right]$. On the other hand, since $(X_n)_{n \in \mathbb{N}_0}$ is a martingale, we have

$$\begin{aligned} X_{n-1} &= \mathbf{E}[X_n \mid X_0, \dots, X_{n-1}] \\ &= \sum_{i=-1, +1} \mathbf{P}[D_n = i \mid X_0, \dots, X_{n-1}] f_n(X_1, \dots, X_{n-1}, i) \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbf{P}[D_n = 1 \mid X_1, \dots, X_{n-1}] \\ = \frac{X_{n-1} - f_n(X_1, \dots, X_{n-1}, -1)}{f_n(X_1, \dots, X_{n-1}, +1) - f_n(X_1, \dots, X_{n-1}, -1)} \end{aligned}$$

This implies $(X_n)_{n \in \mathbb{N}_0} \stackrel{\mathcal{D}}{=} (\tilde{X}_n)_{n \in \mathbb{N}_0}$. Since $\mathbf{E}[\tau_n - \tau_{n-1}] = \mathbf{E}[(X_n - X_{n-1})^2]$, and since the martingale differences $(X_i - X_{i-1}), i \in \mathbb{N}$, are uncorrelated, we get $\mathbf{E}[\tau_n] = \mathbf{E}[X_n^2]$.

Finally, if (X_n) is bounded in L^2 , then by the martingale convergence theorem there is a square integrable centered random variable X_∞ such that $X_n \xrightarrow{n \rightarrow \infty} X_\infty$ almost surely and in L^2 . In particular, we have $\mathbf{E}[X_n^2] \xrightarrow{n \rightarrow \infty} \mathbf{E}[X_\infty^2]$. Clearly, $(\tau_n)_{n \in \mathbb{N}}$ is monotone increasing and thus converges to some stopping time τ . By the monotone convergence theorem, $\mathbf{E}[\tau] = \lim_{n \rightarrow \infty} \mathbf{E}[\tau_n] = \lim_{n \rightarrow \infty} \mathbf{E}[X_n^2] = \mathbf{E}[X_\infty^2] < \infty$. Hence $\tau < \infty$ a.s. As Brownian motion is continuous, we conclude

$$B_\tau = \lim_{n \rightarrow \infty} B_{\tau_n} = \lim_{n \rightarrow \infty} \tilde{X}_n \stackrel{\mathcal{D}}{=} X_\infty$$

We have shown the statement of Remark 22.6 in the case where the random variable X is the limit of a binary splitting martingale. The general case is now implied by the following theorem.

Theorem 22.10 Let X be a square integrable centered random variable. Then there exists a binary splitting martingale $(X_n)_{n \in \mathbb{N}_0}$ with $X_0 = 0$ and such that $X_n \xrightarrow{n \rightarrow \infty} X$ almost surely and in L^2 .

Proof We follow the idea of the proof in [118]. Let $X_0 := \mathbf{E}[X] = 0$. Inductively, for $n \in \mathbb{N}$, define

$$D_n := \begin{cases} 1, & \text{if } X \geq X_{n-1}, \\ -1, & \text{if } X < X_{n-1}, \end{cases}$$

$$\mathcal{F}_n := \sigma(D_1, \dots, D_n)$$

and

$$X_n := \mathbf{E}[X \mid \mathcal{F}_n]$$

Hence there exists a map $g_n : \{-1, +1\}^n \rightarrow \mathbb{R}$ such that $g_n(D_1, \dots, D_n) = X_n$. Clearly, $\mathbb{1}_{D_k=1} = \mathbb{1}_{X_k \geq X_{k-1}}$ almost surely for all $k \in \mathbb{N}$. Hence the $D_1, \dots, D_k$ can be computed from the $X_1, \dots, X_k$. Thus there exists a map $f_n : \mathbb{R}^{n-1} \times \{-1, +1\} \rightarrow \mathbb{R}$ such that $f_n(X_1, \dots, X_{n-1}, D_n) = X_n$. Therefore, (X_n) is binary splitting.

Manifestly, $(X_n)_{n \in \mathbb{N}_0}$ is a martingale. By Jensen's inequality, we have $\mathbf{E}[X_n^2] \leq \mathbf{E}[X^2] < \infty$ for all $n \in \mathbb{N}$. Hence $(X_n)_{n \in \mathbb{N}_0}$ is bounded in L^2 and thus converges almost surely and in L^2 to some square integrable X_∞ . It remains to show that $X_\infty = X$ holds almost surely. To this end, we first show

$$\lim_{n \rightarrow \infty} D_n(\omega) (X(\omega) - X_n(\omega)) = |X(\omega) - X_\infty(\omega)| \quad \text{for almost all } \omega. \quad (22.9)$$

If $X(\omega) = X_\infty(\omega)$, then (22.9) holds trivially. If $X(\omega) > X_\infty(\omega)$, then $X(\omega) > X_n(\omega)$ and thus $D_n(\omega) = 1$ for all sufficiently large n ; hence (22.9) holds. Similarly, we get (22.9) if $X(\omega) < X_\infty(\omega)$.

Evidently, we have

$$\mathbf{E}[D_n(X - X_n)] = \mathbf{E}[D_n \mathbf{E}[X - X_n \mid \mathcal{F}_n]] = 0$$

As $(D_n(X - X_n))_{n \in \mathbb{N}}$ is bounded in L^2 (and is thus uniformly integrable), we get $\mathbf{E}[|X - X_\infty|] = \lim_{n \rightarrow \infty} \mathbf{E}[D_n(X - X_n)] = 0$; hence $X = X_\infty$ a.s.

5.22.3 22.3 Hartman-wintner Theorem

The goal of this section is to prove the law of the iterated logarithm for i.i.d. centered square integrable random variables $X_n, n \in \mathbb{N}$, that goes back to Hartman and Wintner (1941) (see [69]). For the special case of Rademacher random variables, the upper bound was found earlier by Khinchin in 1923 (see [97]).

Theorem 22.11 (Hartman-Wintner, law of the iterated logarithm) Let $X_1, X_2, \dots$ be i.i.d. real random variables with $\mathbf{E}[X_1] = 0$ and $\text{Var}[X_1] = 1$. Let $S_n = X_1 + \dots + X_n, n \in \mathbb{N}$. Then

$$\limsup_{n \rightarrow \infty} \frac{S_n}{\sqrt{2n \log \log n}} = 1 \quad \text{a.s.} \quad (22.10)$$

The strategy of the proof is to embed the partial sums S_n of the random variables in a Brownian motion and then use the law of the iterated logarithm for Brownian motion. The Skorohod embedding theorem ensures that this works. We follow the exposition in [39, Section 8.8].

Proof By Corollary 22.7, on a suitable probability space there exists a filtration $\mathbb{F}$, a Brownian motion B that is an $\mathbb{F}$ -martingale, and stopping times $\tau_1 \leq \tau_2 \leq \dots$ such that $(S_n)_{n \in \mathbb{N}} \stackrel{\mathcal{D}}{=} (B_{\tau_n})_{n \in \mathbb{N}}$. Furthermore, the $(\tau_n - \tau_{n-1})_{n \in \mathbb{N}}$ are i.i.d. with $\mathbf{E}[\tau_n - \tau_{n-1}] = \text{Var}[X_1] = 1$.

By the law of the iterated logarithm for Brownian motion (see Theorem 22.1), we have

$$\limsup_{t \rightarrow \infty} \frac{B_t}{\sqrt{2t \log \log t}} = 1 \quad \text{a.s.}$$

Hence, it is enough to show that

$$\limsup_{t \rightarrow \infty} \frac{|B_t - B_{\tau_{[t]}}|}{\sqrt{2t \log \log t}} = 0 \quad \text{a.s.}$$

By the strong law of large numbers (Theorem 5.17), we have $\frac{1}{n}\tau_n \xrightarrow{n \rightarrow \infty} 1$ a.s., so let $\varepsilon > 0$ and let $t_0 = t_0(\omega)$ be large enough that

$$\frac{1}{1+\varepsilon} \leq \frac{\tau_{[t]}}{t} \leq 1+\varepsilon \quad \text{for all } t \geq t_0$$

Define

$$M_t := \sup_{s \in [t/(1+\varepsilon), t(1+\varepsilon)]} |B_s - B_t|.$$

It is enough to show that

$$\limsup_{t \rightarrow \infty} \frac{M_t}{\sqrt{2t \log \log t}} = 0$$

Consider the sequence $t_n = (1+\varepsilon)^n$, $n \in \mathbb{N}$, and define

$$M'_n := \sup_{s \in [t_{n-1}, t_{n+2}]} |B_s - B_{t_{n-1}}|.$$

Then (by the triangle inequality), for $t \in [t_n, t_{n+1}]$,

$$M_t \leq 2M'_n.$$

Let $\delta := (1+\varepsilon)^3 - 1$. Then $t_{n+2} - t_{n-1} = \delta t_{n-1}$. Brownian scaling and the reflection principle (Theorem 21.19) now yield

$$\begin{aligned} & \mathbf{P} \left[M'_n > \sqrt{3\delta t_{n-1} \log \log t_{n-1}} \right] \\ &= \mathbf{P} \left[\sup_{s \in [0,1]} |B_s| > \sqrt{3 \log \log t_{n-1}} \right] \\ &\leq 2\mathbf{P} \left[\sup_{s \in [0,1]} B_s > \sqrt{3 \log \log t_{n-1}} \right] \\ &= 4\mathbf{P} \left[B_1 > \sqrt{3 \log \log t_{n-1}} \right] \\ &\leq \frac{2}{\sqrt{3 \log \log t_{n-1}}} \exp \left(-\frac{3}{2} \log \log t_{n-1} \right) \quad (\text{Lemma 22.2}) \\ &\leq n^{-3/2} \text{ for } n \text{ sufficiently large.} \end{aligned}$$

The probabilities can be summed over n ; hence the Borel-Cantelli lemma yields

$$\begin{aligned} \limsup_{t \rightarrow \infty} \frac{M_t}{\sqrt{t \log \log t}} &\leq \limsup_{n \rightarrow \infty} \frac{2M'_n}{\sqrt{t_{n-1} \log \log t_{n-1}}} \\ &\leq 2\sqrt{3\delta} \end{aligned}$$

Letting $\varepsilon \rightarrow 0$, we get $\delta = (1 + \varepsilon)^3 - 1 \rightarrow 0$, and hence the proof is complete. $\square$

5.23 23 Large Deviations

Except for the law of the iterated logarithm, so far we have encountered two types of limit theorems for partial sums $S_n = X_1 + \dots + X_n, n \in \mathbb{N}$, of identically distributed, real random variables $(X_i)_{i \in \mathbb{N}}$ with distribution function F :

(1) (Weak) laws of large numbers state that (under suitable assumptions on the family $(X_i)_{i \in \mathbb{N}}$), for every $x > 0$,

$$\mathbf{P} [|S_n - n\mathbf{E}[X_1]| \geq xn] \xrightarrow{n \rightarrow \infty} 0. \quad (23.1)$$

From this we get immediately that the empirical distribution functions

$$F_n : x \mapsto \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{(-\infty, x]}(X_i)$$

converge in probability; that is, $\|F_n - F\|_\infty \xrightarrow{n \rightarrow \infty} 0$. In other words, for any distribution function $G \neq F$ and any $\varepsilon > 0$ with $\varepsilon < \|F - G\|_\infty$, we have

$$\mathbf{P} [\|F_n - G\|_\infty < \varepsilon] \xrightarrow{n \rightarrow \infty} 0. \quad (23.2)$$

(2) Central limit theorems state that (under different assumptions on the family $(X_i)_{i \in \mathbb{N}}$) for every $x \in \mathbb{R}$

$$\mathbf{P} [S_n - n\mathbf{E}[X_1] \geq x\sqrt{n}] \xrightarrow{n \rightarrow \infty} 1 - \Phi\left(\frac{x}{\sqrt{\text{Var}[X_1]}}\right). \quad (23.3)$$

Here $\Phi : t \mapsto \mathcal{N}_{0,1}((-\infty, t])$ is the distribution function of the standard normal distribution.

In each case, the typical value of S_n is $n\mathbf{E}[X_1]$. Equation (23.3) makes a precise statement about the average size of the deviations (which are of order $\sqrt{n}$) from the typical value. A simple consequence is of course that the probability of large deviations (of order n) from the typical value goes to 0 ; that is, (23.1) holds.

In this chapter, we compute the speed of convergence in (23.1) (Cramér's theorem) and in (23.2) (Sanov's theorem).

We follow in part the expositions in [31, 74].

5.23.1 23.1 Cramér's Theorem

Let $X_1, X_2, \dots$ be i.i.d. with $\mathbf{P}_{X_i} = \mathcal{N}_{0,1}$. Then, for every $x > 0$,

$$\mathbf{P} [S_n \geq xn] = \mathbf{P} [X_1 \geq x\sqrt{n}] = 1 - \Phi(x\sqrt{n}) = (1 + \varepsilon_n) \frac{1}{x\sqrt{2\pi n}} e^{-nx^2/2}$$

where $\varepsilon_n \xrightarrow{n \rightarrow \infty} 0$ (by Lemma 22.2). Taking logarithms, we get

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P} [S_n \geq xn] = -\frac{x^2}{2} \quad \text{for every } x > 0 \quad (23.4)$$

It might be tempting to believe that a central limit theorem could be used to show (23.4) for all centered i.i.d. sequences (X_i) with finite variance. However, in general, the limit might be infinite or might be a different function of x , as we will show below. The moral is that large deviations depend more subtly on the tails of the distribution of X_i than the average-sized fluctuations do (which are determined by the variance only). The following theorem shows this for Bernoulli random variables.

Theorem 23.1 Let $X_1, X_2, \dots$ be i.i.d. with $\mathbf{P}[X_1 = -1] = \mathbf{P}[X_1 = 1] = \frac{1}{2}$. Then, for every $x \geq 0$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P}[S_n \geq xn] = -I(x) \quad (23.5)$$

where the rate function I is given by

$$I(z) = \begin{cases} \frac{1+z}{2} \log(1+z) + \frac{1-z}{2} \log(1-z), & \text{if } z \in [-1, 1] \\ \infty, & \text{if } |z| > 1 \end{cases} \quad (23.6)$$

Remark 23.2 Here we agree that $0 \log 0 = 0$. This makes the restriction of I to $[-1, 1]$ a continuous function with $I(-1) = I(1) = \log 2$. Note that I is strictly convex on $[-1, 1]$ with $I(0) = 0$ and I is monotone increasing on $[0, 1]$ and is monotone decreasing on $[-1, 0]$.

Proof of Theorem 23.1 For $x = 0$ and $x > 1$, the claim is trivial. For $x = 1$, we have $\mathbf{P}[S_n \geq n] = 2^{-n}$, and thus again (23.5) holds trivially. Hence, it is enough to consider $x \in (0, 1)$. Since $\frac{S_n + n}{2} \sim b_{n, 1/2}$ is binomially distributed, we have

$$\mathbf{P}[S_n \geq xn] = 2^{-n} \sum_{k \geq (1+x)n/2} \binom{n}{k}.$$




Fig. 23.1 Rate function $I(z)$ from (23.6)

Define $a_n(x) = \lceil n(1+x)/2 \rceil$ for $n \in \mathbb{N}$. Since $k \mapsto \binom{n}{k}$ is monotone decreasing for $k \geq \frac{n}{2}$, we get

$$Q_n(x) := \max \left\{ \binom{n}{k} : a_n(x) \leq k \leq n \right\} = \binom{n}{a_n(x)}. \quad (23.7)$$

We make the estimate

$$2^{-n} Q_n(x) \leq \mathbf{P}[S_n \geq xn] \leq (n+1) 2^{-n} Q_n(x). \quad (23.8)$$

By Stirling's formula

$$\lim_{n \rightarrow \infty} \frac{1}{n!} n^n e^{-n} \sqrt{2\pi n} = 1,$$

we obtain

$$\begin{aligned}
& \lim_{n \rightarrow \infty} \frac{1}{n} \log Q_n(x) \\
&= \lim_{n \rightarrow \infty} \frac{1}{n} \log \frac{n!}{a_n(x)! \cdot (n - a_n(x))!} \\
&= \lim_{n \rightarrow \infty} \frac{1}{n} \log \frac{n^n}{a_n(x)^{a_n(x)} \cdot (n - a_n(x))^{n - a_n(x)}} \\
&= \lim_{n \rightarrow \infty} \left[\log(n) - \frac{a_n(x)}{n} \log(a_n(x)) - \frac{n - a_n(x)}{n} \log(n - a_n(x)) \right] \\
&= \lim_{n \rightarrow \infty} \left[\log(n) - \frac{1+x}{2} \left(\log\left(\frac{1+x}{2}\right) + \log(n) \right) \right. \\
&\quad \left. - \frac{1-x}{2} \left(\log\left(\frac{1-x}{2}\right) + \log(n) \right) \right] \\
&= -\frac{1+x}{2} \log\left(\frac{1+x}{2}\right) - \frac{1-x}{2} \log\left(\frac{1-x}{2}\right) = -I(x) + \log 2
\end{aligned}$$

Together with (23.8), this implies (23.5). $\square$

Under certain assumptions on the distribution of X_1 , Cramér's theorem [29] provides a general principle to compute the rate function I .

Theorem 23.3 (Cramér (1938)) Let $X_1, X_2, \dots$ be i.i.d. real random variables with finite logarithmic moment generating function

$$\Lambda(t) := \log \mathbf{E} [e^{tX_1}] < \infty \quad \text{for all } t \in \mathbb{R} \quad (23.9)$$

Let

$$\Lambda^*(x) := \sup_{t \in \mathbb{R}} (tx - \Lambda(t)) \quad \text{for } x \in \mathbb{R}$$

the Legendre transform of Λ . Then, for every $x > \mathbf{E}[X_1]$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P} [S_n \geq xn] = -I(x) := -\Lambda^*(x) \quad (23.10)$$

Proof By passing to $X_i - x$ if necessary, we may assume $\mathbf{E}[X_i] < 0$ and $x = 0$. (In fact, if $\tilde{X}_i := X_i - x$, and $\tilde{\Lambda}$ and $\tilde{\Lambda}^*$ are defined as Λ and Λ^* above but for $\tilde{X}_i$ instead of X_i , then $\tilde{\Lambda}(t) = \Lambda(t) - t \cdot x$ and thus $\tilde{\Lambda}^*(0) = \sup_{t \in \mathbb{R}} (-\tilde{\Lambda}(t)) = \Lambda^*(x)$.)

Define $\varphi(t) := e^{\Lambda(t)}$ and

$$\varrho := e^{-\Lambda^*(0)} = \inf_{t \in \mathbb{R}} \varphi(t)$$

By (23.9) and the differentiation lemma (Theorem 6.28), φ is differentiable infinitely often and the first two derivatives are

$$\varphi'(t) = \mathbf{E} [X_1 e^{tX_1}] \quad \text{and} \quad \varphi''(t) = \mathbf{E} [X_1^2 e^{tX_1}]$$

Hence φ is strictly convex and $\varphi'(0) = \mathbf{E}[X_1] < 0$.

First consider the case $\mathbf{P} [X_1 \leq 0] = 1$. Then $\varphi'(t) < 0$ for every $t \in \mathbb{R}$ and $\varrho = \lim_{t \rightarrow \infty} \varphi(t) = \mathbf{P} [X_1 = 0]$. Therefore,

$$\mathbf{P} [S_n \geq 0] = \mathbf{P} [X_1 = \dots = X_n = 0] = \varrho^n$$

and thus the claim follows.

Now let $\mathbf{P}[X_1 < 0] > 0$ and $\mathbf{P}[X_1 > 0] > 0$. Then $\lim_{t \rightarrow \infty} \varphi(t) = \infty = \lim_{t \rightarrow -\infty} \varphi(t)$. As φ is strictly convex, there is a unique $\tau \in \mathbb{R}$ at which φ assumes its minimum; hence

$$\varphi(\tau) = \varrho \quad \text{and} \quad \varphi'(\tau) = 0$$

Since $\varphi'(0) < 0$, we have $\tau > 0$. Using Markov's inequality (Theorem 5.11), we estimate

$$\mathbf{P}[S_n \geq 0] = \mathbf{P}[e^{\tau S_n} \geq 1] \leq \mathbf{E}[e^{\tau S_n}] = \varphi(\tau)^n = \varrho^n$$

Thus we get the upper bound

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P}[S_n \geq 0] \leq \log \varrho = -\Lambda^*(0)$$

The remaining part of the proof is dedicated to verifying the reverse inequality:

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P}[S_n \geq 0] \geq \log \varrho \quad (23.11)$$

We use the method of an exponential size-biasing of the distribution $\mu := \mathbf{P}_{X_1}$ of X_1 , which turns the atypical values that are of interest here into typical values. That is, we define the Cramér transform $\hat{\mu} \in \mathcal{M}_1(\mathbb{R})$ of μ by

$$\hat{\mu}(dx) = \varrho^{-1} e^{\tau x} \mu(dx) \quad \text{for } x \in \mathbb{R}$$

Let $\hat{X}_1, \hat{X}_2, \dots$ be independent and identically distributed with $\mathbf{P}_{\hat{X}_i} = \hat{\mu}$. Then

$$\hat{\varphi}(t) := \mathbf{E}[e^{t \hat{X}_1}] = \frac{1}{\varrho} \int_{\mathbb{R}} e^{tx} e^{\tau x} \mu(dx) = \frac{1}{\varrho} \varphi(t + \tau)$$

Hence

$$\begin{aligned} \mathbf{E}[\hat{X}_1] &= \hat{\varphi}'(0) = \frac{1}{\varrho} \varphi'(\tau) = 0 \\ \text{Var}[\hat{X}_1] &= \hat{\varphi}''(0) = \frac{1}{\varrho} \varphi''(\tau) \in (0, \infty) \end{aligned}$$

Defining $\hat{S}_n = \hat{X}_1 + \dots + \hat{X}_n$, we get

$$\begin{aligned} \mathbf{P}[S_n \geq 0] &= \int_{\{x_1 + \dots + x_n \geq 0\}} \mu(dx_1) \dots \mu(dx_n) \\ &= \int_{\{x_1 + \dots + x_n \geq 0\}} (\varrho e^{-\tau x_1}) \hat{\mu}(dx_1) \dots (\varrho e^{-\tau x_n}) \hat{\mu}(dx_n) \\ &= \varrho^n \mathbf{E}[e^{-\tau \hat{S}_n} \mathbb{1}_{\{\hat{S}_n \geq 0\}}] \end{aligned}$$

Thus, in order to show (23.11), it is enough to show

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{E}[e^{-\tau \hat{S}_n} \mathbb{1}_{\{\hat{S}_n \geq 0\}}] \geq 0 \quad (23.12)$$

However, by the central limit theorem (Theorem 15.37), for every $c > 0$,

$$\begin{aligned}
\frac{1}{n} \log \mathbf{E} \left[e^{-\tau \hat{S}_n} \mathbb{1}_{\{\hat{S}_n \geq 0\}} \right] &\geq \frac{1}{n} \log \mathbf{E} \left[e^{-\tau \hat{S}_n} \mathbb{1}_{\{0 \leq \hat{S}_n \leq c\sqrt{n}\}} \right] \\
&\geq \frac{1}{n} \log \left(e^{-\tau c \sqrt{n}} \mathbf{P} \left[\frac{\hat{S}_n}{\sqrt{n}} \in [0, c] \right] \right) \\
&\xrightarrow{n \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{-\tau c \sqrt{n}}{n} + \lim_{n \rightarrow \infty} \frac{1}{n} \log \left(\mathcal{N}_{0, \text{Var}[\hat{X}_1]}([0, c]) \right) \\
&= 0
\end{aligned}$$

Example 23.4 If $\mathbf{P}_{X_1} = \mathcal{N}_{0,1}$, then

$$\Lambda(t) = \log \left(\mathbf{E} \left[e^{tX_1} \right] \right) = \log \left(\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{tx} e^{-x^2/2} dx \right) = \frac{t^2}{2}$$

Furthermore,

$$\Lambda^*(z) = \sup_{t \in \mathbb{R}} (tz - \Lambda(t)) = \sup_{t \in \mathbb{R}} \left(tz - \frac{t^2}{2} \right) = \frac{z^2}{2}$$

Hence the rate function coincides with that of (23.4).

Example 23.5 If $\mathbf{P}_{X_1} = \frac{1}{2}\delta_{-1} + \frac{1}{2}\delta_1$, then $\Lambda(t) = \log \cosh(t)$. The maximizer $t^* = t^*(z)$ of the variational problem for Λ^* solves the equation $z = \Lambda'(t^*) = \tanh(t^*)$. Hence

$$\Lambda^*(z) = z t^* - \Lambda(t^*) = z \operatorname{arctanh}(z) - \log(\cosh(\operatorname{arctanh}(z))).$$

Now $\operatorname{arctanh}(z) = \frac{1}{2} \log \frac{1+z}{1-z}$ for $z \in (-1, 1)$ and

$$\cosh(\operatorname{arctanh}(z)) = \frac{1}{\sqrt{1-z^2}} = \frac{1}{\sqrt{(1-z)(1+z)}}$$

Therefore,

$$\begin{aligned}
\Lambda^*(z) &= \frac{z}{2} \log(1+z) - \frac{z}{2} \log(1-z) + \frac{1}{2} \log(1-z) + \frac{1}{2} \log(1+z) \\
&= \frac{1+z}{2} \log(1+z) + \frac{1-z}{2} \log(1-z)
\end{aligned}$$

However, this is the rate function from Theorem 23.1.

Exercise 23.1.1 Let X be a real random variable with density

$$f(x) = c^{-1} \frac{e^{-|x|}}{1+|x|^3}$$

where $c = \int_{-\infty}^{\infty} \frac{e^{-|x|}}{1+|x|^3} dx$. Check if the logarithmic moment generating function Λ is continuous and sketch the graph of Λ .

5.23.2 23.2 Large Deviations Principle

The basic idea of Cramér's theorem is to quantify the probabilities of rare events by an exponential rate and a rate function. In this section, we develop a formal framework for the quantification of probabilities of rare events in which the complete theory of large deviations can be developed. For further reading, consult, e.g., [31, 32] or [74].

Let E be a Polish space with complete metric d . Recall that

$$B_\varepsilon(x) = \{y \in E : d(x, y) < \varepsilon\}$$

denotes the open ball of radius $\varepsilon > 0$ that is centered at $x \in E$.

A map $f : E \rightarrow \overline{\mathbb{R}} = [-\infty, \infty]$ is called lower semicontinuous if, for every $a \in \overline{\mathbb{R}}$, the level set $f^{-1}([-\infty, a]) \subset E$ is closed. (In particular, continuous maps are lower semicontinuous. On the other hand, $\mathbb{1}_{(0,1)} : \mathbb{R} \rightarrow \mathbb{R}$ is lower semicontinuous but not continuous.) An equivalent condition for lower semicontinuity is that

$$\liminf_{\varepsilon \downarrow 0} f(B_\varepsilon(x)) = f(x) \quad \text{for all } x \in E$$

(Recall that $\inf f(A) = \inf\{f(x) : x \in A\}$.) If $K \subset E$ is compact and nonempty, then f assumes its infimum on K . Indeed, for the case where $f(x) = \infty$ for all $x \in K$, the statement is trivial. Now assume $\inf f(K) < \infty$. If $a_n \downarrow \inf f(K)$ is strictly monotone decreasing, then $K \cap f^{-1}([-\infty, a_n]) \neq \emptyset$ is compact for every $n \in \mathbb{N}$ and hence the infinite intersection also is nonempty:

$$f^{-1}(\inf f(K)) = K \cap \bigcap_{n=1}^{\infty} f^{-1}([-\infty, a_n]) \neq \emptyset.$$

Definition 23.6 (Rate function) A lower semicontinuous function $I : E \rightarrow [0, \infty]$ is called a rate function. If all level sets $I^{-1}([-\infty, a])$, $a \in [0, \infty)$, are compact, then I is called a good rate function.

Definition 23.7 (Large deviations principle) Let I be a rate function and $(\mu_\varepsilon)_{\varepsilon > 0}$ be a family of probability measures on E . We say that $(\mu_\varepsilon)_{\varepsilon > 0}$ satisfies a large deviations principle (LDP) with rate function I if

(LDP 1) $\liminf_{\varepsilon \rightarrow 0} \varepsilon \log(\mu_\varepsilon(U)) \geq -\inf I(U)$ for every open $U \subset E$,

(LDP 2) $\limsup_{\varepsilon \rightarrow 0} \varepsilon \log(\mu_\varepsilon(C)) \leq -\inf I(C)$ for every closed $C \subset E$.

We say that a family $(P_n)_{n \in \mathbb{N}}$ of probability measures on E satisfies an LDP with rate $r_n \uparrow \infty$ and rate function I if (LDP 1) and (LDP 2) hold with $\varepsilon_n = 1/r_n$ and $\mu_{1/r_n} = P_n$.

Often (LDP 1) and (LDP 2) are referred to as lower bound and upper bound. In many cases, the lower bound is a lot easier to show than the upper bound.

Before we show that Cramér's theorem is essentially an LDP, we make two technical statements.

Theorem 23.8 The rate function in an LDP is unique.

Proof Assume that $(\mu_\varepsilon)_{\varepsilon > 0}$ satisfies an LDP with rate functions I and J . Then, for every $x \in E$ and $\delta > 0$,

$$\begin{aligned} I(x) &\geq \inf I(B_\delta(x)) \\ &\geq -\liminf_{\varepsilon \rightarrow 0} \varepsilon \log(\mu_\varepsilon(B_\delta(x))) \\ &\geq -\limsup_{\varepsilon \rightarrow 0} \varepsilon \log\left(\mu_\varepsilon\left(\overline{B_\delta(x)}\right)\right) \\ &\geq \inf J\left(\overline{B_\delta(x)}\right) \xrightarrow{\delta \rightarrow 0} J(x) \end{aligned}$$

Hence $I(x) \geq J(x)$. Similarly, we get $J(x) \geq I(x)$.

Lemma 23.9 Let $N \in \mathbb{N}$ and let $a_\varepsilon^i, i = 1, \dots, N, \varepsilon > 0$, be nonnegative numbers. Then

$$\limsup_{\varepsilon \rightarrow 0} \varepsilon \log \sum_{i=1}^N a_\varepsilon^i = \max_{i=1, \dots, N} \limsup_{\varepsilon \rightarrow 0} \varepsilon \log(a_\varepsilon^i).$$

Proof The sum and maximum differ at most by a factor N :

$$\max_{i=1, \dots, N} \varepsilon \log(a_\varepsilon^i) \leq \varepsilon \log \sum_{i=1}^N a_\varepsilon^i \leq \varepsilon \log(N) + \max_{i=1, \dots, N} \varepsilon \log(a_\varepsilon^i).$$

The maximum and limit (superior) can be interchanged and hence

$$\begin{aligned}
\max_{i=1,\dots,N} \limsup_{\varepsilon \rightarrow 0} \varepsilon \log(a_\varepsilon^i) &= \limsup_{\varepsilon \rightarrow 0} \varepsilon \log \left(\max_{i=1,\dots,N} a_\varepsilon^i \right) \\
&\leq \limsup_{\varepsilon \rightarrow 0} \varepsilon \log \left(\sum_{i=1}^N a_\varepsilon^i \right) \\
&\leq \limsup_{\varepsilon \rightarrow 0} \varepsilon \log(N) + \max_{i=1,\dots,N} \limsup_{\varepsilon \rightarrow 0} \varepsilon \log(a_\varepsilon^i) \\
&= \max_{i=1,\dots,N} \limsup_{\varepsilon \rightarrow 0} \varepsilon \log(a_\varepsilon^i)
\end{aligned}$$

Example 23.10 Let $X_1, X_2, \dots$ be i.i.d. real random variables that satisfy the condition of Cramér's theorem (Theorem 23.3); i.e., $\Lambda(t) = \log(\mathbf{E}[e^{tX_1}]) < \infty$ for every $t \in \mathbb{R}$. Furthermore, let $S_n = X_1 + \dots + X_n$ for every n . We will show that Cramér's theorem implies that $P_n := \mathbf{P}_{S_n/n}$ satisfies an LDP with rate n and with good rate function $I(x) = \Lambda^*(x) := \sup_{t \in \mathbb{R}} (tx - \Lambda(t))$. Without loss of generality, we can assume that $\mathbf{E}[X_1] = 0$. The function I is lower semicontinuous, strictly convex (in the interval where it is finite) and has its unique minimum at $I(0) = 0$. By convexity, we have $I(y) > I(x)$ whenever $y > x \geq 0$ or $y < x \leq 0$.

Cramér's theorem says that $\lim_{n \rightarrow \infty} \frac{1}{n} \log(P_n([x, \infty))) = -I(x)$ for $x > 0$ and (by symmetry) $\lim_{n \rightarrow \infty} \frac{1}{n} \log(P_n((-\infty, x])) = -I(x)$ for $x < 0$. Clearly, for $x > 0$,

$$\begin{aligned}
-I(x) &\geq \liminf_{n \rightarrow \infty} \frac{1}{n} \log P_n((x, \infty)) \\
&\geq \sup_{y > x} \liminf_{n \rightarrow \infty} \frac{1}{n} \log P_n([y, \infty)) = -\inf_{y > x} I(y)
\end{aligned}$$

Similarly, $\liminf_{n \rightarrow \infty} \frac{1}{n} \log P_n((-\infty, x)) \geq -\inf_{y < x} I(y)$ for $x < 0$. Furthermore, by the law of large numbers, for any $x > 0$, we have

$$\begin{aligned}
\lim_{n \rightarrow \infty} \frac{1}{n} \log P_n((-\infty, x)) &= \lim_{n \rightarrow \infty} \frac{1}{n} \log P_n([-x, \infty)) \\
&= \lim_{n \rightarrow \infty} \frac{1}{n} \log P_n((-\infty, x)) = \lim_{n \rightarrow \infty} \frac{1}{n} \log P_n((-\infty, x]) \\
&= 0 = -I(0)
\end{aligned}$$

The main work has been done by showing that the family $(P_n)_{n \in \mathbb{N}}$ satisfies conditions (LDP 1) and (LDP 2) at least for unbounded intervals. It remains to show by some standard arguments (LDP 1) and (LDP 2) for arbitrary open and closed sets, respectively.

First assume that $C \subset \mathbb{R}$ is closed. Define $x_+ := \inf(C \cap [0, \infty))$ as well as $x_- := \sup(C \cap (-\infty, 0])$. By monotonicity of I , on $(-\infty, 0]$ and $[0, \infty)$, we get $\inf I(C) = I(x_-) \wedge I(x_+)$ (with the convention $I(-\infty) = I(\infty) = \infty$). If $x_- = 0$ or $x_+ = 0$, then $\inf(I(C)) = 0$, and (LDP 2) holds trivially. Now let $x_- < 0 < x_+$.

Using Lemma 23.9, we get

$$\begin{aligned}
&\limsup_{n \rightarrow \infty} \frac{1}{n} \log P_n(C) \\
&\leq \limsup_{n \rightarrow \infty} \frac{1}{n} \log (P_n((-\infty, x_-]) + P_n([x_+, \infty))) \\
&= \max \left\{ \limsup_{n \rightarrow \infty} \frac{1}{n} \log P_n((-\infty, x_-]), \limsup_{n \rightarrow \infty} \frac{1}{n} \log P_n([x_+, \infty)) \right\} \\
&= \max \{-I(x_-), -I(x_+)\} = -\inf I(C)
\end{aligned}$$

This shows (LDP 2).

Now let $U \subset \mathbb{R}$ be open. Let $x \in U \cap [0, \infty)$ with $I(x) < \infty$ (if such an x exists). Then there exists an $\varepsilon > 0$ with $(x - \varepsilon, x + \varepsilon) \subset U$. Now

$$\begin{aligned} \liminf_{n \rightarrow \infty} \frac{1}{n} \log P_n((x - \varepsilon, \infty)) &\geq -I(x) > -I(x + \varepsilon) \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \log P_n([x + \varepsilon, \infty)) \end{aligned}$$

Therefore,

$$\begin{aligned} \liminf_{n \rightarrow \infty} \frac{1}{n} \log P_n(U) &\geq \liminf_{n \rightarrow \infty} \frac{1}{n} \log P_n((x - \varepsilon, x + \varepsilon)) \\ &= \liminf_{n \rightarrow \infty} \frac{1}{n} \log (P_n((x - \varepsilon, \infty)) - P_n([x + \varepsilon, \infty))) \\ &= \liminf_{n \rightarrow \infty} \frac{1}{n} \log (P_n((x - \varepsilon, \infty))) \geq -I(x) \end{aligned}$$

Similarly, this also holds for $x \in U \cap (-\infty, 0)$ with $I(x) < \infty$; hence

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log P_n(U) \geq -\inf I(U)$$

This shows the lower bound (LDP 1).

In fact, the condition $\Lambda(t) < \infty$ for all $t \in \mathbb{R}$ can be dropped. Since $\Lambda(0) = 0$, we have $\Lambda^*(x) \geq 0$ for every $x \in \mathbb{R}$. The map Λ^* is a convex rate function but is, in general, not a good rate function. We quote the following strengthening of Cramér's Theorem (see [31, Theorem 2.2.3]).

Theorem 23.11 (Cramér) If $X_1, X_2, \dots$ are i.i.d. real random variables, then $(\mathbf{P}_{S_n/n})_{n \in \mathbb{N}}$ satisfies an LDP with rate function Λ^* .

Exercise 23.2.1 Let $E = \mathbb{R}$. Show that $\mu_\varepsilon := \mathcal{N}_{0,\varepsilon}$ satisfies an LDP with good rate function $I(x) = x^2/2$. Further, show that strict inequality can hold in the upper bound (LDP 2).

Exercise 23.2.2 Let $E = \mathbb{R}$. Show that $\mu_\varepsilon := \mathcal{N}_{0,\varepsilon^2}$ satisfies an LDP with good rate function $I(x) = \infty \cdot \mathbb{1}_{\mathbb{R} \setminus \{0\}}(x)$. Further, show that strict inequality can hold in the lower bound (LDP 1).

Exercise 23.2.3 Let $E = \mathbb{R}$. Show that $\mu_\varepsilon := \frac{1}{2}\mathcal{N}_{-1,\varepsilon} + \frac{1}{2}\mathcal{N}_{1,\varepsilon}$ satisfies an LDP with good rate function $I(x) = \frac{1}{2} \min((x+1)^2, (x-1)^2)$.

Exercise 23.2.4 Compute Λ and Λ^* in the case $X_1 \sim \exp_\theta$ for $\theta > 0$. Interpret the statement of Theorem 23.11 in this case. Check that Λ^* has its unique zero at $\mathbf{E}[X_1]$. (Result: $\Lambda^*(x) = \theta x - \log(\theta x) - 1$ if $x > 0$ and $= \infty$ otherwise.)

Exercise 23.2.5 Compute Λ and Λ^* for the case where X_1 is Cauchy distributed and interpret the statement of Theorem 23.11.

Exercise 23.2.6 Let $X_\lambda \sim \text{Poi}_\lambda$ for every $\lambda > 0$. Show that $\mu_\varepsilon := \mathbf{P}_{\varepsilon X_{1/\varepsilon}}$ satisfies an LDP with good rate function $I(x) = x \log(x/\lambda) + \lambda - x$ for $x \geq 0$ (and $= \infty$ otherwise).

Exercise 23.2.7 Let $(X_t)_{t \geq 0}$ be a random walk on $\mathbb{Z}$ in continuous time that makes a jump to the right with rate $\frac{1}{2}$ and a jump to the left also with rate $\frac{1}{2}$. Show that $(\mathbf{P}_{\varepsilon X_{1/\varepsilon}})_{\varepsilon > 0}$ satisfies an LDP with convex good rate function

$$I(x) = 1 + x \operatorname{arcsinh}(x) - \sqrt{1 + x^2}.$$

5.23.3 23.3 Sanov's Theorem

This section is close to the exposition in [31].

We present a large deviations principle that, unlike Cramér's theorem, is not based on a linear

space. Rather, we consider empirical distributions of independent random variables with values in a finite set Σ , which often is called an alphabet.

Let μ be a probability measure on Σ with $\mu(\{x\}) > 0$ for any $x \in \Sigma$. Further, let $X_1, X_2, \dots$ be i.i.d. random variables with values in Σ and with distribution $\mathbf{P}_{X_1} = \mu$. We will derive a large deviations principle for the empirical measures

$$\xi_n(X) := \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$$

Note that by the law of large numbers, $\mathbf{P}$ -almost surely $\xi_n(X) \xrightarrow{n \rightarrow \infty} \mu$. Hence, as the state space we get $E = \mathcal{M}_1(\Sigma)$, equipped with the metric of total variation $d(\mu, \nu) = \|\mu - \nu\|_{TV}$. (As Σ is finite, in E vague convergence, weak convergence and convergence in total variation coincide.) Further, let

$$E_n := \{\mu \in \mathcal{M}_1(\Sigma) : n\mu(\{x\}) \in \mathbb{N}_0 \text{ for every } x \in \Sigma\}$$

be the range of the random variables $\xi_n(X)$.

Recall that the entropy of μ is defined by

$$H(\mu) := - \int \log(\mu(\{x\})) \mu(dx)$$

If $\nu \in \mathcal{M}_1(\Sigma)$, then we define the relative entropy (or Kullback-Leibler information, see [104]) of ν given μ by

$$H(\nu | \mu) := \int \log \left(\frac{\nu(\{x\})}{\mu(\{x\})} \right) \nu(dx) \quad (23.13)$$

Since $\mu(\{x\}) > 0$ for all $x \in \Sigma$, the integrand ν -a.s. is finite and hence the integral also is finite. A simple application of Jensen's inequality yields $H(\mu) \geq 0$ and $H(\nu | \mu) \geq 0$ (see Lemma 5.26 and Exercise 5.3.3). Furthermore, $H(\nu | \mu) = 0$ if and only if $\nu = \mu$. In addition, clearly,

$$H(\nu | \mu) + H(\nu) = - \int \log(\mu(\{x\})) \nu(dx) \quad (23.14)$$

Since the map $\nu \mapsto I_\mu(\nu) := H(\nu | \mu)$ is continuous, I_μ is a rate function.

Lemma 23.12 For every $n \in \mathbb{N}$ and $\nu \in E_n$, we have

$$(n+1)^{-\#\Sigma} e^{-nH(\nu|\mu)} \leq \mathbf{P}[\xi_n(X) = \nu] \leq e^{-nH(\nu|\mu)} \quad (23.15)$$

Proof We consider the set of possible values for the n -tuple $(X_1, \dots, X_n)$ such that $\xi_n(X) = \nu$:

$$A_n(\nu) := \left\{ k = (k_1, \dots, k_n) \in \Sigma^n : \frac{1}{n} \sum_{i=1}^n \delta_{k_i} = \nu \right\}$$

For every $k \in A_n(\nu)$, we have (compare (23.14))

$$\begin{aligned} \mathbf{P}[\xi_n(X) = \nu] &= \#A_n(\nu) \mathbf{P}[X_1 = k_1, \dots, X_n = k_n] \\ &= \#A_n(\nu) \prod_{x \in \Sigma} \mu(\{x\})^{n\nu(\{x\})} \\ &= \#A_n(\nu) \exp \left(n \int \nu(dx) \log \mu(\{x\}) \right) \\ &= \#A_n(\nu) \exp(-n[H(\nu) + H(\nu | \mu)]) \end{aligned}$$

Now let $Y_1, Y_2, \dots$ be i.i.d. random variables with values in Σ and with distribution $\mathbf{P}_{Y_1} = \nu$. As in the calculation for X , we obtain (since $H(\nu | \nu) = 0$)

$$1 \geq \mathbf{P}[\xi_n(Y) = v] = \#A_n(v)e^{-nH(v)}$$

hence $\#A_n(v) \leq e^{nH(v)}$. This implies the second inequality in (23.15).

The random variable $n\xi_n(Y)$ has the multinomial distribution with parameters $(nv(\{x\}))_{x \in \Sigma}$. Hence the map $E_n \rightarrow [0, 1], v' \mapsto \mathbf{P}[\xi_n(Y) = v']$ is maximal at $v' = v$. Therefore,

$$\#A_n(v) = e^{nH(v)} \mathbf{P}[\xi_n(Y) = v] \geq \frac{e^{nH(v)}}{\#E_n} \geq (n+1)^{-\#\Sigma} e^{nH(v)}$$

This implies the first inequality in (23.15).

We come to the main theorem of this section, Sanov's theorem (see [149, 150]).

Theorem 23.13 (Sanov (1957)) Let $X_1, X_2, \dots$ be i.i.d. random variables with values in the finite set Σ and with distribution μ . Then the family $(\mathbf{P}_{\xi_n(X)})_{n \in \mathbb{N}}$ of distributions of empirical measures satisfies an LDP with rate n and rate function $I_\mu := \bar{H}(\cdot | \mu)$.

Proof By Lemma 23.12, for every $A \subset E$,

$$\begin{aligned} \mathbf{P}[\xi_n(X) \in A] &= \sum_{v \in A \cap E_n} \mathbf{P}[\xi_n(X) = v] \\ &\leq \sum_{v \in A \cap E_n} e^{-nH(v|\mu)} \\ &\leq \#(A \cap E_n) \exp(-n \inf I_\mu(A \cap E_n)) \\ &\leq (n+1)^{\#\Sigma} \exp(-n \inf I_\mu(A)) \end{aligned}$$

Therefore,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P}[\xi_n(X) \in A] \leq -\inf I_\mu(A)$$

Hence the upper bound in the LDP holds (even for arbitrary A).

Similarly, we can use the first inequality in Lemma 23.12 to get

$$\mathbf{P}[\xi_n(X) \in A] \geq (n+1)^{-\#\Sigma} \exp(-n \inf I_\mu(A \cap E_n))$$

and thus

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P}[\xi_n(X) \in A] \geq -\limsup_{n \rightarrow \infty} \inf I_\mu(A \cap E_n) \quad (23.16)$$

Note that, in this inequality, in the infimum we cannot simply replace $A \cap E_n$ by A . However, we show that, for open A this can be done at least asymptotically. Hence, let $A \subset E$ be open. For $v \in A$, there is an $\varepsilon > 0$ with $B_\varepsilon(v) \subset A$. For $n \geq (2\#\Sigma)/\varepsilon$, we have $E_n \cap B_\varepsilon(v) \neq \emptyset$ and hence there exists a sequence $v_n \xrightarrow{n \rightarrow \infty} v$ with $v_n \in E_n \cap A$ for large $n \in \mathbb{N}$. As I_μ is continuous, we have

$$\limsup_{n \rightarrow \infty} \inf I_\mu(A \cap E_n) \leq \lim_{n \rightarrow \infty} I_\mu(v_n) = I_\mu(v)$$

Since $v \in A$ is arbitrary, we get $\limsup_{n \rightarrow \infty} \inf I_\mu(A \cap E_n) = \inf I_\mu(A)$.

Example 23.14 Let $\Sigma = \{-1, 1\}$ and let $\mu = \frac{1}{2}\delta_{-1} + \frac{1}{2}\delta_1$ be the uniform distribution on Σ . Define $m = m(v) := v(\{1\}) - v(\{-1\})$. Then the relative entropy of $v \in \mathcal{M}_1(\Sigma)$ is

$$H(v | \mu) = \frac{1+m}{2} \log(1+m) + \frac{1-m}{2} \log(1-m)$$

Note that this is the rate function from Theorem 23.1.

Next we describe formally the connection between the LDPs of Sanov and Cramér that was indicated in the previous example. To this end, we use Sanov's theorem to derive a version of Cramér's theorem for $\mathbb{R}^d$ -valued random variables taking only finitely many different values.

Example 23.15 Let $\Sigma \subset \mathbb{R}^d$ be finite and let μ be a probability measure on Σ . Further, let $X_1, X_2, \dots$ be i.i.d. random variables with values in Σ and distribution $\mathbf{P}_{X_1} = \mu$. Define $S_n = X_1 + \dots + X_n$ for every $n \in \mathbb{N}$. Let $\Lambda(t) = \log \mathbf{E} [e^{\langle t, X_1 \rangle}]$ for $t \in \mathbb{R}^d$ (which is finite since Σ is finite) and $\Lambda^*(x) = \sup_{t \in \mathbb{R}^d} (\langle t, x \rangle - \Lambda(t))$ for $x \in \mathbb{R}^d$.

We show that $(\mathbf{P}_{S_n/n})_{n \in \mathbb{N}}$ satisfies an LDP with rate n and rate function Λ^* . Let $\xi_n(X)$ be the empirical measure of $X_1, \dots, X_n$. Let $E := \mathcal{M}_1(\Sigma)$. Define the map

$$m : E \rightarrow \mathbb{R}^d, \quad v \mapsto \int x v(dx) = \sum_{x \in \Sigma} x v(\{x\})$$

That is, m maps v to its first moment. Clearly, $\frac{1}{n} S_n = m(\xi_n(X))$. For $x \in \mathbb{R}^d$ and $A \subset \mathbb{R}^d$, define

$$E_x := m^{-1}(\{x\}) = \{v \in E : m(v) = x\}$$

and

$$E_A = m^{-1}(A) = \{v \in E : m(v) \in A\}.$$

The map $v \mapsto m(v)$ is continuous; hence E_A is open (respectively closed) if A is open (respectively closed). Let $\tilde{I}(x) := \inf I_\mu(E_x)$ (where $I_\mu(v) = H(v | \mu)$ is the relative entropy). Then, by Sanov's theorem for open $U \subset \mathbb{R}^d$,

$$\begin{aligned} \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P}_{S_n/n}(U) &= \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P}_{\xi_n(X)}(m^{-1}(U)) \\ &\geq -\inf I_\mu(m^{-1}(U)) = -\inf \tilde{I}(U) \end{aligned}$$

Similarly, for closed $C \subset \mathbb{R}^d$, we have

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P}_{S_n/n}(C) \leq -\inf \tilde{I}(C)$$

In other words, $(\mathbf{P}_{S_n/n})_{n \in \mathbb{N}}$ satisfies an LDP with rate n and rate function $\tilde{I}$. Hence, it only remains to show that $\tilde{I} = \Lambda^*$.

Note that $t \mapsto \Lambda(t)$ is differentiable (with derivative Λ') and is strictly convex. Hence the variational problem for $\Lambda^*(x)$ admits a unique maximizer $t^*(x)$. More precisely,

$$\Lambda^*(x) = \langle t^*(x), x \rangle - \Lambda(t^*(x)),$$

$\Lambda^*(x) > \langle t, x \rangle - \Lambda(t)$ for all $t \neq t^*(x)$, and $\Lambda'(t^*(x)) = x$. By Jensen's inequality, for every $v \in \mathcal{M}_1(\Sigma)$,

$$\begin{aligned} \Lambda(t) &= \log \int e^{\langle t, y \rangle} \mu(dy) \\ &= \log \int \left(e^{\langle t, y \rangle} \frac{\mu(\{y\})}{v(\{y\})} \right) v(dy) \\ &\geq \int \log \left(e^{\langle t, y \rangle} \frac{\mu(\{y\})}{v(\{y\})} \right) v(dy) \\ &= \langle t, m(v) \rangle - H(v | \mu) \end{aligned}$$

with equality if and only if $v = v_t$, where $v_t(\{y\}) = \mu(\{y\})e^{\langle t, y \rangle - \Lambda(t)}$. Hence,

$$\langle t, x \rangle - \Lambda(t) \leq \inf_{v \in E_x} H(v \mid \mu)$$

with equality if $v_t \in E_x$. However, we now know that $m(v_t) = \Lambda'(t)$; hence we have $v_{t^*(x)} \in E_x$ and thus

$$\Lambda^*(x) = \langle t^*(x), x \rangle - \Lambda(t^*(x)) = \inf_{v \in E_x} H(v \mid \mu) = \tilde{I}(x)$$

The method of the proof that we applied in the last example to derive the LDP with rate function $\tilde{I}$ is called a contraction principle. We formulate this principle as a theorem.

Theorem 23.16 (Contraction principle) Assume the family $(\mu_\varepsilon)_{\varepsilon > 0}$ of probability measures on E satisfies an LDP with rate function I . If F is a topological space and $m : E \rightarrow F$ is continuous, then the image measures $(\mu_\varepsilon \circ m^{-1})_{\varepsilon > 0}$ satisfy an LDP with rate function $\tilde{I}(x) = \inf I(m^{-1}(\{x\}))$.

5.23.4 23.4 Varadhan's Lemma and Free Energy

Assume that $(\mu_\varepsilon)_{\varepsilon > 0}$ is a family of probability measures that satisfies an LDP with rate function I . In particular, we know that, for small $\varepsilon > 0$, the mass of μ_ε is concentrated around the zeros of I . In statistical physics, one is often interested in integrating with respect to μ_ε (where $1/\varepsilon$ is interpreted as "size of the system") functions that attain their maximal values away from the zeros of I . In addition, these functions are exponentially scaled with $1/\varepsilon$. Hence the aim is to study the asymptotics of $Z_\varepsilon^\phi := \int e^{\phi(x)/\varepsilon} \mu_\varepsilon(dx)$ as $\varepsilon \rightarrow 0$. Under some mild conditions on the continuity of ϕ , the main contribution to the integral comes from those points x that are not too unlikely (for μ_ε) and for which at the same time $\phi(x)$ is large. That is, those x for which $\phi(x) - I(x)$ is close to its maximum. These contributions are quantified in terms of the tilted probability measures $\mu_\varepsilon^\phi(dx) = (Z_\varepsilon^\phi)^{-1} e^{\phi(x)/\varepsilon} \mu_\varepsilon(dx)$, $\varepsilon > 0$, for which we derive an LDP. As an application, we get the statistical physics principle of minimising the free energy. As an example, we analyze the Weiss ferromagnet.

We start with a lemma that is due to Varadhan [166].

Theorem 23.17 (Varadhan's Lemma (1966)) Let I be a good rate function and let $(\mu_\varepsilon)_{\varepsilon > 0}$ be a family of probability measures on E that satisfies an LDP with rate function I . Further, let $\phi : E \rightarrow \mathbb{R}$ be continuous and assume that

$$\inf_{M > 0} \limsup_{\varepsilon \rightarrow 0} \varepsilon \log \int e^{\phi(x)/\varepsilon} \mathbb{1}_{\{\phi(x) \geq M\}} \mu_\varepsilon(dx) = -\infty \quad (23.17)$$

Then

$$\lim_{\varepsilon \rightarrow 0} \varepsilon \log \int e^{\phi(x)/\varepsilon} \mu_\varepsilon(dx) = \sup_{x \in E} (\phi(x) - I(x)). \quad (23.18)$$

Remark 23.18 (Moment condition) The tail condition (23.17) holds if there exists an $\alpha > 1$ such that

$$\limsup_{\varepsilon \rightarrow 0} \varepsilon \log \int e^{\alpha \phi/\varepsilon} d\mu_\varepsilon < \infty \quad (23.19)$$

Indeed, for every $M \in \mathbb{R}$, we have

$$\begin{aligned}
\varepsilon \log \int e^{\phi(x)/\varepsilon} \mathbb{1}_{\{\phi(x) \geq M\}} \mu_\varepsilon(dx) &= M + \varepsilon \log \int e^{(\phi(x)-M)/\varepsilon} \mathbb{1}_{\{\phi(x) \geq M\}} \mu_\varepsilon(dx) \\
&\leq M + \varepsilon \log \int e^{\alpha(\phi(x)-M)/\varepsilon} \mu_\varepsilon(dx) \\
&= -(\alpha-1)M + \varepsilon \log \int e^{\alpha\phi(x)/\varepsilon} \mu_\varepsilon(dx)
\end{aligned}$$

Together with (23.19), this implies (23.17).

Proof We use different arguments to show that the right-hand side of (23.18) is a lower and an upper bound for the left-hand side.

Lower bound. For any $x \in E$ and $r > 0$, we have

$$\begin{aligned}
\liminf_{\varepsilon \rightarrow 0} \varepsilon \log \int e^{\phi/\varepsilon} d\mu_\varepsilon &\geq \liminf_{\varepsilon \rightarrow 0} \varepsilon \log \int_{B_r(x)} e^{\phi/\varepsilon} d\mu_\varepsilon \\
&\geq \inf \phi(B_r(x)) - I(x) \xrightarrow{r \rightarrow 0} \phi(x) - I(x)
\end{aligned}$$

Upper bound. For $M > 0$ and $\varepsilon > 0$, define

$$F_M^\varepsilon := \int_{\{\phi \geq M\}} e^{\phi(x)/\varepsilon} \mu_\varepsilon(dx) \quad \text{and} \quad G_M^\varepsilon := \int_{\{\phi < M\}} e^{\phi(x)/\varepsilon} \mu_\varepsilon(dx)$$

Define

$$F_M := \limsup_{\varepsilon \rightarrow 0} \varepsilon \log F_M^\varepsilon \quad \text{and} \quad G_M := \limsup_{\varepsilon \rightarrow 0} \varepsilon \log G_M^\varepsilon$$

By Lemma 23.9, for any $M > 0$,

$$\limsup_{\varepsilon \rightarrow 0} \varepsilon \log \int e^{\phi(x)/\varepsilon} \mu_\varepsilon(dx) = F_M \vee G_M$$

As by assumption $\inf_{M>0} F_M = -\infty$, it is enough to show that

$$\sup_{M>0} G_M \leq \sup_{x \in E} (\phi(x) - I(x)) \quad (23.20)$$

Let $\delta > 0$. For any $x \in E$ there is an $r(x) > 0$ with

$$\inf I(B_{2r(x)}(x)) \geq I(x) - \delta \quad \text{and} \quad \sup \phi(B_{2r(x)}(x)) \leq \phi(x) + \delta.$$

Let $a \geq 0$. Since I is a good rate function, the level set $K := I^{-1}([0, a])$ is compact. Thus we can find finitely many $x_1, \dots, x_N \in I^{-1}([0, a])$ such that $\bigcup_{i=1}^N B_{r(x_i)}(x_i) \supset K$. Therefore,

$$\begin{aligned}
G_M^\varepsilon &\leq \int_{\{\phi < M\} \cap K^c} e^{\phi(x)/\varepsilon} \mu_\varepsilon(dx) + \sum_{i=1}^N \int_{\{\phi < M\} \cap B_{r(x_i)}(x_i)} e^{\phi(x)/\varepsilon} \mu_\varepsilon(dx) \\
&\leq e^{M/\varepsilon} \mu_\varepsilon(K^c) + \sum_{i=1}^N e^{(\phi(x_i) \wedge M + \delta)/\varepsilon} \mu_\varepsilon(B_{r(x_i)}(x_i)) \\
&= e^{(M + \varepsilon \log(\mu_\varepsilon(K^c)))/\varepsilon} + \sum_{i=1}^N e^{(\phi(x_i) \wedge M + \delta + \varepsilon \log(\mu_\varepsilon(B_{r(x_i)}(x_i))))/\varepsilon}.
\end{aligned}$$

Using Lemma 23.9 and the LDP, we infer

$$\begin{aligned}
G_M &\leq (M - a) \vee \max_{i=1, \dots, N} (\phi(x_i) - I(x_i) + 2\delta) \\
&\leq (M - a) \vee \sup_{x \in E} (\phi(x) - I(x)) + 2\delta.
\end{aligned}$$

By letting first $\delta \downarrow 0$ and then $a \uparrow \infty$, we obtain (23.20).

Theorem 23.19 (Tilted LDP) Assume that $(\mu_\varepsilon)_{\varepsilon>0}$ satisfies an LDP with good rate function I . Further, let $\phi : E \rightarrow \mathbb{R}$ be a continuous function that satisfies condition (23.17). Define $Z_\varepsilon^\phi := \int e^{\phi/\varepsilon} d\mu_\varepsilon$ and $\mu_\varepsilon^\phi \in \mathcal{M}_1(E)$ by

$$\mu_\varepsilon^\phi(dx) = (Z_\varepsilon^\phi)^{-1} e^{\phi(x)/\varepsilon} \mu_\varepsilon(dx)$$

Further, define $I^\phi : E \rightarrow [0, \infty]$ by

$$I^\phi(x) = \sup_{z \in E} (\phi(z) - I(z)) - (\phi(x) - I(x)). \quad (23.21)$$

Then $(\mu_\varepsilon^\phi)_{\varepsilon>0}$ satisfies an LDP with rate function I^ϕ .

Proof This is left as an exercise. (Compare [32, Exercise 2.1.24], see also [43, Section II.7].)

Varadhan's lemma has various applications in statistical physics. Consider a Polish space Σ that is interpreted as the space of possible states of a particle. Further, let $\lambda \in \mathcal{M}_1(\Sigma)$ be a distribution that is understood as the a priori distribution of this particle if the influence of energy could be neglected. If Σ is finite or is a bounded subset of an $\mathbb{R}^d$, then by symmetry, typically λ is the uniform distribution on Σ . If we place n indistinguishable particles independently according to λ on the random positions $z_1, \dots, z_n \in \Sigma$, then the state of this ensemble can be described by $x := \frac{1}{n} \sum_{i=1}^n \delta_{z_i}$. Denote by $\mu_n^0 \in \mathcal{M}_1(\mathcal{M}_1(\Sigma))$ the corresponding a priori distribution of x ; that is, of the n -particle system.

Now we introduce the hypothesis that the energy $U_n(x)$ of a state has the form $U_n(x) = nU(x)$, where $U(x)$ is the average energy of one particle of the ensemble in state x .

Let $T > 0$ be the temperature of the system and let $\beta := 1/T$ be the so-called inverse temperature. In statistical physics, a key quantity is the so-called partition function

$$Z_n^\beta := \int e^{-\beta U_n} d\mu_n^0$$

A postulate of statistical physics is that the distribution of the state x is the Boltzmann distribution:

$$\mu_n^\beta(dx) = (Z_n^\beta)^{-1} e^{-\beta U_n(x)} \mu_n^0(dx) \quad (23.22)$$

Varadhan's lemma (more precisely, the tilted LDP) and Sanov's theorem are the keys to building a connection to the variational principle for the free energy. For simplicity, assume that Σ is a finite set and $\lambda = \mathcal{U}_\Sigma$ is the uniform distribution on Σ . By Sanov's theorem, $(\mu_n^0)_{n \in \mathbb{N}}$ satisfies an LDP with rate n and rate function $I(x) = H(x | \lambda)$, where $H(x | \lambda)$ is the relative entropy of x with respect to λ . By (23.14), we have $H(x | \lambda) = \log(\#\Sigma) - H(x)$, where $H(x)$ is the entropy of x .




Fig. 23.2 The shifted free energy $F^{\wedge}(\beta)$ of the Weiss ferromagnet without exterior field ($h = 0$)

Define the free energy (or Helmholtz potential) per particle as

$$F^\beta(x) := U(x) - \beta^{-1}H(x).$$

The theorem on the tilted LDP yields that the sequence of Boltzmann distributions $(\mu_n^\beta)_{n \in \mathbb{N}}$ satisfies an LDP with rate n and rate function

$$I^\beta(x) = F^\beta(x) - \inf_{y \in \mathcal{M}_1(\Sigma)} F^\beta(y).$$

Thus, for large n , the Boltzmann distribution is concentrated on those x that minimize the free energy. For different temperatures (that is, for different values of β) these can be very different states. This is the reason for phase transitions at critical temperatures (e.g., melting ice).

Example 23.20 We consider the Weiss ferromagnet. This is a microscopic model for a magnet that assumes that each of n indistinguishable magnetic particles has one of two possible orientations $\sigma_i \in \Sigma = \{-1, +1\}$. The mean magnetization $m = \frac{1}{n} \sum_{i=1}^n \sigma_i$ describes the state of the system completely (as the particles are indistinguishable). Macroscopically, this is the quantity that can be measured. The basic idea is that it is energetically favorable for particles to be oriented in the same direction. We ignore the spatial structure and assume that any particle interacts with any other particle in the same way. This is often called the mean field assumption. In addition, we assume that there is an exterior magnetic field of strength h . Thus up to constants the average energy of a particle is

$$U(m) = -\frac{1}{2}m^2 - hm.$$




Fig. 23.3 Shifted free energy $F^\beta(m) - F^\beta(0)$ of the Weiss ferromagnet with exterior field $h = 0.04$

The entropy of the state m is

$$H(m) = -\frac{1+m}{2} \log \left(\frac{1+m}{2} \right) - \frac{1-m}{2} \log \left(\frac{1-m}{2} \right)$$

Hence the average free energy of a particle is

$$F^\beta(m) = -\frac{1}{2}m^2 - hm + \beta^{-1} \left[\frac{1+m}{2} \log \left(\frac{1+m}{2} \right) + \frac{1-m}{2} \log \left(\frac{1-m}{2} \right) \right].$$

In order to obtain the minima of F^β , we compute the derivative

$$0 \stackrel{!}{=} \frac{d}{dm} F^\beta(m) = -m - h + \beta^{-1} \operatorname{arctanh}(m).$$

Hence, m solves the equation

$$m = \tanh(\beta(m + h)). \quad (23.23)$$

In the case $h = 0$, $m = 0$ is a solution of (23.23) for any β . If $\beta \leq 1$, then this is the only solution and F^β attains its global minimum at $m = 0$. If $\beta > 1$, then (23.23) has two other solutions, $m_-^{\beta,0} \in (-1, 0)$ and $m_+^{\beta,0} = -m_-^{\beta,0}$, whose values can only be computed numerically.

In this case, F^β has a local maximum at 0 and has global minima $m_\pm^{\beta,0}$. For large n , only those values of m for which F^β is close to its minimal value can be attained and thus the distribution is concentrated around 0 if $\beta \leq 1$ and around $m_\pm^{\beta,0}$ if $\beta > 1$. In the latter case, the absolute value of the mean magnetization is $|m_\pm^{\beta,0}| = m_+^{\beta,0} > 0$. Hence, there is a phase transition between the high temperature phase ($\beta \leq 1$) without magnetization and the low temperature phase ($\beta > 1$) where so-called spontaneous magnetization occurs (that is, magnetization without an exterior field).




Fig. 23.4 Weiss ferromagnet: magnetization $\hat{m}(\beta, h)$ as a function of β

If $h \neq 0$, then F^β does not have a minimum at $m = 0$. Rather, F^β is asymmetric and has a global minimum $m^{\beta,h}$ with the same sign as h . Furthermore, for large β , there is another minimum with the opposite sign. Again, the exact values can only be computed numerically. However, for high temperatures (small β), we can approximate $m^{\beta,h}$ using the approximation $\tanh(\beta(m + h)) \approx \beta(m + h)$. Hence we get

$$m^{\beta,h} \approx \frac{h}{\beta^{-1} - 1} = \frac{h}{T - T_c} \quad \text{for } T \rightarrow \infty, \quad (23.24)$$

where the Curie temperature $T_c = 1$ is the critical temperature for spontaneous magnetization. The relation (23.24) is called the Curie-Weiss law. $\square$

5.24 24 The Poisson Point Process

Poisson point processes can be used as a cornerstone in the construction of very different stochastic objects such as, for example, infinitely divisible distributions, Markov processes with complex dynamics, objects of stochastic geometry and so forth.

In this chapter, we briefly develop the general framework of random measures and construct the Poisson point process and characterize it in terms of its Laplace transform. As an application we construct a certain subordinator and show that the Poisson point process is the invariant measure of systems of independent random walks. Via the connection with subordinators,

in the third section, we construct two distributions that play prominent roles in population genetics: the Poisson-Dirichlet distribution and the GEM distribution.

For a nice exposition including many examples, see also [99].

5.24.1 24.1 Random Measures

In the following, let E be a locally compact Polish space (for example, $E = \mathbb{R}^d$ or $E = \mathbb{Z}^d$) with Borel σ -algebra $\mathcal{B}(E)$. Let

$$\mathcal{B}_b(E) = \{B \in \mathcal{B}(E) : B \text{ is relatively compact}\}$$

be the system of bounded Borel sets and $\mathcal{M}(E)$ the space of Radon measures on E (see Definition 13.3).

Definition 24.1 Denote by $\mathbb{M} = \sigma(I_A : A \in \mathcal{B}_b(E))$ the smallest σ -algebra on $\mathcal{M}(E)$ with respect to which all maps

$$I_A : \mu \mapsto \mu(A), \quad A \in \mathcal{B}_b(E)$$

are measurable.

Denote by $\mathcal{B}^+(E)$ the set of measurable maps $E \rightarrow [0, \infty]$ and by $\mathcal{B}_b^{\mathbb{R}}(E)$ the set of bounded measurable maps $E \rightarrow \mathbb{R}$ with compact support. For every $f \in \mathcal{B}^+(E)$, the integral $I_f(\mu) := \int f d\mu$ is well-defined and for every $f \in \mathcal{B}_b^{\mathbb{R}}(E)$, $I_f(\mu)$ is well-defined and finite.

Theorem 24.2 Let τ_v be the vague topology on $\mathcal{M}(E)$. Then

$$\mathbb{M} = \mathcal{B}(\tau_v) = \sigma(I_f : f \in C_c(E)) = \sigma(I_f : f \in C_c^+(E))$$

Proof This is left as an exercise. (See [82, Lemma 4.1].)

Let $\widetilde{\mathcal{M}}(E)$ be the space of all measures on E endowed with the σ -algebra

$$\widetilde{\mathbb{M}} = \sigma(I_A : A \in \mathcal{B}_b(E)).$$

Choose a countable dense set $F \subset E$, and for any $x \in F$ choose a compact neighborhood K_x . Then we get (compare Exercise 13.1.8)

$$\mathcal{M}(E) = \bigcap_{x \in F} \left\{ \mu \in \widetilde{\mathcal{M}}(E) : \mu(K_x) < \infty \right\}$$

Hence $\mathcal{M}(E) \in \widetilde{\mathbb{M}}$. Clearly, $\mathbb{M} = \widetilde{\mathbb{M}}|_{\mathcal{M}(E)}$ is the trace σ -algebra of $\widetilde{\mathbb{M}}$ on $\mathcal{M}(E)$. Here we need the slightly larger space in order to define random measures in such a way that all almost surely well-defined operations on random measures again yield random measures.

Definition 24.3 A random measure on E is a random variable X on some probability space $(\Omega, \mathcal{A}, \mathbf{P})$ with values in $(\widetilde{\mathcal{M}}(E), \widetilde{\mathbb{M}})$ and with $\mathbf{P}[X \in \mathcal{M}(E)] = 1$.

Theorem 24.4 Let X be a random measure on E . Then the set function $\mathbf{E}[X] : \mathcal{B}(E) \rightarrow [0, \infty]$, $A \mapsto \mathbf{E}[X(A)]$ is a measure. We call $\mathbf{E}[X]$ the intensity measure of X . We say that X is integrable if $\mathbf{E}[X] \in \mathcal{M}(E)$.

Proof Clearly, $\mathbf{E}[X]$ is finitely additive. Let $A, A_1, A_2, \dots \in \mathcal{B}(E)$ with $A_n \uparrow A$. Consider the random variables $Y_n := X(A_n)$ and $Y = X(A)$. Then $Y_n \uparrow Y$ and hence, by monotone convergence, $\mathbf{E}[X](A_n) = \mathbf{E}[Y_n] \xrightarrow{n \rightarrow \infty} \mathbf{E}[Y] = \mathbf{E}[X](A)$. Hence $\mathbf{E}[X]$ is lower semicontinuous and is thus a measure (by Theorem 1.36).

Theorem 24.5 Let $\mathbf{P}_X$ be the distribution of a random measure X . Then $\mathbf{P}_X$ is uniquely determined by the distributions of either of the families

$$((I_{f_1}, \dots, I_{f_n}) : n \in \mathbb{N}; f_1, \dots, f_n \in C_c^+(E)) \quad (24.1)$$

or

$$((I_{A_1}, \dots, I_{A_n}) : n \in \mathbb{N}; A_1, \dots, A_n \in \mathcal{B}_b(E) \text{ pairwise disjoint}) . \quad (24.2)$$

Proof The class of sets

$$\mathcal{I} = \{(I_{f_1}, \dots, I_{f_n})^{-1}(A) : n \in \mathbb{N}; f_1, \dots, f_n \in C_c^+(E), A \in \mathcal{B}([0, \infty)^n)\}$$

is a π -system and by Theorem 24.2 it generates $\mathbb{M}$. Hence the measure $\mathbf{P}_X$ is characterized by its values on $\mathcal{I}$.

Similarly, the claim follows for

$$((I_{A_1}, \dots, I_{A_n}) : n \in \mathbb{N}; A_1, \dots, A_n \in \mathcal{B}_b(E)) .$$

If $A_1, \dots, A_n \in \mathcal{B}_b(E)$ are arbitrary, then there exist $2^n - 1$ pairwise disjoint sets $B_1, \dots, B_{2^n-1}$ with $A_i = \bigcup_{k: B_k \subset A_i} B_k$ for all $i = 1, \dots, n$. The distribution of $(I_{A_1}, \dots, I_{A_n})$ can be computed from that of $(I_{B_1}, \dots, I_{B_{2^n-1}})$.

In the following, let $i = \sqrt{-1}$ be the imaginary unit.

Definition 24.6 Let X be a random measure on E . Denote by

$$\mathcal{L}_X(f) = \mathbf{E} \left[\exp \left(- \int f dX \right) \right], \quad f \in \mathcal{B}^+(E)$$

the Laplace transform of X and by

$$\varphi_X(f) = \mathbf{E} \left[\exp \left(i \int f dX \right) \right], \quad f \in \mathcal{B}_b^{\mathbb{R}}(E)$$

the characteristic function of X .

Theorem 24.7 The distribution $\mathbf{P}_X$ of a random measure X is characterized by its Laplace transform $\mathcal{L}_X(f)$, $f \in C_c^+(E)$, as well as by its characteristic function $\varphi_X(f)$, $f \in C_c(E)$.

Proof This is a consequence of Theorem 24.5 and the uniqueness theorem for characteristic functions (Theorem 15.8) and for Laplace transforms (Exercise 15.1.2) of random variables on $[0, \infty)^n$.

Definition 24.8 We say that a random measure X on E has independent increments if, for any choice of finitely many pairwise disjoint measurable sets $A_1, \dots, A_n$, the random variables $X(A_1), \dots, X(A_n)$ are independent.

Corollary 24.9 The distribution of a random measure X on E with independent increments is uniquely determined by the family $(\mathbf{P}_{X(A)}, A \in \mathcal{B}_b(E))$.

Proof This is an immediate consequence of Theorem 24.5.




Fig. 24.1 Poisson point process on the unit square with intensity measure 50λ

Definition 24.10 Let $\mu \in \mathcal{M}(E)$. A random measure X with independent increments is called a Poisson point process (PPP) with intensity measure μ if, for any $A \in \mathcal{B}_b(E)$, we have $\mathbf{P}_{X(A)} = \text{Poi}_{\mu(A)}$. In this case, we write $\text{PPP}_{\mu} := \mathbf{P}_X \in \mathcal{M}_1(\mathcal{M}(E))$ and say that X is a PPP_{μ} .

See Fig. 24.1 for a simulation of a Poisson point process on the unit square.

Remark 24.11 The definition of the PPP (and its construction in the following theorem) still works if $(E, \mathcal{E}, \mu)$ is only assumed to be a σ -finite measure space. However, the characterization in terms of Laplace transforms is a bit simpler in the case of locally compact Polish spaces considered here.

Theorem 24.12 For every $\mu \in \mathcal{M}(E)$, there exists a Poisson point process X with intensity measure μ .

Proof μ is σ -finite since $\mu \in \mathcal{M}(E)$. Hence there exist $E_n \uparrow E$ with $\mu(E_n) < \infty$ for every $n \in \mathbb{N}$. Define $\mu_1 = \mu(E_1 \cap \cdot)$ and $\mu_n = \mu((E_n \setminus E_{n-1}) \cap \cdot)$ for $n \geq 2$. If $X_1, X_2, \dots$ are independent Poisson point processes with intensity measures $\mu_1, \mu_2, \dots$, then $X = \sum_{n=1}^{\infty} X_n$ has intensity measure $\mathbf{E}[X] = \mu$ and hence X is a random measure (see Exercise 24.1.1). Furthermore, it is easy to see that X has independent increments and that

$$\mathbf{P}_{X(A)} = \mathbf{P}_{X_1(A)} * \mathbf{P}_{X_2(A)} * \dots = \text{Poi}_{\mu_1(A)} * \text{Poi}_{\mu_2(A)} * \dots = \text{Poi}_{\mu(A)}.$$

Hence we have $X \sim \text{PPP}_{\mu}$.

Therefore, it is enough to consider the case $\mu(E) \in (0, \infty)$. Define

$$v = \frac{\mu(\cdot)}{\mu(E)} \in \mathcal{M}_1(E).$$

Let $N, Y_1, Y_2, \dots$ be independent random variables with $N \sim \text{Poi}_{\mu(E)}$ and $\mathbf{P}_{Y_i} = v$ for all $i \in \mathbb{N}$. Define

$$X(A) = \sum_{n=1}^N \mathbb{1}_A(Y_n) \quad \text{for } A \in \mathcal{B}(E)$$

The random variables $\mathbb{1}_A(Y_1), \mathbb{1}_A(Y_2), \dots$ are independent and $\text{Ber}_{v(A)}$ -distributed; hence we have $X(A) \sim \text{Poi}_{\mu(A)}$ (see Theorem 15.14(iii)). Let $n \in \mathbb{N}$ and let $A_1, \dots, A_n \in \mathcal{B}(E)$ be pairwise disjoint. Then

$$\psi(t) := \mathbf{E} \left[\exp \left(i \sum_{l=1}^n t_l \mathbb{1}_{A_l}(Y_1) \right) \right] = 1 + \sum_{l=1}^n v(A_l) (e^{it_l} - 1), \quad t \in \mathbb{R}^n,$$

is the characteristic function of $(\mathbb{1}_{A_1}(Y_1), \dots, \mathbb{1}_{A_n}(Y_1))$. Further, let φ be the characteristic function of $(X(A_1), \dots, X(A_n))$ and let φ_l be the characteristic function of $X(A_l)$ for $l = 1, \dots, n$. Hence

$$\varphi_l(t_l) = \exp(\mu(A_l)(e^{it_l} - 1)).$$

By Theorem 15.14(iii), we have

$$\begin{aligned} \varphi(t) &= \mathbf{E} \left[\exp \left(i \sum_{l=1}^n t_l X(A_l) \right) \right] \\ &= \exp(\mu(E)(\psi(t) - 1)) \\ &= \exp \left(\sum_{l=1}^n \mu(A_l)(e^{it_l} - 1) \right) = \prod_{l=1}^n \varphi_l(t_l). \end{aligned}$$

Thus $X(A_1), \dots, X(A_n)$ are independent. This implies $X \sim \text{PPP}_\mu$.

Exercise 24.1.1 Let $X_1, X_2, \dots$ be random measures and $\lambda_1, \lambda_2, \dots \in [0, \infty)$. Define $X := \sum_{n=1}^\infty \lambda_n X_n$. Show that X is a random measure if and only if we have $\mathbf{P}[X(B) < \infty] = 1$ for all $B \in \mathcal{B}_b(E)$. Infer that if X is a random variable with values in $(\widetilde{\mathcal{M}}(E), \widetilde{\mathbb{M}}(E))$ and $\mathbf{E}[X] \in \mathcal{M}(E)$, then X is a random measure.

Exercise 24.1.2 Let τ_w be the topology of weak convergence on $\mathcal{M}_1(E)$ and let $\sigma(\tau_w)$ be the Borel σ -algebra on $\mathcal{M}_1(E)$. Show that $\mathbb{M}|_{\mathcal{M}_1(E)} = \sigma(\tau_w)$.

5.24.2 24.2 Properties of the Poisson Point Process

Theorem 24.13 Let $\mu \in \mathcal{M}(E)$ be atom-free; that is, $\mu(\{x\}) = 0$ for every $x \in E$. Let X be a random measure on E with $\mathbf{P}[X(A) \in \mathbb{N}_0 \cup \{\infty\}] = 1$ for every $A \in \mathcal{B}(E)$. Then the following are equivalent:

- (i) $X \sim \text{PPP}_\mu$.
- (ii) X almost surely has no double points; that is,

$$\mathbf{P}[X(\{x\}) \geq 2 \quad \text{for some } x \in E] = 0$$

and

$$\mathbf{P}[X(A) = 0] = e^{-\mu(A)} \quad \text{for all } A \in \mathcal{B}_b(E) \quad (24.3)$$

Proof (i) $\implies$ (ii) This is obvious.

(ii) $\implies$ (i) If $A_1, \dots, A_n \in \mathcal{B}_b(E)$ are pairwise disjoint, then

$$\begin{aligned} \mathbf{P}[X(A_1) = 0, \dots, X(A_n) = 0] &= \mathbf{P}[X(A_1 \cup \dots \cup A_n) = 0] \\ &= e^{-\mu(A_1 \cup \dots \cup A_n)} \\ &= \prod_{l=1}^n e^{-\mu(A_l)} = \prod_{l=1}^n \mathbf{P}[X(A_l) = 0]. \end{aligned}$$

Hence the random variables $\widetilde{X}(A) := X(A) \wedge 1$ are independent for disjoint sets A . The rest of the proof is similar to that of Theorem 5.34. Let $A \in \mathcal{B}_b(E)$. Choose an $A_0 \subset A$ with $\mu(A_0) = \mu(A)/2$ (this is possible by Exercise 8.3.1 since μ is atom-free) and define $A_1 = A \setminus A_0$.

Similarly, choose $A_{i,0}, A_{i,1} \subset A_i$ for $i = 0, 1$ and inductively define disjoint sets $A_{i,0}, A_{i,1} \subset A_i$ for $i \in \{0, 1\}^{n-1}$ with $\mu(A_i) = 2^{-n}\mu(A)$ for every $i \in \{0, 1\}^n$. Define

$$N_n(A) := \sum_{i \in \{0,1\}^n} \tilde{X}(A_i).$$

As X does not have double points, we have $N_n(A) \uparrow X(A)$ almost surely. On the other hand, by assumption, $N_n(A) \sim b_{2^n, 1 - \exp(-2^{-n}\mu(A))}$ for $n \in \mathbb{N}$; hence the characteristic functions converge:

$$\varphi_{N_n(A)}(t) = \left(1 + \left(1 - e^{-2^{-n}\mu(A)}\right)(e^{it} - 1)\right)^{2^n} \xrightarrow{n \rightarrow \infty} \exp(\mu(A)(e^{it} - 1)) = \varphi_{\text{Poi}_{\mu(A)}}(t)$$

Therefore, we have $\mathbf{P}_{N_n(A)} \xrightarrow{n \rightarrow \infty} \text{Poi}_{\mu(A)}$ and thus $X(A) \sim \text{Poi}_{\mu(A)}$.

If $A_1, \dots, A_k \in \mathcal{B}_b(E)$ are pairwise disjoint, then the sets $N_n(A_1), \dots, N_n(A_k)$ (constructed in a way similar to that above) are independent. Hence also the limits $X(A_l) = \lim_{n \rightarrow \infty} N_n(A_l)$, $l = 1, \dots, k$ are independent.

Theorem 24.14 Let $\mu \in \mathcal{M}(E)$ and let X be a Poisson point process with intensity measure μ . Then X has Laplace transform

$$\mathcal{L}_X(f) = \exp\left(\int \mu(dx) (e^{-f(x)} - 1)\right), \quad f \in \mathcal{B}^+(E)$$

and characteristic function

$$\varphi_X(f) = \exp\left(\int \mu(dx) (e^{if(x)} - 1)\right), \quad f \in \mathcal{B}_b^{\mathbb{R}}(E)$$

Proof It is enough to show the claim for simple functions $f = \sum_{l=1}^n \alpha_l \mathbb{1}_{A_l}$ with complex numbers $\alpha_1, \dots, \alpha_n$ and with pairwise disjoint sets $A_1, \dots, A_n \in \mathcal{B}_b(E)$. (For general f , the claim follows by the usual approximation arguments.) For such f , however,

$$\begin{aligned} \mathbf{E}[\exp(-I_f(X))] &= \mathbf{E}\left[\prod_{l=1}^n e^{-\alpha_l X(A_l)}\right] = \prod_{l=1}^n \mathbf{E}[e^{-\alpha_l X(A_l)}] \\ &= \prod_{l=1}^n \exp(\mu(A_l)(e^{-\alpha_l} - 1)) \\ &= \exp\left(\sum_{l=1}^n \mu(A_l)(e^{-\alpha_l} - 1)\right) \\ &= \exp\left(\int \mu(dx) (e^{-f(x)} - 1)\right) \end{aligned}$$

Corollary 24.15 (Moments of the PPP) Let $\mu \in \mathcal{M}(E)$ and $X \sim \text{PPP}_{\mu}$.

- (i) If $f \in \mathcal{L}^1(\mu)$, then $\mathbf{E}[\int f dX] = \int f d\mu$.
- (ii) If $f \in \mathcal{L}^2(\mu) \cap \mathcal{L}^1(\mu)$, then $\text{Var}[\int f dX] = \int f^2 d\mu$.

Recall that only for finite μ , we have the inclusion $\mathcal{L}^2(\mu) \subset \mathcal{L}^1(\mu)$.

Proof If $f = f^+ - f^- \in \mathcal{L}^1(\mu)$, then for the characteristic function, integral and differentiation interchange, $\frac{d}{dt}\varphi_X(tf^+) = i\varphi_X(tf^+) \int f(x)e^{itf^+(x)}\mu(dx)$ and hence (by Exercise 15.4.4(iii))

$$\mathbf{E}[I_{f^+}(X)] = \frac{1}{i} \frac{d}{dt} \varphi_X(tf^+) \Big|_{t=0} = \int f^+ d\mu$$

Arguing similarly with f^- and adding up, we get (i).

If $f \in \mathcal{L}^1(\mu) \cap \mathcal{L}^2(\mu)$, then the argument can be iterated (using Theorem 15.34)

$$\frac{d^2}{dt^2} \varphi_X(tf) = -\varphi_X(tf) \left[\int f^2(x) e^{itf(x)} \mu(dx) + \left(\int f(x) e^{itf(x)} \mu(dx) \right)^2 \right]$$

hence we have $\mathbf{E}[I_f(X)^2] = -\frac{d^2}{dt^2} \varphi_X(tf) \Big|_{t=0} = I_{f^2}(\mu) + I_f(\mu)^2$.

Theorem 24.16 (Mapping theorem) Let E and F be locally compact Polish spaces and let $\phi : E \rightarrow F$ be a measurable map. Let $\mu \in \mathcal{M}(E)$ with $\mu \circ \phi^{-1} \in \mathcal{M}(F)$ and let X be a PPP on E with intensity measure μ . Then $X \circ \phi^{-1}$ is a PPP on F with intensity measure $\mu \circ \phi^{-1}$.

Proof For $f \in \mathcal{B}^+(F)$,

$$\begin{aligned} \mathcal{L}_{X \circ \phi^{-1}}(f) &= \mathcal{L}_X(f \circ \phi) = \exp \left(\int (e^{-f(\phi(x))} - 1) \mu(dx) \right) \\ &= \exp \left(\int (e^{-f(y)} - 1) (\mu \circ \phi^{-1})(dy) \right) \end{aligned}$$

Now, Theorem 24.14 and Theorem 24.7 yield the claim.

Theorem 24.17 Let $v \in \mathcal{M}((0, \infty))$ and let $X \sim \text{PPP}_v$ on $(0, \infty)$. Further, define $Y := \int xX(dx)$. Then the following are equivalent.

- (i) $\mathbf{P}[Y < \infty] > 0$.
- (ii) $\mathbf{P}[Y < \infty] = 1$.
- (iii) $\int v(dx)(1 \wedge x) < \infty$.

If (i)-(iii) hold, then Y is an infinitely divisible nonnegative random variable with Lévy measure v .

Proof Let $Y_\infty = \int_{[1, \infty)} xX(dx)$ and $Y_t := \int_{(t, 1)} xX(dx)$ for $t \in [0, 1)$. Evidently, $Y = Y_0 + Y_\infty$. Furthermore, it is clear that

$$\mathbf{P}[Y_\infty < \infty] > 0 \iff \mathbf{P}[Y_\infty < \infty] = 1 \iff v([1, \infty)) < \infty. \quad (24.4)$$

If (iii) holds, then $\mathbf{E}[Y_0] = \int_{(0, 1)} xv(dx) < \infty$; hence $Y_0 < \infty$ a.s. (and thus $Y < \infty$ a.s. by (24.4)). On the other hand, if (iii) does not hold, then $Y_\infty = \infty$ a.s. or $\mathbf{E}[Y_0] = \infty$. While Y_∞ can have infinite expectation even if $Y_\infty < \infty$ a.s., for Y_0 this is impossible since, in contrast with Y_∞ , Y_0 is composed not of a few large contributions but many small ones so that a law of large numbers is in force. More precisely, by Corollary 24.15, we have

$$\text{Var}[Y_t] = \int_{(t, 1)} x^2 v(dx) \leq \int_{(t, 1)} xv(dx) = \mathbf{E}[Y_t] < \infty \quad \text{for all } t \in (0, 1)$$

Hence, by Chebyshev's inequality,

$$\mathbf{P} \left[Y_t < \frac{\mathbf{E}[Y_t]}{2} \right] \leq \frac{4\text{Var}[Y_t]}{\mathbf{E}[Y_t]^2} \xrightarrow{t \rightarrow 0} 0.$$

Thus $Y_0 = \sup_{t \in (0, 1)} Y_t \geq \mathbf{E}[Y_0]/2 = \infty$ almost surely.

Now assume that (i)-(iii) hold. By Theorem 24.14, Y has the Laplace transform

$$\mathbf{E}[e^{-tY}] = \exp \left(\int v(dx) (e^{-tx} - 1) \right)$$

By the Lévy-Khinchin formula (Theorem 16.14), Y is infinitely divisible with Lévy measure ν .

Corollary 24.18 Let $\mu_i \in \mathcal{M}_1([0, \infty))$, $i = 1, 2$, be infinitely divisible distributions with canonical measures $\nu_i \in \mathcal{M}((0, \infty))$ and deterministic parts $\alpha_i \geq 0$ (compare Theorem 16.14). If we have

$$\alpha_1 \leq \alpha_2 \quad \text{and} \quad \nu_1([x, \infty)) \leq \nu_2([x, \infty)) \quad \text{for all } x > 0, \quad (24.5)$$

then μ_1 is stochastically smaller than μ_2 ; i.e., $\mu_1 \leq_{st} \mu_2$.

Proof (The proof follows [100, Proof of Lemma 6.1].) The idea is to use a coupling argument where based on one Poisson point process we construct the two random variables Y_1, Y_2 with $Y_i \sim \mu_i$, $i = 1, 2$, such that $Y_1 \leq Y_2$ almost surely. By Theorem 17.58, this yields the claim.

Let $G_i(x) := \nu_i([x, \infty))$, $i = 1, 2$, $x > 0$, and

$$\phi_i(y) := G_i^{-1}(y) = \inf \{x \geq 0 : G_i(x) \leq y\} \quad \text{for } y > 0.$$

If ν_i is finite, then $\phi_i(y) = 0$ for $y \geq \nu_i((0, \infty))$. Let λ denote the Lebesgue measure on $[0, \infty)$. By construction, for the image measure restricted to the positive reals, we have

$$(\lambda \circ \phi_i^{-1})|_{(0, \infty)} = \nu_i, \quad i = 1, 2$$

Now assume that X is a PPP on $(0, \infty)$ with intensity measure λ . By Theorem 24.16, the random measures

$$X_i := \left(\int \phi_i(x) X(dx) \right) \Big|_{(0, \infty)} = X \circ \phi_i^{-1}$$

are PPPs with intensity measures ν_i , $i = 1, 2$. By Theorem 24.17, we thus have

$$Y_i := \alpha_i + \int \phi_i(x) X(dx) \sim \mu_i \quad \text{for } i = 1, 2$$

However, by assumption, we have $G_1 \geq G_2$ which implies $\phi_1 \leq \phi_2$ and thus $Y_1 \leq Y_2$ a.s.

Example 24.19 By Corollary 16.10, for every nonnegative infinitely divisible distribution μ with Lévy measure ν , there exists a stochastic process $(Y_t)_{t \geq 0}$ with independent stationary increments and $Y_t \sim \mu^{*t}$ (hence with Lévy measure $t\nu$). Here we give a direct construction of this process. Let X be a PPP on $(0, \infty) \times [0, \infty)$ with intensity measure $\nu \otimes \lambda$ (here λ is the Lebesgue measure). Define $Y_0 = 0$ and

$$Y_t := \int_{(0, \infty) \times (0, t]} x X(d(x, s))$$

By the mapping theorem, we have $X(\cdot \times (s, t]) \sim \text{PPP}_{(t-s)\nu}$; hence $Y_t - Y_s$ is infinitely divisible with Lévy measure $(t-s)\nu$. The independence of the increments is evident. Note that $t \mapsto Y_t$ is right continuous and monotone increasing.

The process Y that we have just constructed is called a subordinator with Lévy measure ν .

The procedure in the previous example can be generalized by allowing time sets more general than $[0, \infty)$.

Definition 24.20 A random measure Y is called infinitely divisible if, for any $n \in \mathbb{N}$, there exist i.i.d. random measures $Y_1, \dots, Y_n$ with $Y = Y_1 + \dots + Y_n$.

Theorem 24.21 Let $\nu \in \mathcal{M}((0, \infty) \times E)$ with

$$\int \mathbb{1}_A(t)(1 \wedge x)\nu(d(x, t)) < \infty \quad \text{for all } A \in \mathcal{B}_b(E)$$

and let $\alpha \in \mathcal{M}(E)$. Let X be a PPP_ν and

$$Y(A) := \alpha(A) + \int x \mathbb{1}_A(t) X(d(x, t)) \quad \text{for } A \in \mathcal{B}(E)$$

Then Y is an infinitely divisible random measure with independent increments. For $A \in \mathcal{B}(E)$, $Y(A)$ has the Lévy measure $v(\cdot \times A)$.

We call ν the canonical measure and α the deterministic part of Y .

Proof This is a direct consequence of Theorem 24.16 and Theorem 24.17.

Remark 24.22 We can write Y as $Y = \alpha + \int x \delta_t X(d(x, t))$, where δ_t is the Dirac measure at $t \in E$. If instead of $x \delta_t$, we allow more general measures $\chi \in \mathcal{M}(E)$, then we get a representation

$$Y = \alpha + \int_{\mathcal{M}(E)} \chi X(d\chi)$$

where $X \sim \text{PPP}_\nu$ on $\mathcal{M}(E)$ and $\nu \in \mathcal{M}(\mathcal{M}(E))$ with

$$\int v(d\chi)(\chi(A) \wedge 1) < \infty$$

for any $A \in \mathcal{B}_b(E)$. It can be shown that this is the most general form of an infinitely divisible measure on E . We call ν the canonical measure of Y and α the deterministic part. Y is characterized by its Laplace transform which obeys the Lévy-Khinchin formula:

$$\mathcal{L}_Y(f) = \exp \left(- \int f d\alpha + \int v(d\chi) \left(e^{-\int f d\chi} - 1 \right) \right)$$

Theorem 24.23 (Coloring theorem) Let F be a further locally compact Polish space, let $\mu \in \mathcal{M}(E)$ be atom-free and let $(Y_x)_{x \in E}$ be i.i.d. random variables, independent of X , with values in F and distribution $v \in \mathcal{M}_1(F)$. Then

$$Z(A) := \int \mathbb{1}_A(x, Y_x) X(dx), \quad A \in \mathcal{B}(E \times F)$$

is a $\text{PPP}_{\mu \otimes v}$ on $E \times F$.

Proof This is left as an exercise.

There is an obvious generalization of the coloring theorem: The assumption that μ is atom-free was needed in order that X have no double points. That is, for every unit mass that X produces, there is a different random variable Y_x . However, this can also be achieved by different means and in somewhat greater generality.

Accordingly, let E, F be locally compact Polish spaces, let $\mu \in \mathcal{M}(E)$ and let κ be a stochastic kernel from E to F with $\mu\kappa := \int \mu(dx) \kappa(x, \cdot) \in \mathcal{M}(F)$. Let $(Y_{x,t})_{x \in E, t \in [0,1]}$ be independent random variables with distributions $\mathbf{P}_{Y_{x,t}} = \kappa(x, \cdot)$ for $x \in E$ and $t \in [0, 1]$.

For $X \sim \text{PPP}_\mu$, define the lifting $\tilde{X}$ as that PPP on $E \times [0, 1]$ with intensity measure $\mu \otimes \lambda|_{[0,1]}$, where λ is the Lebesgue measure. Clearly, $X \stackrel{\mathcal{D}}{=} \tilde{X}(\cdot \times [0, 1])$. The random measure $\tilde{X}$ can be understood as a realization of X in which the different points of X are assigned arbitrary $[0, 1]$ -valued labels to distinguish them. Now let

$$X^\kappa(A) := \int \tilde{X}(d(x, t)) \mathbb{1}_A(Y_{x,t}) \quad \text{for } A \in \mathcal{B}(F)$$

Theorem 24.24 X^κ is a random measure with $\mathbf{P}_{X^\kappa} = \text{PPP}_{\mu\kappa}$.

Proof Clearly, almost surely X^κ is a measure. For $A \in \mathcal{B}_b(F)$, we have by assumption

$$\mathbf{E}[X^\kappa(A)] = \mathbf{E} \left[\int \tilde{X}(d(x, t)) \kappa(x, A) \right] = (\mu\kappa)(A) < \infty$$

Hence $X^\kappa(A) < \infty$ almost surely, and thus X^κ is a random measure. We compute the Laplace transform of X^κ . Let $g(x) := -\log \mathbf{E}[e^{-f(Y_{x,t})}]$. Then (since $\tilde{X}$ has no double points)

$$\begin{aligned}
 \mathcal{L}_{X^\kappa}(f) &= \mathbf{E} \left[\exp \left(- \int \tilde{X}(d(x,t)) f(Y_{x,t}) \right) \right] \\
 &= \mathbf{E} \left[\prod_{(x,t): \tilde{X}(\{(x,t)\})=1} e^{-f(Y_{x,t})} \right] \\
 &= \mathbf{E} \left[\prod_{(x,t): \tilde{X}(\{(x,t)\})=1} \mathbf{E}[e^{-f(Y_{x,t})}] \right] \\
 &= \mathbf{E} \left[\prod_{(x,t): \tilde{X}(\{(x,t)\})=1} e^{-g(x)} \right] = \mathcal{L}_X(g) \\
 &= \exp \left(\int \mu(dx) (\mathbf{E}[e^{-f(Y_{x,t})}] - 1) \right) \\
 &= \exp \left(\int \mu(dx) \int \kappa(x, dy) (e^{-f(y)} - 1) \right) \\
 &= \exp \left(\int \mu \kappa(dy) (e^{-f(y)} - 1) \right)
 \end{aligned}$$

Example 24.25 (PPP as invariant distribution) As an application of the previous theorem, consider a stochastic process on $E = \mathbb{Z}^d$ or $E = \mathbb{R}^d$ that consists of a system of independent random walks. Hence assume that we are given i.i.d. random variables $Z_n^i, i, n \in \mathbb{N}$ with distribution $v \in \mathcal{M}_1(E)$. Further, assume that, at time n , the position of the i th particle of our system of random walks is $S_n^i := S_0^i + \sum_{l=1}^n Z_l^i$, where S_0^i is an arbitrary, possibly random, starting point. Assume that the particles are indistinguishable. Hence we simply add the particles at each site:

$$X_n(A) := \sum_{i=1}^{\infty} \mathbb{1}_A(S_n^i) \quad \text{for } A \subset E$$

Each X_n is a measure on E and, if at the beginning the particles are not too concentrated locally, it is a locally finite measure and hence a random measure. Assume that $X_0 \sim \text{PPP}_\mu$ for some $\mu \in \mathcal{M}(E)$. Define $\kappa(x, \cdot) = \delta_x * \nu$, and write κ^n for the n -fold application of κ ; that is, $\kappa^n(x, \cdot) = \delta_x * \nu^{*n}$. We thus get $X_0^\kappa \stackrel{\mathcal{D}}{=} X_1$. Indeed, independence of the motions of the individual particles in the definition of X_0^κ is exactly independence of the random walks. As X_1 is also a PPP, we get inductively $X_n^\kappa \stackrel{\mathcal{D}}{=} X_{n+1}$ and thus $X_n \sim \text{PPP}_{\mu \kappa^n} = \text{PPP}_{\mu * \nu^{*n}}$. In particular, $X_0 \stackrel{\mathcal{D}}{=} X_n$ if and only if $\mu * \nu = \mu$. Clearly, this is true if we have $E = \mathbb{Z}^d$ and μ the counting measure or if $E = \mathbb{R}^d$ and μ is the Lebesgue measure. For example, if $E = \mathbb{Z}^d$, then under rather mild assumptions on ν one can show that the counting measure $\mu = \lambda$ is the unique (up to multiples) solution of $\mu * \nu = \mu$. In this case, every invariant measure is a convex combination of PPPs with different intensity measures $\theta\lambda$.

Exercise 24.2.1 Use an approximation with simple functions in order to show the claim of Corollary 24.15 without using characteristic functions.

Exercise 24.2.2 Prove the coloring theorem (Theorem 24.23).

Exercise 24.2.3 Let $p_1, p_2 \in (0, 1]$ and $r_1, r_2 > 0$. Show the following statement about the stochastic order of negative binomial distributions: $b_{r_1, p_1}^- \leq_{\text{st}} b_{r_2, p_2}^-$ if and only if

$$p_1 \geq p_2 \quad \text{and} \quad p_1^{r_1} \geq p_2^{r_2}$$

5.24.3 24.3 The Poisson-dirichlet Distribution*

The goal of this section is to solve the following problem. Take a stick of length 1. Choose a point of the stick uniformly at random and break the stick at this point. Put the left part of the stick (with length, say, W_1) aside. With the remaining part of the stick proceed just as with the original stick. Break it in two and put the left part (of length W_2) aside. Successively, we thus collect fractions of the stick of lengths $W_1, W_2, W_3, \dots$. What is the joint distribution of $(W_1, W_2, \dots)$? Furthermore, if we order the numbers $W_1, W_2, \dots$ in decreasing order $W_{(1)} \geq W_{(2)} \geq \dots$, what is the distribution of $(W_{(1)}, W_{(2)}, \dots)$? And finally, why do we ask these questions in a chapter on Poisson point processes?

Answering these questions requires some preparation. We saw that the Beta distribution occurs naturally in Pólya's urn model as the limiting distribution of the fraction of balls of a given color. Clearly, Pólya's urn model can be considered for any number $n \geq 2$ of colors. The limiting distribution is then the n -dimensional generalization of the Beta distribution, namely the so-called Dirichlet distribution.

Define the $(n-1)$ -dimensional simplex

$$\Delta_n := \{(x_1, \dots, x_n) \in [0, 1]^n : x_1 + \dots + x_n = 1\}.$$

Definition 24.26 Let $n \in \{2, 3, \dots\}$ and $\theta_1, \dots, \theta_n > 0$. The Dirichlet distribution $\text{Dir}_{\theta_1, \dots, \theta_n}$ is the distribution on Δ_n that is defined for measurable $A \subset \Delta_n$ by

$$\text{Dir}_{\theta_1, \dots, \theta_n}(A) = \int \mathbb{1}_A(x_1, \dots, x_n) f_{\theta_1, \dots, \theta_n}(x_1, \dots, x_n) dx_1 \dots dx_{n-1}$$

Here

$$f_{\theta_1, \dots, \theta_n}(x_1, \dots, x_n) = \frac{\Gamma(\theta_1 + \dots + \theta_n)}{\Gamma(\theta_1) \dots \Gamma(\theta_n)} x_1^{\theta_1-1} \dots x_n^{\theta_n-1}.$$

If the parameters $\theta_1, \dots, \theta_n$ are integer-valued, they correspond to the numbers of balls of the different colors that are originally in the urn. Assume that the colors $n-1$ and n are light green and green and that in the dim light we cannot distinguish them. Then we should still end up with a Dirichlet distribution in the limit but with $n-1$ instead of n and with $\theta_{n-1} + \theta_n$ instead of θ_{n-1} and θ_n ; that is, $\text{Dir}_{\theta_1, \dots, \theta_{n-2}, \theta_{n-1} + \theta_n}$. Let $(M_t)_{t \geq 0}$ be the Moran Gamma subordinator, the stochastic process with right continuous, monotone increasing paths $t \mapsto M_t$ and independent, stationary, Gamma-distributed increments: $M_t - M_s \sim \Gamma_{1, t-s}$ for $t > s \geq 0$. An important connection between M and the Dirichlet distribution is revealed by the corollaries of the following theorem and by Theorem 24.32.

Theorem 24.27 Let $n \in \mathbb{N}$, $\theta_1, \dots, \theta_n > 0$ and $\Theta := \theta_1 + \dots + \theta_n$. Let $X \sim \text{Dir}_{\theta_1, \dots, \theta_n}$ and let $Z \sim \Gamma_{1, \Theta}$ be independent random variables. Then the random variables $S_i := Z \cdot X_i, i = 1, \dots, n$ are independent and $S_i \sim \Gamma_{1, \theta_i}$.

Proof In the following, always let $x_n := 1 - \sum_{i=1}^{n-1} x_i$ and $s = \sum_{j=1}^n s_j$. Let

$$\Delta'_n := \left\{ (x_1, \dots, x_{n-1}) \in (0, 1)^{n-1} : \sum_{i=1}^{n-1} x_i < 1 \right\}$$

For $x \in \Delta'_n$ and $z \geq 0$, the distribution of $(X_1, \dots, X_{n-1}, Z)$ has the density

$$f(x_1, \dots, x_{n-1}, z) = \prod_{j=1}^n \left(x_j^{\theta_j-1} / \Gamma(\theta_j) \right) z^{\Theta-1} e^{-z}$$

Consider the map

$$F : \Delta'_n \times (0, \infty) \rightarrow (0, \infty)^n, \quad (x_1, \dots, x_{n-1}, z) \mapsto (zx_1, \dots, zx_n)$$

This map is invertible with inverse map

$$F^{-1} : (s_1, \dots, s_n) \mapsto (s_1/s, \dots, s_{n-1}/s, s).$$

The Jacobian determinant of F is $\det(F'(x_1, \dots, x_{n-1}, z)) = z^{n-1}$. By the transformation formula for densities (Theorem 1.101), $(S_1, \dots, S_n)$ has density

$$\begin{aligned} g(s_1, \dots, s_n) &= \frac{f(F^{-1}(s_1, \dots, s_n))}{|\det(F'(F^{-1}(s_1, \dots, s_n)))|} \\ &= \frac{s^{\Theta-1} e^{-s}}{s^{n-1}} \prod_{j=1}^n \left((s_j/s)^{\theta_j-1} / \Gamma(\theta_j) \right) \\ &= \prod_{j=1}^n \left(s_j^{\theta_j-1} e^{-s_j} / \Gamma(\theta_j) \right) \end{aligned}$$

However, this is the density for independent Gamma distributions.

Corollary 24.28 If $t_i := \sum_{j=1}^i \theta_j$ for $i = 0, \dots, n$, then the random variables $X = ((M_{t_i} - M_{t_{i-1}}) / M_{t_n}, i = 1, \dots, n)$ and $S := M_{t_n}$ are independent with distributions $X \sim \text{Dir}_{\theta_1, \dots, \theta_n}$ and $S \sim \Gamma_{1, t_n}$.

Corollary 24.29 Let $(X_1, \dots, X_n) \sim \text{Dir}_{\theta_1, \dots, \theta_n}$. Then $X_1 \sim \beta_{\theta_1, \sum_{i=2}^n \theta_i}$ and $(X_2 / (1 - X_1), \dots, X_n / (1 - X_1)) \sim \text{Dir}_{\theta_2, \dots, \theta_n}$ are independent.

Proof Let M be as in Corollary 24.28. Then $X_1 = M_{t_1} / M_{t_n} \sim \beta_{\theta_1, t_n - \theta_1}$. Since $X_1 = \left(\frac{M_{t_n} - M_{t_1}}{M_{t_1}} + 1 \right)^{-1}$, we see that X_1 depend only on M_{t_1} and $M_{t_n} - M_{t_1}$. On the other hand,

$$\left(\frac{X_2}{1 - X_1}, \dots, \frac{X_n}{1 - X_1} \right) = \left(\frac{M_{t_2} - M_{t_1}}{M_{t_n} - M_{t_1}}, \dots, \frac{M_{t_n} - M_{t_{n-1}}}{M_{t_n} - M_{t_1}} \right)$$

is independent of M_{t_1} . By Corollary 24.28, it is also independent of $M_{t_n} - M_{t_1}$ and is $\text{Dir}_{\theta_2, \dots, \theta_n}$ -distributed. $\square$

Corollary 24.30 Let $V_1, \dots, V_{n-1}$ be independent, $V_i \sim \beta_{\theta_i, \theta_{i+1} + \dots + \theta_n}$ and $V_n = 1$. Then

$$\left(V_1, (1 - V_1) V_2, (1 - V_1) (1 - V_2) V_3, \dots, \left(\prod_{i=1}^{n-1} (1 - V_i) \right) V_n \right) \sim \text{Dir}_{\theta_1, \dots, \theta_n}.$$

Proof This follows by iterating the claim of Corollary 24.29. $\square$

It is natural to ask what happens if we distinguish more and more colors (instead of pooling them). For simplicity, consider a symmetric situation where we have $\theta_1 = \dots = \theta_n = \theta/n$ for some $\theta > 0$. Hence we consider

$$\text{Dir}_{\theta; n} := \text{Dir}_{\theta, \dots, \theta} \quad \text{for } \theta > 0.$$

If $X^n = (X_1^n, \dots, X_n^n) \sim \text{Dir}_{\theta/n; n}$, then, by symmetry, we have $\mathbf{E}[X_i^n] = 1/n$ for every $n \in \mathbb{N}$ and $i = 1, \dots, n$. Hence, clearly $(X_1^n, \dots, X_k^n) \xrightarrow{n \rightarrow \infty} 0$ for any $k \in \mathbb{N}$. In order to obtain a nontrivial limit, one possibility is to reorder the values by decreasing size: $X_{(1)}^n \geq X_{(2)}^n \geq \dots$

Definition 24.31 Let $\theta > 0$ and let $(M_t)_{t \in [0, \theta]}$ be a Moran Gamma subordinator. Let $m_1 \geq m_2 \geq \dots \geq 0$ be the jump sizes of M in decreasing order and let $\tilde{m}_i = m_i/M_\theta, i = 1, 2, \dots$. The distribution of the random variables $(\tilde{m}_1, \tilde{m}_2, \dots)$ on $S := \{(x_1 \geq x_2 \geq \dots \geq 0) : x_1 + x_2 + \dots = 1\}$ is called the Poisson-Dirichlet distribution PD_θ with parameter $\theta > 0$.

To be honest, we still have to show that $\sum_{i=1}^{\infty} \tilde{m}_i = 1$. To this end, let Y be a PPP on $(0, \infty) \times (0, \theta]$ with intensity measure $\nu \otimes \lambda$, where λ is the Lebesgue measure and $\nu(dx) = e^{-x}x^{-1}dx$ is the Lévy measure of the $\Gamma_{1,1}$ distribution. We can define M by $M_t := \sum_{(x,s): Y(\{x,s\})=1, s \leq t} x$. Now we have $m_1 = \sup\{x \in (0, \infty) : Y(\{x\} \times (0, \theta]) = 1\}$. Inductively, we get $m_n = \sup\{x < m_{n-1} : Y(\{x\} \times (0, \theta]) = 1\}$ for $n \geq 2$. Interchanging the order of summations, we obtain $M_\theta = \sum_{n=1}^{\infty} m_n$.

Theorem 24.32 If $X^n \sim \text{Dir}_{\theta/n;n}$ for $n \in \mathbb{N}$, then $\mathbf{P}_{(X_{(1)}^n, X_{(2)}^n, \dots)} \xrightarrow{n \rightarrow \infty} \text{PD}_\theta$.

Proof The idea is to express the random variables $X^n, n \in \mathbb{N}$, in terms of the increments of the Moran Gamma subordinator $(M_t)_{t \in [0, \theta]}$ in such a way that convergence of distributions implies almost sure convergence. Hence, let $X_i^n = (M_{\theta i/n} - M_{\theta(i-1)/n})/M_\theta$. By Corollary 24.28, we have $X^n \sim \text{Dir}_{\theta/n;n}$. Let $t_1, t_2, \dots \in (0, \theta]$ be the positions of the jumps $m_1 \geq m_2 \geq \dots$. Evidently, $X_{(1)}^n \geq \tilde{m}_1$ for every n . If n is large enough that $|t_1 - t_2| > \theta/n$, then $X_{(2)}^n \geq \tilde{m}_2$. Inductively, we get

$\liminf_{n \rightarrow \infty} X_{(i)}^n \geq \tilde{m}_i$ almost surely. Using the convention $X_{(i)}^n = 0$ for $i > n$, we have $\sum_{i=1}^{\infty} X_{(i)}^n = 1$ for every $n \in \mathbb{N}$. By Fatou's lemma, we thus get

$$1 = \sum_{i=1}^{\infty} \tilde{m}_i \leq \sum_{i=1}^{\infty} \liminf_{n \rightarrow \infty} X_{(i)}^n \leq \liminf_{n \rightarrow \infty} \sum_{i=1}^{\infty} X_{(i)}^n = 1.$$

Therefore, $\lim_{n \rightarrow \infty} X_{(i)}^n = \tilde{m}_i$ almost surely.

Instead of ordering the values of X^n by their sizes, there is a different way of arranging the terms so that the distributions converge. Think of a biological population in which a certain phenotypical property can be measured with different levels of precision. If we distinguish n different values of this property, then we write X_i^n for the proportion of the population that has type $i \in \{1, \dots, n\}$.

Now successively choose individuals from the population at random. Let I_1^n be the type of the first individual. Denote by I_2^n the type of the first individual that is not of type I_1^n . That is, I_2^n is the second type that we see. Now inductively define I_k^n as the k th type that we see; that is, the type of the first individual that has none of the types $I_1^n, \dots, I_{k-1}^n$. Consider the vector $\hat{X}^n = (\hat{X}_1^n, \dots, \hat{X}_n^n)$, where $\hat{X}_k^n = X_{I_k^n}^n$. Since the probability of the event $\{I_1^n = i\}$ is proportional to the size of the subpopulation of type i , we say that $\hat{X}^n$ is the successively size-biased vector.

The distribution of $\hat{X}^n$ does not change if we change the order of the $X_1^n, \dots, X_n^n$. For example, instead of $X_1^n, \dots, X_n^n$, we can use the order statistics $(X_{(1)}^n, \dots, X_{(n)}^n)$ and again end up with $\hat{X}^n$ as the successively size-biased vector. Hence we can define the successively size-biased vector $\hat{X}$ for the infinite vector $X \sim \text{PD}_\theta$. If $X^n \sim \text{Dir}_{\theta/n;n}$, then by Theorem 24.32, we have $\hat{X}^n \xrightarrow{n \rightarrow \infty} \hat{X}$. Hence we can compute the distribution of $\hat{X}$.

Theorem 24.33 Let $\theta > 0$ and $X^n \sim \text{Dir}_{\theta/n;n}, n \in \mathbb{N}$. Let $X \sim \text{PD}_\theta$. Further, let $V_1, V_2, \dots$ be i.i.d. random variables on $[0, 1]$ with density $x \mapsto \theta(1-x)^{\theta-1}$. Define $Z_1 = V_1$ and $Z_k = \left(\prod_{i=1}^{k-1} (1 - V_i)\right) V_k$ for $k \geq 2$. Then:

- (i) $\hat{X}^n \xrightarrow{n \rightarrow \infty} \hat{X}$.
- (ii) $\hat{X} \stackrel{\mathcal{D}}{=} Z$.

The distribution of Z is called the GEM_θ distribution (Griffiths-Engen-McCloskey).

Proof Statement (i) was shown in the discussion preceding the theorem. In order to show (ii),

we compute the distribution of $\hat{X}^n$ and show that it converges to the distribution of Z .

Let $\hat{X}^{n,1}$ be the vector $X^{n,1} = (X_{I_1^n}^n, X_2, \dots, X_{I_1^n-1}^n, X_{I_1^n+1}^n, \dots, X_n^n)$, in which only the first coordinate is sampled size-biasedly. We show that

$$\hat{X}^{n,1} \sim \text{Dir}_{(\theta/n)+1, \theta/n, \dots, \theta/n}. \quad (24.6)$$

Let $f(x) = (\Gamma(\theta)/\Gamma(\theta/n)^n) \cdot \prod_{k=1}^n x_k^{(\theta/n)-1}$ be the density of $\text{Dir}_{\theta/n; n}$. We compute the density $f^{n,1}$ of $X^{n,1}$ by decomposing according to the value i of I_1^n :

$$\begin{aligned} f^{n,1}(x) &= \sum_{i=1}^n x_1 f(x_2, \dots, x_i, x_1, x_{i+1}, \dots, x_n) = n x_1 f(x) \\ &= \frac{n \Gamma(\theta)}{\Gamma(\theta/n)^n} x_1^{\theta/n} \prod_{i=2}^n x_i^{(\theta/n)-1} \\ &= \frac{\Gamma(\theta+1)}{\Gamma((\theta/n)+1) \Gamma(\theta/n)^{n-1}} x_1^{\theta/n} \prod_{i=2}^n x_i^{(\theta/n)-1}. \end{aligned}$$

However, this is the density of $\text{Dir}_{(\theta/n)+1, \theta/n, \dots, \theta/n}$. By Corollary 24.29, we have

$$\hat{X}^{n,1} \stackrel{\mathcal{D}}{=} (V_1^n, (1 - V_1^n) Y_1, \dots, (1 - V_1^n) Y_{n-1}),$$

where

$$V_1^n \sim \beta_{(\theta/n)+1, \theta(n-1)/n} \quad \text{and} \quad Y = (Y_1, \dots, Y_{n-1}) \sim \text{Dir}_{\theta/n; n-1}$$

are independent. Applying this to Y , we get inductively

$$\hat{X}^n \stackrel{\mathcal{D}}{=} Z^n \quad (24.7)$$

where

$$Z_1^n = V_1^n \quad \text{and} \quad Z_k^n = \left(\prod_{i=1}^{k-1} (1 - V_i^n) \right) V_k^n \quad \text{for } k \geq 2$$

and where $V_1^n, \dots, V_{n-1}^n$ are independent and $V_i^n \sim \beta_{(\theta/n)+1, \theta(n-i)/n}$. Now it is easy to check that $\beta_{(\theta/n)+1, \theta(n-i)/n} \xrightarrow{n \rightarrow \infty} \beta_{1, \theta}$ for every $i \in \mathbb{N}$. Recall that $\beta_{1, \theta}$ has the density $x \mapsto \theta(1-x)^{\theta-1}$. Hence $V_i^n \xrightarrow{n \rightarrow \infty} V_i$ for every i and thus $Z^n \xrightarrow{n \rightarrow \infty} Z$ and $\hat{X}^n \xrightarrow{n \rightarrow \infty} Z$. Together with (i), this proves claim (ii).

At the beginning of this chapter, we raised the question of how the sizes $W_1, W_2, \dots$ of the stick lengths are distributed if at each step, we break the remaining part of the stick at a point chosen uniformly at random. The preceding theorem gives the answer: The vector $(W_{(1)}, W_{(2)}, \dots)$ has distribution PD_1 , and $(W_1, W_2, \dots)$ has distribution GEM_1 .

5.24.4 The Chinese Restaurant Process

We will study a further situation in which the Poisson-Dirichlet distribution arises naturally. As the technical details get a bit tricky, we content ourselves with the description of the problem and with stating (but not proving) two fundamental theorems. An excellent reference for this type of problem is [130].

Consider a Chinese restaurant with countably many enumerated round tables. At each table, there is enough space for arbitrarily many guests. Initially, the restaurant is empty. One by one an infinite number of guests arrive. The first guest sits down at table number one. If

there are already n guests sitting at k tables, then the $(n+1)$ th guest can choose between sitting down at any of the k occupied tables or at the free table with the smallest number (that is, $k+1$). Assume that the guest makes his choice at random (and independently of the previous choices of the other guests). For $l \leq k$, denote by N_l^n the number of guests at the l th table and assume that the probability of choosing the l th table is $(N_l^n - \alpha)/(n + \theta)$. Then the probability of choosing the first free table is $(\theta + k\alpha)/(n + \theta)$. Here $\alpha \in [0, 1]$ and $\theta > -\alpha$ are parameters. We say that $(N^n)_{n \in \mathbb{N}} = (N_1^n, N_2^n, \dots)_{n \in \mathbb{N}}$ is the Chinese restaurant process with parameters (α, θ) .

In the special case $\alpha = 0$, there is a nice interpretation: Assume that the new guest can also choose his seating position at the table (that is, his neighbor to the right). Then, for any of the present guests, the probability of being chosen as a right neighbor is $1/(n + \theta)$. The probability of starting a new table is $\theta/(n + \theta)$.

In order to study the large n behavior of $N^n/n = (N_1^n/n, N_2^n/n, \dots)$, we extend the Poisson-Dirichlet distribution and the GEM distribution by a further parameter.

Definition 24.34 Let $\alpha \in [0, 1]$ and $\theta > -\alpha$. Let $V_1, V_2, \dots$ be independent and $V_i \sim \beta_{1-\alpha, \theta+i\alpha}$. Define $Z = (Z_1, Z_2, \dots)$ by $Z_1 = V_1$ and

$$Z_k = V_k \prod_{i=1}^{k-1} (1 - V_i) \quad \text{for } k \geq 2$$

Then $\text{GEM}_{\alpha, \theta} := \mathbf{P}_Z$ is called the GEM distribution with parameters (α, θ) . The distribution of the size-biased vector $(Z_{(1)}, Z_{(2)}, \dots)$ is called the Poisson-Dirichlet distribution with parameters (α, θ) , or briefly $\text{PD}_{\alpha, \theta}$.

Explicit formulas for the densities of the finite-dimensional marginals of $\text{PD}_{\alpha, \theta}$ can be found in [132]. Note that, for $\alpha = 0$, we recover the classical distributions $\text{GEM}_\theta = \text{GEM}_{0, \theta}$ and $\text{PD}_\theta = \text{PD}_{0, \theta}$.

Theorem 24.35 Let $\alpha \in [0, 1], \theta > -\alpha$ and let $(N^n)_{n \in \mathbb{N}}$ be the Chinese restaurant process with parameters (α, θ) . Then $\mathbf{P}_{N^n/n} \xrightarrow{n \rightarrow \infty} \text{PD}_{\alpha, \theta}$.

Proof See [129] or [130, Theorem 25].

As for the one-parameter Poisson-Dirichlet distribution, there is a representation of $\text{PD}_{\alpha, \theta}$ in terms of the size-ordered jumps of a certain subordinator. In the following, let $\alpha \in (0, 1)$ and let $(M_t)_{t \in [0, 1]}$ be an α -stable subordinator; that is, a subordinator with Lévy measure $v(dx) = x^{-\alpha-1}dx$. Further, let $m_1 \geq m_2 \geq \dots \geq 0$ be the jumps of M , $\tilde{m}_i = m_i/M_1$ for $i \in \mathbb{N}$, and $\tilde{m} = (\tilde{m}_1, \tilde{m}_2, \dots)$. We quote the following theorem from [130, Section 4.2].

Theorem 24.36 Let $\alpha \in (0, 1)$.

- (i) $\tilde{m} \sim \text{PD}_{\alpha, 0}$.
- (ii) If $\theta > -\alpha$, then $\text{PD}_{\alpha, \theta} \ll \text{PD}_{\alpha, 0} = \mathbf{P}[\tilde{m} \in \cdot]$ and the density is

$$\text{PD}_{\alpha, \theta}(dx) = \frac{M_1^{-\theta}}{\mathbf{E}[M_1^{-\theta}]} \mathbf{P}[\tilde{m} \in dx]$$

Exercise 24.3.1 Let $(X, 1 - X) \sim \text{Dir}_{\theta_1, \theta_2}$. Show that $X \sim \beta_{\theta_1, \theta_2}$ is Beta-distributed.

Exercise 24.3.2 Let $X = (X_1, \dots, X_n) \sim \text{Dir}_{\theta_1, \dots, \theta_n}$. Show the following.

- (i) For any permutation σ on $\{1, \dots, n\}$, we have

$$(X_{\sigma(1)}, \dots, X_{\sigma(n)}) \sim \text{Dir}_{\theta_{\sigma(1)}, \dots, \theta_{\sigma(n)}}$$

- (ii) $(X_1, \dots, X_{n-2}, X_{n-1} + X_n) \sim \text{Dir}_{\theta_1, \dots, \theta_{n-2}, \theta_{n-1} + \theta_n}$.

Exercise 24.3.3 Let $(N^n)_{n \in \mathbb{N}}$ be the Chinese restaurant process with parameters $(0, \theta)$.

- (i) Let $\theta = 1$.

- (a) Show that $\mathbf{P}[N_1^n = k] = 1/n$ for any $k = 1, \dots, n$,
 (b) Show that, for $k_l = 1, \dots, n - (k_1 + \dots + k_{l-1})$,

$$\mathbf{P}[N_l^n = k_l \mid N_1^n = k_1, \dots, N_{l-1}^n = k_{l-1}] = \frac{1}{n - (k_1 + \dots + k_{l-1})}.$$

- (c) Infer the claim of Theorem 24.35 in the case $\alpha = 0$ and $\theta = 1$.
 (ii) Let $\theta > 0$.
 (a) Show that $n\mathbf{P}[N_1^n = \lfloor nx \rfloor] \xrightarrow{n \rightarrow \infty} \theta(1-x)^{\theta-1}$ for $x \in (0, 1)$.
 (b) Show that

$$n\mathbf{P}[N_l^n = \lfloor nx_l \rfloor \mid N_1^n = \lfloor nx_1 \rfloor, \dots, N_{l-1}^n = \lfloor nx_{l-1} \rfloor] \xrightarrow{n \rightarrow \infty} (\theta/y_l)(1 - x_l/y_l)^{\theta-1}$$

for $x_1, \dots, x_l \in (0, 1)$ with $y_l = 1 - (x_1 + \dots + x_{l-1}) > x_l$.

- (c) As in (i), infer the claim of Theorem 24.35 for $\alpha = 0$ and $\theta > 0$.

5.25 25 The Itô Integral

The Itô integral allows us to integrate stochastic processes with respect to the increments of a Brownian motion or a somewhat more general stochastic process. We develop the Itô integral first for Brownian motion and then for generalized diffusion processes (so called Itô processes). In the third section, we derive the celebrated Itô formula. This is the chain rule for the Itô integral that enables us to do explicit calculations with the Itô integral. In the fourth section, we use the Itô formula to obtain a stochastic solution of the classical Dirichlet problem. This in turn is used in the fifth section in order to show that like symmetric simple random walk, Brownian motion is recurrent in low dimensions and transient in high dimensions.

5.25.1 25.1 Itô Integral with Respect to Brownian Motion

Let $W = (W_t)_{t \geq 0}$ be a Brownian motion on the space $(\Omega, \mathcal{F}, \mathbf{P})$ with respect to the filtration $\mathbb{F}$ that satisfies the usual conditions (see Definition 21.22). That is, W is a Brownian motion and an $\mathbb{F}$ -martingale. The aim of this section is to construct an integral

$$I_t^W(H) = \int_0^t H_s dW_s$$

for a large class of integrands $H : \Omega \times [0, \infty) \rightarrow \mathbb{R}$, $(\omega, t) \mapsto H_t(\omega)$ in such a way that $(I_t^W(H))_{t \geq 0}$ is a continuous $\mathbb{F}$ -martingale. Since almost all paths $s \mapsto W_s(\omega)$ of Brownian motion are of locally infinite variation, $W(\omega)$ is not the distribution function of a signed Lebesgue-Stieltjes measure on $[0, \infty)$. Hence $I_t^W(H)$ cannot be defined in the framework of classical integration theory. The basic new idea is to establish the integral as an L^2 -limit. We start with an elementary example to illustrate this.

Example 25.1 Assume that $X_1, X_2, \dots$ are i.i.d. $\text{Rad}_{1/2}$ random variables; that is, $\mathbf{P}[X_n = 1] = \mathbf{P}[X_n = -1] = \frac{1}{2}$. Let $(h_n)_{n \in \mathbb{N}}$ be a sequence of real numbers. Under which assumptions on $(h_n)_{n \in \mathbb{N}}$ is the series

$$R := \sum_{n \in \mathbb{N}} h_n X_n \tag{25.1}$$

well-defined? If $\sum_{n \in \mathbb{N}} |h_n| < \infty$, then the series converges absolutely for every ω . In this case, there is no problem. Now assume that only the weaker condition $\sum_{n \in \mathbb{N}} h_n^2 < \infty$ holds. In this case, the series (25.1) does not necessarily converge any more for every ω . However, we have

$\mathbf{E}[h_n X_n] = 0$ for each $n \in \mathbb{N}$ and $\sum_{n=1}^{\infty} \text{Var}[h_n X_n] = \sum_{n=1}^{\infty} h_n^2 < \infty$. Hence $R_N := \sum_{k=1}^N h_k X_k$ converges in L^2 (for $N \rightarrow \infty$). We can thus define the series R in (25.1) as the L^2 -limit of the partial sums R_N . Note that (at least formally) for the approximating sums the order of summation matters. In a sense, we have constructed $\sum_{n=1}^{\infty}$ instead of $\sum_{n \in \mathbb{N}}$.

An equivalent formulation that gives a flavor of what is to come is the following. Denote by ℓ^2 the Hilbert space of square summable sequences of real numbers with inner product $\langle h, g \rangle = \sum_{n=1}^{\infty} h_n g_n$ and norm $\|g\| = \langle g, g \rangle^{1/2}$. Let ℓ^f be the subspace of those sequences with only finitely many nonzero entries. Then $R(h) = \sum_{n \in \mathbb{N}} h_n X_n$ for $h \in \ell^f$ is well-defined (since it is a finite sum). Since

$$\mathbf{E}[R(h)^2] = \text{Var}[R(h)] = \sum_{n \in \mathbb{N}} \text{Var}[h_n X_n] = \sum_{n \in \mathbb{N}} h_n^2 = \|h\|^2,$$

the map $R : \ell^f \rightarrow \mathcal{L}^2(\mathbf{P})$ is an isometry. As $\ell^f \subset \ell^2$ is dense, there is a unique continuous extension of R to ℓ^2 . Hence, if $h \in \ell^2$ and $(h^N)_{N \in \mathbb{N}}$ is a sequence in ℓ^f with $\|h^N - h\| \xrightarrow{N \rightarrow \infty} 0$, then $R(h^N) \xrightarrow{N \rightarrow \infty} R(h)$ in the L^2 sense. In particular, $h_n^N := h_n \mathbb{1}_{\{n \leq N\}}$, $n \in \mathbb{N}$, $N \in \mathbb{N}$, is an approximating sequence for h , and we have $R(h^N) = \sum_{n=1}^N h_n X_n$. Thus the approximation of R with the partial sums R_N that we described above is a special case of this construction.

The programme for the construction of the Itô integral $I_t^W(H)$ is the following. First consider simple functions as integrands H ; that is, the map $t \mapsto H_t(\omega)$ is a step function. For these H , the integral can easily be defined as a finite sum. The next step is to extend the integral, as in Example 25.1, to integrands that can be approximated in a certain L^2 -space by simple integrands.

Definition 25.2 Denote by $\mathcal{E}$ the vector space of maps $H : \Omega \times [0, \infty) \rightarrow \mathbb{R}$ of the form

$$H_t(\omega) = \sum_{i=1}^n h_{i-1}(\omega) \mathbb{1}_{(t_{i-1}, t_i]}$$

where $n \in \mathbb{N}$, $0 = t_0 < t_1 < \dots < t_n$ and h_{i-1} is bounded and $\mathcal{F}_{t_{i-1}}$ -measurable for every $i = 1, \dots, n$. $\mathcal{E}$ is called the vector space of predictable simple processes.

We equip $\mathcal{E}$ with a (pseudo) norm $\|\cdot\|_{\mathcal{E}}$ by defining

$$\|H\|_{\mathcal{E}}^2 = \sum_{i=1}^n \mathbf{E}[h_{i-1}^2] (t_i - t_{i-1}) = \mathbf{E} \left[\int_0^{\infty} H_s^2 ds \right]$$

Definition 25.3 For $H \in \mathcal{E}$ and $t \geq 0$, define

$$I_t^W(H) = \sum_{i=1}^n h_{i-1} (W_{t_i \wedge t} - W_{t_{i-1} \wedge t})$$

and

$$I_{\infty}^W(H) = \sum_{i=1}^n h_{i-1} (W_{t_i} - W_{t_{i-1}})$$

Clearly, for every bounded stopping time τ ,

$$\begin{aligned} \mathbf{E}[I_{\tau}^W(H)] &= \sum_{i=1}^n \mathbf{E} \left[h_{i-1} (W_{t_i}^{\tau} - W_{t_{i-1}}^{\tau}) \right] \\ &= \sum_{i=1}^n \mathbf{E} \left[h_{i-1} \mathbf{E} [W_{t_i}^{\tau} - W_{t_{i-1}}^{\tau} \mid \mathcal{F}_{t_{i-1}}] \right] = 0 \end{aligned}$$

since, by the optional stopping theorem (OST), the stopped Brownian motion W^τ is an $\mathbb{F}$ -martingale. Hence (again by the OST) $(I_t^W(H))_{t \geq 0}$ is an $\mathbb{F}$ -martingale. In particular, we have $\mathbf{E} \left[\left(I_{t_{i+1}}^W(H) - I_{t_i}^W(H) \right) \left(I_{t_{j+1}}^W(H) - I_{t_j}^W(H) \right) \right] = 0$ for $i \neq j$. Therefore,

$$\begin{aligned} \mathbf{E} [I_\infty^W(H)^2] &= \sum_{i=1}^n \mathbf{E} \left[\left(I_{t_i}^W(H) - I_{t_{i-1}}^W(H) \right)^2 \right] \\ &= \sum_{i=1}^n \mathbf{E} \left[h_{i-1}^2 (W_{t_i} - W_{t_{i-1}})^2 \right] \\ &= \sum_{i=1}^n \mathbf{E} [h_{i-1}^2] (t_i - t_{i-1}) = \|H\|_{\mathcal{E}}^2 \end{aligned} \quad (25.2)$$

From these considerations, the following statement is immediate.

Theorem 25.4

- (i) The map $I_\infty^W : \mathcal{E} \rightarrow \mathcal{L}^2(\Omega, \mathcal{F}, \mathbf{P})$ is an isometric linear map (with respect to $\|\cdot\|_{\mathcal{E}}$ and $\|\cdot\|_2$).
- (ii) The process $(I_t^W(H))_{t \geq 0}$ is an L^2 -bounded continuous $\mathbb{F}$ -martingale.

Proof Only the linearity remains to be shown. However, this is trivial.

The idea is to extend the map I_∞^W continuously from $\mathcal{E}$ to a suitable closure $\bar{\mathcal{E}}$ of $\mathcal{E}$. Now as a subspace of what space should we close $\mathcal{E}$? A minimal requirement is that $(\omega, t) \mapsto H_t(\omega)$ be measurable (with respect to $\mathcal{F} \otimes \mathcal{B}([0, \infty))$) and that H be adapted.

Definition 25.5 A stochastic process $X = (X_t)_{t \geq 0}$ with values in a Polish space E is called

- (i) product measurable if $(\omega, t) \mapsto X_t(\omega)$ is measurable with respect to $\mathcal{F} \otimes \mathcal{B}([0, \infty)) - \mathcal{B}(E)$,
- (ii) progressively measurable if, for every $t \geq 0$, the map $\Omega \times [0, t] \rightarrow E, (\omega, s) \mapsto X_s(\omega)$ is measurable with respect to $\mathcal{F}_t \otimes \mathcal{B}([0, t]) - \mathcal{B}(E)$,
- (iii) predictable (or previsible) if $(\omega, t) \mapsto X_t(\omega)$ is measurable with respect to the predictable σ -algebra $\mathcal{P}$ on $\Omega \times [0, \infty)$:

$$\mathcal{P} := \sigma(X : X \text{ is a left continuous adapted process}).$$

Remark 25.6 Any $H \in \mathcal{E}$ is predictable. This property ensures that $I^M(H)$ is a martingale for every (even discontinuous) martingale M . The notion of predictability is important only for integration with respect to discontinuous martingales. As we will not develop that calculus in this book, predictability will not be central for us.

Remark 25.7 If H is progressively measurable, then H is evidently also product measurable and adapted. With a little work, the converse can also be shown: If H is adapted and product measurable, then there is a progressively measurable modification of H (see, e.g., [115, pp. 68ff]).

Theorem 25.8 If H is adapted and right continuous or left continuous, then H is progressively measurable. If H is adapted and a.s. right continuous or left continuous, then there exists a version of H that is progressively measurable.

In particular, every predictable process is progressively measurable.

Proof See Exercise 21.1.4.

We consider $\mathcal{E}$ as a subspace of

$$\mathcal{E}_0 := \left\{ H : \text{product measurable, adapted and } \|H\|^2 := \mathbf{E} \left[\int_0^\infty H_t^2 dt \right] < \infty \right\}$$

Let $\bar{\mathcal{E}}$ denote the closure of $\mathcal{E}$ in $\mathcal{E}_0$.

Theorem 25.9 If H is progressively measurable (for instance, left continuous or right continuous and adapted) and $\mathbf{E} \left[\int_0^\infty H_t^2 dt \right] < \infty$, then $H \in \bar{\mathcal{E}}$.

Proof Let H be progressively measurable and $\mathbf{E} \left[\int_0^\infty H_t^2 dt \right] < \infty$. It is enough to show that, for any $T > 0$, there exists a sequence $(H^n)_{n \in \mathbb{N}}$ in $\mathcal{E}$ such that

$$\mathbf{E} \left[\int_0^T (H_s - H_s^n)^2 ds \right] \xrightarrow{n \rightarrow \infty} 0 \quad (25.3)$$

Step 1. First assume that H is continuous and bounded. Define $H_0^n = 0$ and

$$H_t^n = H_{i2^{-n}T} \quad \text{if } i2^{-n}T < t \leq (i+1)2^{-n}T \text{ for some } i = 0, \dots, 2^n - 1$$

and $H_t^n = 0$ for $t > T$. Then $H^n \in \mathcal{E}$, and we have $H_t^n(\omega) \xrightarrow{n \rightarrow \infty} H_t(\omega)$ for all $t > 0$ and $\omega \in \Omega$. By the dominated convergence theorem, we get (25.3).

Step 2. Now let H be progressively measurable and bounded. It is enough to show that there exist continuous adapted processes $H^n, n \in \mathbb{N}$, for which (25.3) holds. Let

$$H_t^n := n \int_{(t-1/n) \vee 0}^{t \wedge T} H_s ds \quad \text{for } t \geq 0, n \in \mathbb{N}$$

Then H^n is continuous, adapted and bounded by $\|H\|_\infty$. By the fundamental theorem of calculus (see Exercise 13.1.7), we have

$$H_t^n(\omega) \xrightarrow{n \rightarrow \infty} H_t(\omega) \quad \text{for } \lambda\text{-almost all } t \in [0, T] \text{ and for all } \omega \in \Omega. \quad (25.4)$$

By Fubini's theorem and the dominated convergence theorem, we thus conclude that

$$\mathbf{E} \left[\int_0^T (H_s - H_s^n)^2 ds \right] = \int_{\Omega \times [0, T]} (H_s(\omega) - H_s^n(\omega))^2 (\mathbf{P} \otimes \lambda)(d(\omega, s)) \xrightarrow{n \rightarrow \infty} 0$$

Step 3. Now let H be progressively measurable, and assume $\mathbf{E} \left[\int_0^\infty H_t^2 dt \right] < \infty$. It is enough to show that there exists a sequence $(H^n)_{n \in \mathbb{N}}$ of bounded, progressively measurable processes such that (25.3) holds. Manifestly, we can choose $H_t^n = H_t \mathbb{1}_{\{|H_t| < n\}}$.

Definition 25.10 (Itô integral) For $H \in \bar{\mathcal{E}}$, define the Itô integral

$$\int_0^\infty H_s dW_s := I_\infty^W(H)$$

as the continuous extension of the map $I_\infty^W : \mathcal{E} \rightarrow \mathcal{L}^2(\mathbf{P})$ to the closure $\bar{\mathcal{E}}$ of $\mathcal{E}$. In other words, if $(H^n)_{n \in \mathbb{N}}$ is a sequence in $\mathcal{E}$ with $\|H - H^n\| \xrightarrow{n \rightarrow \infty} 0$, then we define $I_\infty^W(H)$ by

$$I_\infty^W(H) := \lim_{n \rightarrow \infty} I_\infty^W(H^n) \quad \text{in } L^2$$

If τ is a stopping time, then in the following we use the abbreviation

$$H_t^{(\tau)} := H_t \mathbb{1}_{\{t \leq \tau\}} \quad \text{for } t \geq 0.$$

(Note that this is not the stopped process $H_t^\tau = H_{\tau \wedge t}$.)

Theorem 25.11

(i) The map $I_\infty^W : \bar{\mathcal{E}} \rightarrow \mathcal{L}^2(\Omega, \mathcal{F}, \mathbf{P})$ is linear and

$$\mathbf{E} [I_\infty^W(H)^2] = \mathbf{E} \left[\int_0^\infty H_s^2 ds \right]$$

(ii) For every $H \in \bar{\mathcal{E}}$, the process $\tilde{I}^W(H)$ defined by $\tilde{I}_t^W(H) := I_\infty^W(H^{(t)})$ is an L^2 -bounded $\mathbb{F}$ -martingale that has a continuous modification $I^W(H)$.

Definition 25.12 (Itô integral as a process) Let $I^W(H)$ be the continuous version of the martingale $(I_\infty^W(H^{(t)}))_{t \geq 0}$ (see Theorem 25.11(ii)). Denote by

$$\int_s^t H_r dW_r := I_t^W(H) - I_s^W(H) \quad \text{for } 0 \leq s \leq t \leq \infty$$

the Itô integral of H with respect to Brownian motion W on the interval $[s, t]$.

Proof of Theorem 25.11 (i) This is a direct consequence of the definition of $I_\infty^W(H)$.

(ii) Let $(H^n)_{n \in \mathbb{N}}$ be a sequence in $\mathcal{E}$ with $\|H^n - H\| \xrightarrow{n \rightarrow \infty} 0$. By Theorem 25.4(ii), we have

$$I_\infty^W((H^n)^{(t)}) = I_t^W(H^n) = \mathbf{E} [I_\infty^W(H^n) | \mathcal{F}_t] \quad \text{for all } t \geq 0, n \in \mathbb{N}.$$

Since $\|(H^n)^{(t)} - H^{(t)}\| \leq \|H^n - H\| \xrightarrow{n \rightarrow \infty} 0$, this implies (using Corollary 8.21)

$$\tilde{I}_t^W(H) = \lim_{n \rightarrow \infty} I_t^W(H^n) = \lim_{n \rightarrow \infty} \mathbf{E} [I_\infty^W(H^n) | \mathcal{F}_t] = \mathbf{E} [I_\infty^W(H) | \mathcal{F}_t]$$

Hence $\tilde{I}^W(H)$ is an L^2 -bounded martingale and $I_t^W(H^n) \xrightarrow{n \rightarrow \infty} \tilde{I}_t^W(H)$ in L^2 for every $t \geq 0$. By Theorem 25.4(ii), $I^W(H^n)$ is continuous for every $n \in \mathbb{N}$. Thus, by Exercise 21.4.3, there exists a continuous modification $I^W(H)$ of $\tilde{I}^W(H)$.

The last step in the construction of the Itô integral is to weaken the strong integrability condition $\mathbf{E} [\int_0^\infty H_s^2 ds] < \infty$. We start with a simple observation.

Let τ be a stopping time and recall that $\int_0^\tau H_s dW_s$ denotes the random variable that for any ω assumes the value $(\int_0^{\tau(\omega)} H_s dW_s)(\omega)$.

Lemma 25.13 Let τ be a stopping time and let $H \in \bar{\mathcal{E}}$.

(i) We have

$$\int_0^\tau H_s dW_s = \int_0^\infty H_s^{(\tau)} dW_s := \int_0^\infty H_s \mathbb{1}_{\{s \leq \tau\}} dW_s \quad \text{a.s.}$$

(ii) In particular, for any $t \geq 0$, on the event $\{\tau \geq t\}$ we have

$$\int_0^t H_s dW_s = \int_0^t H_s^{(\tau)} dW_s \quad \text{a.s.}$$

(iii) Let $G \in \bar{\mathcal{E}}$ be such that $H_s = G_s$ for all $s \leq \tau$. Then

$$\int_0^\tau H_s dW_s = \int_0^\tau G_s dW_s \quad \text{a.s.}$$

Proof (i) Assume first that τ takes values in $\{k/2^n : k \in \mathbb{N}_0\} \cup \{\infty\}$ for some $n \in \mathbb{N}$. Then $\mathbb{1}_{\{k/2^n \leq \tau\}} \mathbb{1}_{\{t \in ((k-1)/2^n, k/2^n]\}} \in \mathcal{E}$ for all $k \in \mathbb{N}$. If, in addition, $H \in \mathcal{E}$, then also $H^{(\tau)} \in \mathcal{E}$ and the claim follows directly from the definition of the Itô integral (Definition 25.3). Now let $H \in \bar{\mathcal{E}}$ and let $(H^k)_{k \in \mathbb{N}}$ be a sequence in $\mathcal{E}$ such that $\|H^k - H\|_{\mathcal{E}} \xrightarrow{k \rightarrow \infty} 0$. Writing $H_t^{k,(\tau)} := H_t^k \mathbb{1}_{\{t \leq \tau\}}$ we get that $\|H^{k,(\tau)} - H^{(\tau)}\|_{\mathcal{E}} \xrightarrow{k \rightarrow \infty} 0$. By choosing a suitable sequence $k_m \uparrow \infty$, we obtain

$$\begin{aligned}\int_0^\tau H_s dW_s &= \lim_{m \rightarrow \infty} \int_0^\tau H_s^{k_m} dW_s \\ &= \lim_{m \rightarrow \infty} \int_0^\infty H_s^{k_m, (\tau)} dW_s = \int_0^\infty H_s^{(\tau)} dW_s \quad \text{a.s.}\end{aligned}$$

Finally, assume that τ is an arbitrary stopping time and define $\tau_n := 2^{-n} \lceil 2^n \tau \rceil$ for $n \in \mathbb{N}$. Then (τ_n) is a sequence of stopping times with $\tau_n \downarrow \tau$. Recall that $I^W(H)$ is continuous and note that $\|H^{(\tau_n)} - H^{(\tau)}\|_{\mathcal{E}} \xrightarrow{n \rightarrow \infty} 0$. Hence by taking a suitable sequence $n(m) \uparrow \infty$, we get

$$\begin{aligned}\int_0^\tau H_s dW_s &= \lim_{m \rightarrow \infty} \int_0^{\tau_{n(m)}} H_s dW_s \\ &= \lim_{m \rightarrow \infty} \int_0^\infty H_s^{(\tau_{n(m)})} dW_s = \int_0^\infty H_s^{(\tau)} dW_s \quad \text{a.s.}\end{aligned}$$

(ii), (iii) These statements are direct consequences of (i).

Definition 25.14 Let $\mathcal{E}_{\text{loc}}$ be the space of progressively measurable stochastic processes H with

$$\int_0^T H_s^2 ds < \infty \quad \text{a.s. for all } T > 0$$

Lemma 25.15 For every $H \in \mathcal{E}_{\text{loc}}$, there exists a sequence $(\tau_n)_{n \in \mathbb{N}}$ of stopping times with $\tau_n \uparrow \infty$ almost surely and $\mathbf{E} \left[\int_0^{\tau_n} H_s^2 ds \right] < \infty$ and hence such that $H^{(\tau_n)} \in \bar{\mathcal{E}}$ for every $n \in \mathbb{N}$.

Proof Define

$$\tau_n := \inf \left\{ t \geq 0 : \int_0^t H_s^2 ds \geq n \right\}$$

By the definition of $\mathcal{E}_{\text{loc}}$, we have $\tau_n \uparrow \infty$ almost surely. By construction, we have $\|H^{(\tau_n)}\|^2 = \mathbf{E} \left[\int_0^{\tau_n} H_s^2 ds \right] \leq n$.

Definition 25.16 Let $H \in \mathcal{E}_{\text{loc}}$ and let $(\tau_n)_{n \in \mathbb{N}}$ be as in Lemma 25.15. For $t \geq 0$, define the Itô integral as the almost sure limit

$$\int_0^t H_s dW_s := \lim_{n \rightarrow \infty} \int_0^t H_s^{(\tau_n)} dW_s \quad (25.5)$$

Theorem 25.17 Let $H \in \mathcal{E}_{\text{loc}}$.

- (i) The limit in (25.5) is well-defined and continuous at t . Up to a.s. equality, it is independent of the choice of the sequence $(\tau_n)_{n \in \mathbb{N}}$.
- (ii) If τ is a stopping time with $\mathbf{E} \left[\int_0^\tau H_s^2 ds \right] < \infty$, then the stopped Itô integral $\left(\int_0^{\tau \wedge t} H_s dW_s \right)_{t \geq 0}$ is an L^2 -bounded, continuous martingale.
- (iii) If $\mathbf{E} \left[\int_0^T H_s^2 ds \right] < \infty$ for all $T > 0$, then $\left(\int_0^t H_s dW_s \right)_{t \geq 0}$ is a square integrable continuous martingale.

Proof (i) By Lemma 25.13(ii), on the event $\{\tau_n \geq t\}$, we have

$$\int_0^t H_s dW_s = \int_0^t H_s^{(\tau_n)} dW_s$$

Hence the limit exists, is continuous and is independent of the choice of the sequence $(\tau_n)_{n \in \mathbb{N}}$.

(ii) This is immediate by Theorem 25.11.

(iii) As we can choose $\tau_n = n$, this follows from (ii).

Theorem 25.18 Let H be progressively measurable and $\mathbf{E} \left[\int_0^T H_s^2 ds \right] < \infty$ for all $T > 0$. Then

$$M_t := \int_0^t H_s dW_s, \quad t \geq 0$$

defines a square integrable continuous martingale, and

$$(N_t)_{t \geq 0} := \left(M_t^2 - \int_0^t H_s^2 ds \right)_{t \geq 0}$$

is a continuous martingale with $N_0 = 0$.

Proof It is enough to show that N is a martingale. Clearly, N is adapted. Let τ be a bounded stopping time. Then

$$\begin{aligned} \mathbf{E}[N_\tau] &= \mathbf{E} \left[M_\tau^2 - \int_0^\tau H_s^2 ds \right] \\ &= \mathbf{E} \left[\left(\int_0^\tau H_s^{(\tau)} dW_s \right)^2 \right] - \mathbf{E} \left[\int_0^\tau (H_s^{(\tau)})^2 ds \right] = 0 \end{aligned}$$

Thus, by the optional stopping theorem (see Exercise 21.1.3(iii)), N is a martingale.

Recall the notions of local martingales and square variation from Section 21.10.

Corollary 25.19 If $H \in \mathcal{E}_{\text{loc}}$, then the Itô integral $M_t = \int_0^t H_s dW_s$ is a continuous local martingale with square variation process $\langle M \rangle_t = \int_0^t H_s^2 ds$.

Example 25.20

(i) $W_t = \int_0^t 1 dW_s$ is a square integrable martingale, and $(W_t^2 - t)_{t \geq 0}$ is a continuous martingale.

(ii) Since $\mathbf{E} \left[\int_0^T W_s^2 ds \right] = \frac{T^2}{2} < \infty$ for all $T \geq 0$, $M_t := \int_0^t W_s dW_s$ is a continuous, square integrable martingale, and $\left(M_t^2 - \int_0^t W_s^2 ds \right)_{t \geq 0}$ is a continuous martingale.

(iii) Assume that H is progressively measurable and bounded, and let $M_t := \int_0^t H_s dW_s$. Then M is progressively measurable (since it is continuous and adapted) and

$$\mathbf{E} \left[\int_0^T M_s^2 ds \right] = \int_0^T \left(\int_0^s \mathbf{E} [H_r^2] dr \right)^2 ds \leq \frac{T^2 \|H\|_\infty^2}{2}$$

Hence $\tilde{M}_t := \int_0^t M_s dW_s$ is a square integrable, continuous martingale and $\left(\tilde{M}_t^2 - \int_0^t M_s^2 dW_s \right)_{t \geq 0}$ is a continuous martingale.

5.25.2 25.2 Itô Integral with Respect to Diffusions

If

$$H = \sum_{i=1}^n h_{i-1} \mathbb{1}_{(t_{i-1}, t_i]} \in \mathcal{E} \quad (25.6)$$

then the elementary integral

$$I_t^M(H) = \sum_{i=1}^n h_{i-1} (M_{t_i \wedge t} - M_{t_{i-1} \wedge t})$$

is a martingale (respectively local martingale) if M is a martingale (respectively local martingale). Furthermore,

$$\begin{aligned}\mathbf{E} \left[(I_{\infty}^M(H))^2 \right] &= \sum_{i=1}^n \mathbf{E} \left[h_{i-1}^2 (M_{t_i} - M_{t_{i-1}})^2 \right] \\ &= \sum_{i=1}^n \mathbf{E} \left[h_{i-1}^2 (\langle M \rangle_{t_i} - \langle M \rangle_{t_{i-1}}) \right] \\ &= \mathbf{E} \left[\int_0^{\infty} H_t^2 d\langle M \rangle_t \right]\end{aligned}$$

if the expression on the right-hand side is finite. Roughly speaking, the procedure in Section 25.1 by which we defined the Itô integral for Brownian motion and integrands $H \in \bar{\mathcal{E}}$ can be repeated to construct a stochastic integral with respect to M for a large class of integrands H . Essentially, in the definition of the norm on $\mathcal{E}$ we have to replace dt (that is, the square variation of Brownian motion) by the square variation $d\langle M \rangle_t$ of M :

$$\|H\|_M^2 := \mathbf{E} \left[\int_0^{\infty} H_t^2 d\langle M \rangle_t \right]$$

Extending the integral to the closure $\bar{\mathcal{E}}$ works just as for Brownian motion. The tricky point is to check whether a given integrand is in $\bar{\mathcal{E}}$. For example, for discontinuous martingales M the integrands have to be predictable in order for the stochastic integral to be a martingale (not to mention the difficulty of establishing for such M , the existence of the square variation process). For the case of discrete time processes, we saw this in Section 9.3. Now if M is a continuous martingale with continuous square variation $\langle M \rangle$, then the following problem occurs. In the proof of Theorem 25.9 in Step 2, in order to show that progressively measurable processes H are in $\bar{\mathcal{E}}$, we used the fact that $H_t^n(\omega) \xrightarrow{n \rightarrow \infty} H_t(\omega)$ for Lebesgue-almost all t and all ω . Now if $d\langle M \rangle_t$ is not absolutely continuous with respect to the Lebesgue measure, then this is not sufficient to infer convergence of the integrals with respect to $d\langle M \rangle_t$. In the case of absolutely continuous square variation, however, that proof works without change. As in Section 25.1, we obtain the following theorem.

Theorem 25.21 Let M be a continuous local martingale with absolutely continuous square variation $\langle M \rangle$ and let H be a progressively measurable process with $\int_0^T H_s^2 d\langle M \rangle_s < \infty$ a.s. for all $T \geq 0$. Then the Itô integral $N_t := \int_0^t H_s dM_s$ is well-defined and is a continuous local martingale with square variation $\langle N \rangle_t = \int_0^t H_s^2 d\langle M \rangle_s$. For any sequence $(\tau_n)_{n \in \mathbb{N}}$ with $\tau_n \uparrow \infty$ and $\|H^{(\tau_n)}\|_M < \infty$, and for any family $(H^{n,m}, n, m \in \mathbb{N}) \subset \mathcal{E}$ with $\|H^{n,m} - H^{(\tau_n)}\|_M \xrightarrow{m \rightarrow \infty} 0$, we have

$$\int_0^t H_s dM_s = \lim_{n \rightarrow \infty} \lim_{m \rightarrow \infty} I_t^M(H^{m,n}) \quad \text{in probability for all } t \geq 0.$$

The following theorem formulates a certain generalization.

Theorem 25.22 Let M^1 and M^2 be continuous local martingales with absolutely continuous square variation. Let H^i be progressively measurable processes with $\int_0^T (H_s^i)^2 d\langle M^i \rangle_s < \infty$ for $i = 1, 2$ and $T < \infty$. Let $N_t^i := \int_0^t H_s^i dM_s^i$ for $i = 1, 2$. Then N^1 and N^2 are continuous local martingales with quadratic covariation $\langle N^i, N^j \rangle_t = \int_0^t H_s^i H_s^j d\langle M^i, M^j \rangle_s$ for $i, j \in \{1, 2\}$. If M^1 and M^2 are independent, then $\langle N^1, N^2 \rangle \equiv 0$.

Proof First assume $H^1, H^2 \in \mathcal{E}$. Then there are numbers $0 = t_0 < t_1 < \dots < t_n$ and $\mathcal{F}_{t_k}$ -measurable bounded maps $h_k^i, i = 1, 2, k = 0, \dots, n-1$ such that

$$H_t^i(\omega) = \sum_{k=1}^n h_{k-1}^i(\omega) \mathbb{1}_{(t_{k-1}, t_k]}(t)$$

Therefore,

$$N_t^i N_t^j = \sum_{k,l=1}^n h_{k-1}^i h_{l-1}^j \left(M_{t_k \wedge t}^i - M_{t_{k-1} \wedge t}^i \right) \left(M_{t_l \wedge t}^j - M_{t_{l-1} \wedge t}^j \right).$$

Those summands with $k \neq l$ are local martingales. For any of the summands with $k = l$,

$$\begin{aligned} & \left(h_{k-1}^i h_{k-1}^j \left(\left(M_{t_k \wedge t}^i - M_{t_{k-1} \wedge t}^i \right) \left(M_{t_k \wedge t}^j - M_{t_{k-1} \wedge t}^j \right) \right. \right. \\ & \quad \left. \left. - \left(\langle M^i, M^j \rangle_{t_k \wedge t} - \langle M^i, M^j \rangle_{t_{k-1} \wedge t} \right) \right) \right)_{t \geq 0} \end{aligned}$$

is a local martingale. Since

$$\sum_{k=1}^n h_{k-1}^i h_{k-1}^j \left(\langle M^i, M^j \rangle_{t_k \wedge t} - \langle M^i, M^j \rangle_{t_{k-1} \wedge t} \right) = \int_0^t H_s^i H_s^j d \langle M^i, M^j \rangle_s$$

$\left(N_t^i N_t^j - \int_0^t H_s^i H_s^j d \langle M^i, M^j \rangle_s \right)_{t \geq 0}$ is a continuous local martingale.

The case of general progressively measurable H^1, H^2 that satisfy an integrability condition follows by the usual L^2 -approximation arguments.

If M^1 and M^2 are independent, then $\langle M^1, M^2 \rangle \equiv 0$.

In the following, we consider processes that can be expressed as Itô integrals with respect to a Brownian motion. For these processes, we give a different and more detailed proof of Theorem 25.21.

Definition 25.23 Let W be a Brownian motion and let σ and b be progressively measurable stochastic processes with $\int_0^t \sigma_s^2 + |b_s| ds < \infty$ almost surely for all $t \geq 0$. Then we say that the process X defined by

$$X_t = \int_0^t \sigma_s dW_s + \int_0^t b_s ds \quad \text{for } t \geq 0$$

is a generalized diffusion process (or, briefly, generalized diffusion) with diffusion coefficient σ and drift b . Often X is called an Itô process.

In particular, if σ and b are of the form $\sigma_s = \tilde{\sigma}(X_s)$ and $b_s = \tilde{b}(X_s)$ for certain maps $\tilde{\sigma} : \mathbb{R} \rightarrow [0, \infty)$ and $\tilde{b} : \mathbb{R} \rightarrow \mathbb{R}$, then X is called a diffusion (in the proper sense).

In contrast with generalized diffusions, we will see that under certain regularity assumptions on the coefficients, diffusions in the proper sense are Markov processes (compare Theorems 26.8, 26.10 and 26.26).

A diffusion X can always be decomposed as $X = M + A$, where $M_t = \int_0^t \sigma_s dW_s$ is a continuous local martingale with square variation $\langle M \rangle_t = \int_0^t \sigma_s^2 ds$ (by Corollary 25.19) and $A_t = \int_0^t b_s ds$ is a continuous process of locally finite variation.

Clearly, for the H in (25.6), we have

$$\begin{aligned} \int_0^t H_s dM_s &= \sum_{i=1}^n h_{i-1} \left(M_{t_i \wedge t} - M_{t_{i-1} \wedge t} \right) \\ &= \sum_{i=1}^n h_{i-1} \int_{t_{i-1} \wedge t}^{t_i \wedge t} \sigma_s dW_s = \int_0^t (H_s \sigma_s) dW_s \end{aligned}$$

For progressively measurable H with $\int_0^T H_s^2 d\langle M \rangle_s = \int_0^T (H_s \sigma_s)^2 ds < \infty$ for all $T \geq 0$, we thus define the Itô integral as

$$\int_0^t H_s dM_s := \int_0^t (H_s \sigma_s) dW_s$$

Without further work, in particular, without relying on Theorem 25.21, we get the following theorem.

Theorem 25.24 Let $X = M + A$ be a generalized diffusion with σ and let b be as in Definition 25.23. Let H be progressively measurable with

$$\int_0^T H_s^2 \sigma_s^2 ds < \infty \quad \text{a.s. for all } T \geq 0 \quad (25.7)$$

and

$$\int_0^T |H_s b_s| ds < \infty \quad \text{a.s. for all } T \geq 0 \quad (25.8)$$

Then the process Y defined by

$$Y_t := \int_0^t H_s dX_s := \int_0^t H_s dM_s + \int_0^t H_s dA_s := \int_0^t H_s \sigma_s dW_s + \int_0^t H_s b_s ds$$

is a generalized diffusion with diffusion coefficient $(H_s \sigma_s)_{s \geq 0}$ and drift $(H_s b_s)_{s \geq 0}$. In particular, $N_t := \int_0^t H_s dM_s$ is a continuous local martingale with square variation process $\langle N \rangle_t = \int_0^t H_s^2 d\langle M \rangle_s = \int_0^t H_s^2 \sigma_s^2 ds$.

Exercise 25.2.1 Let M be a continuous local martingale with absolutely continuous square variation $\langle M \rangle$ (e.g., a generalized diffusion), and let H be progressively measurable and continuous with $\int_0^T H_s^2 d\langle M \rangle_s < \infty$ for all $T \geq 0$. Further, assume that $\mathcal{P} = (\mathcal{P}^{(n)})_{n \in \mathbb{N}}$ is an admissible sequence of partitions (see Definition 21.56).

(i) Show that for all $T \geq 0$, in the sense of stochastic convergence, we have

$$\int_0^T H_s dM_s = \lim_{n \rightarrow \infty} \sum_{t \in \mathcal{P}_T^n} H_t (M_{t'} - M_t) \quad (25.9)$$

(ii) Show that there exists a subsequence of $\mathcal{P}$ such that almost surely, we have (25.9) for all $T \geq 0$.

5.25.3 25.3 The Itô Formula

This and the following two sections are based on lecture notes of Hans Föllmer.

If $t \mapsto X_t$ is a differentiable map with derivative X' and $F \in C^1(\mathbb{R})$ with derivative F' , then we have the classical substitution rule

$$F(X_t) - F(X_0) = \int_0^t F'(X_s) dX_s = \int_0^t F'(X_s) X'_s ds \quad (25.10)$$

This remains true even if X is continuous and has locally finite variation (see Section 21.10); that is, if X is the distribution function of an absolutely continuous signed measure on $[0, \infty)$. In this case, the derivative X' exists as a Radon-Nikodym derivative almost everywhere, and it is easy to show that (25.10) also holds in this case.

The paths of Brownian motion W are nowhere differentiable (Theorem 21.17 due to Paley, Wiener and Zygmund) and thus have everywhere locally infinite variation. Hence a substitution

rule as simple as (25.10) cannot be expected. Indeed, it is easy to see that such a rule must be false: Choose $F(x) = x^2$. Then the right-hand side in (25.10) (with X replaced by W) is $\int_0^t 2W_s dW_s$ and is hence a martingale. The left-hand side, however, equals W_t^2 , which is a submartingale that only becomes a martingale by subtracting t . Indeed, this t is the additional term that shows up in the substitution rule for Itô integrals, the so-called Itô formula. A somewhat bold

heuristic puts us on the right track: For small t , W_t is of order $\sqrt{t}$. If we formally write $dW_t = \sqrt{dt}$ and carry out a Taylor expansion of $F \in C^2(\mathbb{R})$ up to second order, then we obtain

$$dF(W_t) = F'(W_t) dW_t + \frac{1}{2} F''(W_t) (dW_t)^2 = F'(W_t) dW_t + \frac{1}{2} F''(W_t) dt$$

Rewriting this as an integral yields

$$F(W_t) - F(W_0) = \int_0^t F'(W_s) dW_s + \int_0^t \frac{1}{2} F''(W_s) ds \quad (25.11)$$

(For certain discrete martingales, we derived a similar formula in Example 10.9.) The main goal of this section is to show that this so-called Itô formula is indeed correct.

The subsequent discussion in this section does not explicitly rely on the assumption that we integrate with respect to Brownian motion. All that is needed is that the function with respect to which we integrate have continuous square variation (along a suitable admissible sequence of partitions $\mathcal{P} = (\mathcal{P}^n)_{n \in \mathbb{N}}$). In particular, for Brownian motion, $\langle W \rangle_t = t$.

In the following, let $\mathcal{P} = (\mathcal{P}^n)_{n \in \mathbb{N}}$ be an admissible sequence of partitions (recall the definition of $\mathcal{C}_{\text{qv}} = \mathcal{C}_{\text{qv}}^{\mathcal{P}}, \mathcal{P}_T^n, \mathcal{P}_{S,T}^n, t'$ and so on from Definitions 21.56 and 21.58). Let $X \in C([0, \infty))$ with continuous square variation (along $\mathcal{P}$)

$$T \mapsto \langle X \rangle_T = V_T^2(X) = \lim_{n \rightarrow \infty} \sum_{t \in \mathcal{P}_T^n} (X_{t'} - X_t)^2$$

For Brownian motion, we have $W \in \mathcal{C}_{\text{qv}}^{\mathcal{P}}$ almost surely for any admissible sequence of partitions (Theorem 21.64) and $\langle W \rangle_T = T$. For continuous local martingales M passing to a suitable subsequence $\mathcal{P}'$ of $\mathcal{P}$ ensures that $M \in \mathcal{C}_{\text{qv}}^{\mathcal{P}'}$ almost surely (Theorem 21.70).

Now fix $\mathcal{P}$ and let $X \in \mathcal{C}_{\text{qv}}$ be a (deterministic) function.

Theorem 25.25 (Pathwise Itô formula) Let $X \in \mathcal{C}_{\text{qv}}$ and $F \in C^2(\mathbb{R})$. Then, for all $T \geq 0$, there exists the limit

$$\int_0^T F'(X_s) dX_s := \lim_{n \rightarrow \infty} \sum_{t \in \mathcal{P}_T^n} F'(X_t) (X_{t'} - X_t) \quad (25.12)$$

Furthermore, the Itô formula holds:

$$F(X_T) - F(X_0) = \int_0^T F'(X_s) dX_s + \frac{1}{2} \int_0^T F''(X_s) d\langle X \rangle_s \quad (25.13)$$

Here the right integral in (25.13) is understood as a classical (Lebesgue-Stieltjes) integral.

Remark 25.26 If M is a continuous local martingale, then, by Exercise 25.2.1, the Itô integral $\int_0^T F'(M_s) dM_s$ is the stochastic limit of $\sum_{t \in \mathcal{P}_T^n} F'(M_t) (M_{t'} - M_t)$ as $n \rightarrow \infty$. Thus, in fact, for $X = M(\omega)$, the pathwise integral in (25.12) coincides with the Itô integral (a.s.). In particular, for the Itô integral of Brownian motion, the Itô formula (25.11) holds. $\square$

Proof of Theorem 25.25 We have to show that the limit in (25.12) exists and that (25.13) holds.

For $n \in \mathbb{N}$ and $t \in \mathcal{P}_T^n$ (with successor $t' \in \mathcal{P}_T^n$), the Taylor formula yields

$$F(X_{t'}) - F(X_t) = F'(X_t)(X_{t'} - X_t) + \frac{1}{2}F''(X_t) \cdot (X_{t'} - X_t)^2 + R_t^n, \quad (25.14)$$

where the remainder

$$R_t^n = (F''(\xi) - F''(X_t)) \cdot \frac{1}{2}(X_{t'} - X_t)^2$$

(for a suitable ξ between X_t and $X_{t'}$) can be bounded as follows. As X is continuous, $C := \{X_t : t \in [0, T]\}$ is compact and $F''|_C$ is uniformly continuous. Thus, for every $\varepsilon > 0$, there exists a $\delta > 0$ with

$$|F''(X_r) - F''(X_s)| < \varepsilon \quad \text{for all } r, s \in [0, T] \text{ with } |X_r - X_s| < \delta.$$

Since X is uniformly continuous on $[0, T]$ and since the mesh size $|\mathcal{P}^n|$ of the partition goes to 0 as $n \rightarrow \infty$, for every $\delta > 0$, there exists an N_δ such that

$$\sup_{n \geq N_\delta} \sup_{t \in \mathcal{P}_T^n} |X_{t'} - X_t| < \delta.$$

Hence, for $n \geq N_\delta$ and $t \in \mathcal{P}_T^n$,

$$|R_t^n| \leq \frac{1}{2}\varepsilon (X_{t'} - X_t)^2.$$

Summing over $t \in \mathcal{P}_T^n$ in (25.14) yields

$$\sum_{t \in \mathcal{P}_T^n} (F(X_{t'}) - F(X_t)) = F(X_t) - F(X_0)$$

and

$$\sum_{t \in \mathcal{P}_T^n} |R_t^n| \leq \varepsilon \sum_{t \in \mathcal{P}_T^n} (X_{t'} - X_t)^2 \xrightarrow{n \rightarrow \infty} \varepsilon \langle X \rangle_t < \infty.$$

As $\varepsilon > 0$ was arbitrary, we get $\sum_{t \in \mathcal{P}_T^n} |R_t^n| \xrightarrow{n \rightarrow \infty} 0$. We have (see Exercise 21.10.2)

$$\sum_{t \in \mathcal{P}_T^n} \frac{1}{2}F''(X_t)(X_{t'} - X_t)^2 \xrightarrow{n \rightarrow \infty} \frac{1}{2} \int_0^T F''(X_s) d\langle X \rangle_s.$$

Hence, in (25.14) the sum of the remaining terms also has to converge. That is, the limit in (25.12) exists.

As a direct consequence, we obtain the Itô formula for the Itô integral with respect to diffusions.

Theorem 25.27 (Itô formula for diffusions) Let $Y = M + A$ be a (generalized) diffusion (see Definition 25.23), where $M_t = \int_0^t \sigma_s dW_s$ and $A_t = \int_0^t b_s ds$. Let $F \in C^2(\mathbb{R})$. Then we have the Itô formula

$$\begin{aligned} F(Y_t) - F(Y_0) &= \int_0^t F'(Y_s) dM_s + \int_0^t F'(Y_s) dA_s + \frac{1}{2} \int_0^t F''(Y_s) d\langle M \rangle_s \\ &= \int_0^t F'(Y_s) \sigma_s dW_s + \int_0^t \left(F'(Y_s) b_s + \frac{1}{2} F''(Y_s) \sigma_s^2 \right) ds \end{aligned} \quad (25.15)$$

In particular, for Brownian motion,

$$F(W_t) - F(W_0) = \int_0^t F'(W_s) dW_s + \frac{1}{2} \int_0^t F''(W_s) ds \quad (25.16)$$

As an application of the Itô formula, we characterize Brownian motion as a continuous local martingale with a certain square variation process.

Theorem 25.28 (Lévy's characterization of Brownian motion) Let $X \in \mathcal{M}_{\text{loc},c}$ with $X_0 = 0$. Then the following are equivalent.

- (i) $(X_t^2 - t)_{t \geq 0}$ is a local martingale.
- (ii) $\langle X \rangle_t = t$ for all $t \geq 0$.
- (iii) X is a Brownian motion.

Proof (iii) $\implies$ (i) This is obvious.

(i) $\iff$ (ii) This is clear since the square variation process is unique.

(ii) $\implies$ (iii) It is enough to show that $X_t - X_s \sim \mathcal{N}_{0,t-s}$ given $\mathcal{F}_s$ for $t > s \geq 0$.

Employing the uniqueness theorem for characteristic functions, it is enough to show that (with $i = \sqrt{-1}$) for $A \in \mathcal{F}_s$ and $\lambda \in \mathbb{R}$, we have

$$\varphi_{A,\lambda}(t) := \mathbf{E} [e^{i\lambda(X_t - X_s)} \mathbb{1}_A] = \mathbf{P}[A] e^{-\lambda^2(t-s)/2}$$

Applying Itô's formula separately to the real and the imaginary parts, we obtain

$$e^{i\lambda X_t} - e^{i\lambda X_s} = \int_s^t i\lambda e^{i\lambda X_r} dX_r - \frac{1}{2} \int_s^t \lambda^2 e^{i\lambda X_r} dr$$

Therefore,

$$\begin{aligned} & \mathbf{E} [e^{i\lambda(X_t - X_s)} | \mathcal{F}_s] - 1 \\ &= \mathbf{E} \left[\int_s^t i\lambda e^{i\lambda(X_r - X_s)} dX_r | \mathcal{F}_s \right] - \frac{1}{2} \lambda^2 \mathbf{E} \left[\int_s^t e^{i\lambda(X_r - X_s)} dr | \mathcal{F}_s \right] \end{aligned}$$

Now $M_t := \text{Re} \int_s^t i\lambda e^{i\lambda(X_r - X_s)} dX_r$ and $N_t := \text{Im} \int_s^t i\lambda e^{i\lambda(X_r - X_s)} dX_r$, $t \geq s$ are continuous local martingales with

$$\langle M \rangle_t = \int_s^t \lambda^2 \sin(\lambda(X_r - X_s))^2 dr \leq \lambda^2(t-s)$$

and

$$\langle N \rangle_t = \int_s^t \lambda^2 \cos(\lambda(X_r - X_s))^2 dr \leq \lambda^2(t-s)$$

Thus, by Corollary 21.76, M and N are martingales. Hence we have

$$\mathbf{E} \left[\int_s^t i\lambda e^{i\lambda(X_r - X_s)} dX_r | \mathcal{F}_s \right] = 0$$

Since $A \in \mathcal{F}_s$, Fubini's theorem yields

$$\begin{aligned} \varphi_{A,\lambda}(t) - \varphi_{A,\lambda}(s) &= \mathbf{E} [e^{i\lambda(X_t - X_s)} \mathbb{1}_A] - \mathbf{P}[A] \\ &= -\frac{1}{2} \lambda^2 \int_s^t \mathbf{E} [e^{i\lambda(X_r - X_s)} \mathbb{1}_A] dr = -\frac{1}{2} \lambda^2 \int_s^t \varphi_{A,\lambda}(r) dr \end{aligned}$$

That is, $\varphi_{A,\lambda}$ is the solution of the linear differential equation

$$\varphi_{A,\lambda}(s) = \mathbf{P}[A] \quad \text{and} \quad \frac{d}{dt} \varphi_{A,\lambda}(t) = -\frac{1}{2} \lambda^2 \varphi_{A,\lambda}(t)$$

The unique solution is $\varphi_{A,\lambda}(t) = \mathbf{P}[A]e^{-\lambda^2(t-s)/2}$.

As a consequence of this theorem, we get that any continuous local martingale whose square variation process is absolutely continuous (as a function of time) can be expressed as an Itô integral with respect to some Brownian motion.

Theorem 25.29 (Itô's martingale representation theorem) Let M be a continuous local martingale with $M_0 = 0$ and absolutely continuous square variation $t \mapsto \langle M \rangle_t$. Then, on a suitable extension of the underlying probability space, there exists a Brownian motion W with

$$M_t = \int_0^t \sqrt{\frac{d\langle M \rangle_s}{ds}} dW_s \quad \text{for all } t \geq 0$$

Proof Assume that the probability space is rich enough to carry a Brownian motion $\widetilde{W}$ that is independent of M . Let

$$f_t := \lim_{n \rightarrow \infty} n (\langle M \rangle_t - \langle M \rangle_{t-1/n}) \quad \text{for } t > 0$$

Then f is a progressively measurable version of the Radon-Nikodym derivative $\frac{d\langle M \rangle_t}{dt}$. Clearly, $\int_0^T \mathbb{1}_{\{f_t > 0\}} f_t^{-1} d\langle M \rangle_t \leq T < \infty$ for all $T > 0$. Hence the following integrals are well-defined, and furthermore, as a sum of continuous martingales,

$$W_t := \int_0^t \mathbb{1}_{\{f_s > 0\}} f_s^{-1/2} dM_s + \int_0^t \mathbb{1}_{\{f_s = 0\}} d\widetilde{W}_s$$

is a continuous local martingale. By Theorem 25.22, we have

$$\begin{aligned} \langle W \rangle_t &= \int_0^t \mathbb{1}_{\{f_s > 0\}} f_s^{-1} d\langle M \rangle_s + \int_0^t \mathbb{1}_{\{f_s = 0\}} ds \\ &= \int_0^t \mathbb{1}_{\{f_s > 0\}} f_s^{-1} f_s ds + \int_0^t \mathbb{1}_{\{f_s = 0\}} ds = t \end{aligned}$$

Hence, by Theorem 25.28, W is a Brownian motion. On the other hand, we have

$$\begin{aligned} \int_0^t f_s^{1/2} dW_s &= \int_0^t \mathbb{1}_{\{f_s > 0\}} f_s^{1/2} f_s^{-1/2} dM_s + \int_0^t \mathbb{1}_{\{f_s = 0\}} f_s^{1/2} d\widetilde{W}_s \\ &= \int_0^t \mathbb{1}_{\{f_s > 0\}} dM_s \end{aligned}$$

However,

$$M_t - \int_0^t \mathbb{1}_{\{f_s > 0\}} dM_s = \int_0^t \mathbb{1}_{\{f_s = 0\}} dM_s$$

is a continuous local martingale with square variation $\int_0^t \mathbb{1}_{\{f_s = 0\}} d\langle M \rangle_s = 0$ and hence it almost surely equals zero. Therefore, we have $M_t = \int_0^t f_s^{1/2} dW_s$, as claimed.

We come next to a multidimensional version of the pathwise Itô formula. To this end, let $\mathcal{C}_{\text{qv}}^d$ be the space of continuous maps $X : [0, \infty) \rightarrow \mathbb{R}^d$, $t \mapsto X_t = (X_t^1, \dots, X_t^d)$ such that, for $k, l = 1, \dots, d$, the quadratic covariation (see Definition 21.58) $\langle X^k, X^l \rangle$ exists and is continuous. Further, let $C^2(\mathbb{R}^d)$ be the space of twice continuously differentiable functions F on $\mathbb{R}^d$ with partial derivatives $\partial_k F$ and $\partial_k \partial_l F$, $k, l = 1, \dots, d$. Denote by ∇ the gradient and by $\Delta = (\partial_1^2 + \dots + \partial_d^2)$ the Laplace operator.

Theorem 25.30 (Multidimensional pathwise Itô formula) Let $X \in \mathcal{C}_{\text{qv}}^d$ and $F \in C^2(\mathbb{R}^d)$. Then

$$F(X_T) - F(X_0) = \int_0^T \nabla F dX_s + \frac{1}{2} \sum_{k,l=1}^d \int_0^T \partial_k \partial_l F(X_s) d\langle X^k, X^l \rangle_s$$

where

$$\int_0^T \nabla F(X_s) dX_s := \lim_{n \rightarrow \infty} \sum_{t \in \mathcal{P}_T^n} \sum_{k=1}^d \partial_k F(X_t) (X_t^k - X_{t'}^k)$$

Proof This works just as in the one-dimensional case. We leave the details as an exercise.

Remark 25.31 If each of the integrals $\int_0^T \partial_k F(X_s) dX_s^k$ exists, then

$$\int_0^T \nabla F(X_s) dX_s = \sum_{k=1}^d \int_0^T \partial_k F(X_s) dX_s^k$$

Note that existence of the individual integrals does not follow from the existence of the integral $\int_0^T \nabla F(X_s) dX_s$.

Corollary 25.32 (Product rule) If $X, Y, X - Y, X + Y \in \mathcal{C}_{qv}$, then

$$X_T Y_T = X_0 Y_0 + \int_0^T Y_s dX_s + \int_0^T X_s dY_s + \langle X, Y \rangle_T \quad \text{for all } T \geq 0$$

if both integrals exist. In particular, the product rule holds if X and Y are continuous local martingales.

Proof By assumption (and using the polarization formula), the covariation $\langle X, Y \rangle$ exists. Applying Theorem 25.30 with $F(x, y) = xy$, the claim follows.

For continuous local martingales, by Exercise 25.2.1, there exists a suitable sequence of partitions $\mathcal{P}$ such that the integrals exist (pathwise).

Now let $Y = M + A$ be a d -dimensional generalized diffusion. Hence

$$M_t^k = \sum_{l=1}^d \int_0^t \sigma_s^{k,l} dW_s^l \quad \text{and} \quad A_t^k = \int_0^t b_s^k ds \quad \text{for } t \geq 0, k = 1, \dots, d$$

Here $W = (W^1, \dots, W^d)$ is a d -dimensional Brownian motion and $\sigma^{k,l}$ (respectively b^k) are progressively measurable, locally square integrable (respectively locally integrable) stochastic processes for every $k, l = 1, \dots, d$. Since $\langle W^k, W^l \rangle_t = t \cdot \mathbb{1}_{\{k=l\}}$, we have $\langle Y^k, Y^l \rangle_t = \langle M^k, M^l \rangle_t = \int_0^t a_s^{k,l} ds$, where

$$a_s^{k,l} := \sum_{i=1}^d \sigma_s^{k,i} \sigma_s^{i,l}$$

is the covariance matrix of the diffusion M . In particular, we have $M \in \mathcal{C}_{qv}^d$ almost surely. Note that (by Exercise 25.2.1), there exists a partition sequence $\mathcal{P}$ such that the integrals $\int_0^t \sigma_s^{k,l} \partial_k F(Y_s) dW_s^l$ in (25.17) exist in the pathwise sense. As a corollary of the multidimensional pathwise Itô formula (Theorem 25.30 and Remark 25.31), we thus get the following theorem.

Theorem 25.33 (Multidimensional Itô formula) Let Y be as above and let $F \in C^2(\mathbb{R}^d)$. Then

$$\begin{aligned}
F(Y_T) - F(Y_0) &= \int_0^T \nabla F(Y_s) dY_s + \frac{1}{2} \sum_{k,l=1}^d \int_0^T \partial_k \partial_l F(Y_s) d\langle M^k, M^l \rangle_s \\
&= \sum_{k,l=1}^d \int_0^t \sigma_s^{k,l} \partial_k F(Y_s) dW_s^l + \sum_{k=1}^d \int_0^t b_s^k \partial_k F(Y_s) ds \\
&\quad + \frac{1}{2} \sum_{k,l=1}^d \int_0^t a_s^{k,l} \partial_k \partial_l F(Y_s) ds
\end{aligned} \tag{25.17}$$

In particular, for Brownian motion, we have

$$F(W_t) - F(W_0) = \sum_{k=1}^d \int_0^t \partial_k F(W_s) dW_s^k + \frac{1}{2} \int_0^t \Delta F(W_s) ds \tag{25.18}$$

Corollary 25.34 The process $(F(W_t))_{t \geq 0}$ is a continuous local martingale if and only if F is harmonic (that is, $\Delta F \equiv 0$).

Proof If F is harmonic, then as a sum of Itô integrals, $F(W_t) = F(W_0) + \sum_{k=1}^d \int_0^t \partial_k F(W_s) dW_s^k$ is a continuous local martingale.

On the other hand, if F is a continuous local martingale, then as a difference of continuous local martingales, $\int_0^t \Delta F(W_s) ds$ is also a continuous local martingale. As $t \mapsto \int_0^t \Delta F(W_s) ds$ has finite variation, we have $\int_0^t \Delta F(W_s) ds = 0$ for all $t \geq 0$ almost surely (by Corollary 21.72). Hence $\Delta F \equiv 0$.

Corollary 25.35 (Time-dependent Itô formula) If $F \in C^{2,1}(\mathbb{R}^d \times \mathbb{R})$, then

$$\begin{aligned}
&F(W_T, T) - F(W_0, 0) \\
&= \sum_{k=1}^d \int_0^T \partial_k F(W_s, s) dW_s^k + \int_0^T \left(\partial_{d+1} + \frac{1}{2} (\partial_1^2 + \dots + \partial_d^2) \right) F(W_s, s) ds
\end{aligned}$$

Proof Apply Theorem 25.33 to $Y = (W_t^1, \dots, W_t^d, t)_{t \geq 0}$.

Exercise 25.3.1 (Fubini's theorem for Itô integrals) Let $X \in \mathcal{C}_{\text{qv}}$ and assume that $g : [0, \infty)^2 \rightarrow \mathbb{R}$ is continuous and (in the interior) twice continuously differentiable in the second coordinate with derivative $\partial_2 g$. Use the product rule (Corollary 25.32) to show that

$$\int_0^s \left(\int_0^t g(u, v) du \right) dX_v = \int_0^t \left(\int_0^s g(u, v) dX_v \right) du$$

and

$$\int_0^s \left(\int_0^v g(u, v) du \right) dX_v = \int_0^s \left(\int_u^s g(u, v) dX_v \right) du$$

Exercise 25.3.2 (Stratonovich integral) Let $\mathcal{P}$ be an admissible sequence of partitions, $X \in \mathcal{C}_{\text{qv}}^{\mathcal{P}}$ and $F \in C^2(\mathbb{R})$ with derivative $f = F'$. Show that, for every $t \geq 0$, the Stratonovich integral

$$\int_0^T f(X_t) \circ dX_t := \lim_{n \rightarrow \infty} \sum_{t \in \mathcal{P}_T^n} f\left(\frac{X_{t'} + X_t}{2}\right) (X_{t'} - X_t)$$

is well-defined, and that the classical substitution rule

$$F(X_T) - F(X_0) = \int_0^T F'(X_t) \circ dX_t$$

holds. Show that, in contrast with the Itô integral, the Stratonovich integral with respect to a continuous local martingale is, in general, not a local martingale.

5.25.4 25.4 Dirichlet Problem and Brownian Motion

As for discrete Markov chains (compare Section 19.1), the solutions of the Dirichlet problem in a domain $G \subset \mathbb{R}^d$ can be expressed in terms of a d -dimensional Brownian motion that is stopped upon hitting the boundary G .

In the following, let $G \subset \mathbb{R}^d$ be a bounded open set.

Definition 25.36 (Dirichlet problem) Let $f : \partial G \rightarrow \mathbb{R}$ be continuous. A function $u : \bar{G} \rightarrow \mathbb{R}$ is called a solution of the Dirichlet problem on G with boundary value f if u is continuous, twice differentiable in G and

$$\begin{aligned} \Delta u(x) &= 0 \text{ for } x \in G, \\ u(x) &= f(x) \text{ for } x \in \partial G. \end{aligned} \quad (25.19)$$

For sufficiently smooth domains, there always exists a solution of the Dirichlet problem (see, e.g., [79, Section 4.4]). If there is a solution, then as a consequence of Theorem 25.38, it is unique.

In the following, let $W = (W^1, \dots, W^d)$ be a d -dimensional Brownian motion with respect to a filtration $\mathbb{F}$ that satisfies the usual conditions. We write $\mathbf{P}_x$ and $\mathbf{E}_x$ for probabilities and expectations if W is started at $W_0 = x = (x^1, \dots, x^d) \in \mathbb{R}^d$. If $A \subset \mathbb{R}^d$ is open, then

$$\tau_{A^c} := \inf \{t > 0 : W_t \in A^c\}$$

is an $\mathbb{F}$ -stopping time (see Exercise 21.4.4). Since G is bounded, we have $G \subset (-a, a) \times \mathbb{R}^{d-1}$ for some $a > 0$. Thus $\tau_{G^c} \leq \tau_{((-a,a) \times \mathbb{R}^{d-1})^c}$. By Exercise 21.2.4 (applied to W^1), for $x \in G$,

$$\mathbf{E}_x[\tau_{G^c}] \leq \mathbf{E}_x\left[\tau_{((-a,a) \times \mathbb{R}^{d-1})^c}\right] = (a - x^1)(a + x^1) < \infty. \quad (25.20)$$

In particular, $\tau_{G^c} < \infty$ $\mathbf{P}_x$ -almost surely. Hence $W_{\tau_{G^c}}$ is a $\mathbf{P}_x$ -almost surely welldefined random variable with values in ∂G .

Definition 25.37 For $x \in G$, denote by

$$\mu_{x,G} = \mathbf{P}_x \circ W_{\tau_{G^c}}^{-1}$$

the harmonic measure on ∂G .

Theorem 25.38 If u is a solution of the Dirichlet problem on G with boundary value f , then

$$u(x) = \mathbf{E}_x[f(W_{\tau_{G^c}})] = \int_{\partial G} f(y) \mu_{x,G}(dy) \quad \text{for } x \in G \quad (25.21)$$

In particular, the solution of the Dirichlet problem is always unique.

Proof Let $G_1 \subset G_2 \subset \dots$ be a sequence of open sets with $x \in G_1, G_n \uparrow G$ and $\bar{G}_n \subset G$ for every $n \in \mathbb{N}$. Hence, in particular, every $\bar{G}_n$ is compact and thus ∇u is bounded on $\bar{G}_n$. We abbreviate $\tau := \tau_{G^c}$ and $\tau_n := \tau_{G_n^c}$.

As u is harmonic (that is, $\Delta u = 0$), by the Itô formula, for $t < \tau$,

$$u(W_t) = u(W_0) + \int_0^t \nabla u(W_s) dW_s = u(W_0) + \sum_{k=1}^d \int_0^t \partial_k u(W_s) dW_s^k \quad (25.22)$$

In particular, $M := (u(W_t))_{t \in [0, \tau]}$ is a local martingale up to τ (however, in general, it is not a martingale). For $t < \tau_n$, we have

$$(\partial_k u(W_s))^2 \leq C_n := \sup_{y \in \bar{G}_n} \|\nabla u(y)\|_2^2 < \infty \quad \text{for all } k = 1, \dots, d.$$

Hence, by (25.20),

$$\mathbf{E} \left[\int_0^{\tau_n} (\partial_k u(W_s))^2 ds \right] \leq C_n \mathbf{E}_x [\tau_n] \leq C_n \mathbf{E}[\tau] < \infty$$

Thus, by Theorem 25.17(ii), for every $n \in \mathbb{N}$, the stopped process M^{τ_n} is a martingale. Therefore,

$$\mathbf{E}_x [u(W_{\tau_n})] = \mathbf{E}_x [M_{\tau_n}] = \mathbf{E}_x [M_0] = u(x). \quad (25.23)$$

As W is continuous and $\tau_n \uparrow \tau$, we have $W_{\tau_n} \xrightarrow{n \rightarrow \infty} W_\tau \in \partial G$. Since u is continuous, we get

$$u(W_{\tau_n}) \xrightarrow{n \rightarrow \infty} u(W_\tau) = f(W_\tau). \quad (25.24)$$

Again, since u is continuous and $\bar{G}$ is compact, u is bounded. By the dominated convergence theorem, (25.24) implies convergence of the expectations; that is (also using (25.23)),

$$u(x) = \lim_{n \rightarrow \infty} \mathbf{E}_x [u(W_{\tau_n})] = \mathbf{E}_x [f(W_\tau)]$$

Exercise 25.4.1 Let $G = \mathbb{R} \times (0, \infty)$ be the open upper half plane of $\mathbb{R}^2$ and $x = (x_1, x_2) \in G$. Show that $\tau_{G^c} < \infty$ almost surely and that the harmonic measure $\mu_{x,G}$ on $\mathbb{R} \cong \partial G$ is the Cauchy distribution with scale parameter x_2 that is shifted by x_1 : $\mu_{x,G} = \delta_{x_1} * \text{Cau}_{x_2}$.

Exercise 25.4.2 Let $d \geq 3$ and let $G = \mathbb{R}^{d-1} \times (0, \infty)$ be an open half space of $\mathbb{R}^d$. Let $x = (x_1, \dots, x_d) \in G$. Show that $\tau_{G^c} < \infty$ almost surely and that the harmonic measure $\mu_{x,G}$ on $\mathbb{R}^{d-1} \cong \partial G$ has the density

$$\frac{\mu_{x,G}(dy)}{dy} = \frac{\Gamma(d/2)}{\pi^{d/2}} \frac{x_d}{\sqrt{(x_1 - y_1)^2 + \dots + (x_{d-1} - y_{d-1})^2 + x_d^2}}.$$

Exercise 25.4.3 Let $r > 0$ and let $B_r(0) \subset \mathbb{R}^d$ be the open ball with radius r centered at the origin. For $x \in B_r(0)$, determine the harmonic measure $\mu_{x,B_r(0)}$.

5.25.5 25.5 Recurrence and Transience of Brownian Motion

By Pólya's theorem (Theorem 17.39), symmetric simple random walk $(X_n)_{n \in \mathbb{N}}$ on $\mathbb{Z}^d$ is recurrent (that is, it visits every point infinitely often) if and only if $d \leq 2$. If $d > 2$, then the random walk is transient and eventually leaves every bounded set $A \subset \mathbb{Z}^d$. To give a slightly different (though equivalent) formulation of this,

$$\liminf_{n \rightarrow \infty} \|X_n\| = 0 \quad \text{a.s.} \quad \Longleftrightarrow \quad d \leq 2$$

and

$$\lim_{n \rightarrow \infty} \|X_n\| = \infty \quad \text{a.s.} \quad \Longleftrightarrow \quad d > 2.$$

The main result of this section is that a similar dichotomy also holds for Brownian motion.

Theorem 25.39 Let $W = (W^1, \dots, W^d)$ be a d -dimensional Brownian motion.

(i) If $d \leq 2$, then W is recurrent in the sense that

$$\liminf_{t \rightarrow \infty} \|W_t - y\| = 0 \quad \text{a.s.} \quad \text{for every } y \in \mathbb{R}^d$$

In particular, almost surely the path $\{W_t : t \geq 0\}$ is dense in $\mathbb{R}^d$.

(ii) If $d > 2$, then W is transient in the sense that

$$\lim_{t \rightarrow \infty} \|W_t\| = \infty \quad \text{a.s.}$$

and for any $y \in \mathbb{R}^d \setminus \{0\}$, we have $\inf \{\|W_t - y\| : t \geq 0\} > 0$ almost surely.

The basic idea of the proof is to use a suitable Dirichlet problem (and the result of Section 25.4) to compute the probabilities for W to hit certain balls,

$$B_R(x) := \{y \in \mathbb{R}^d : \|x - y\| < R\}.$$

Let $0 < r < R < \infty$ and let $G_{r,R}$ be the annulus

$$G_{r,R} := B_R(0) \setminus \bar{B}_r(0) = \{x \in \mathbb{R}^d : r < \|x\| < R\}$$

Recall that, for closed $A \subset \mathbb{R}^d$, we write $\tau_A = \inf \{t > 0 : W_t \in A\}$ for the stopping time of first entrance into A . We further write

$$\tau_s := \inf \{t > 0 : \|W_t\| = s\} \quad \text{and} \quad \tau_{r,R} = \inf \{t > 0 : W_t \notin G_{r,R}\}.$$

If we start W at $W_0 \in G_{r,R}$, then clearly $\tau_{r,R} = \tau_r \wedge \tau_R$. On the boundary of $G_{r,R}$, define the function f by

$$f(x) = \begin{cases} 1, & \text{if } \|x\| = r \\ 0, & \text{if } \|x\| = R \end{cases} \quad (25.25)$$

Define $u_{r,R} : \bar{G}_{r,R} \rightarrow \mathbb{R}$ by

$$u_{r,R}(x) = \frac{V(\|x\|) - V(R)}{V(r) - V(R)}$$

where $V : (0, \infty) \rightarrow \mathbb{R}$ is Newton's potential function

$$V(s) = V_d(s) = \begin{cases} s, & \text{if } d = 1 \\ \log(s), & \text{if } d = 2 \\ -s^{2-d}, & \text{if } d > 2 \end{cases} \quad (25.26)$$

It is easy to check that $\varphi : \mathbb{R}^d \setminus \{0\} \rightarrow \mathbb{R}, x \mapsto V_d(\|x\|)$ is harmonic (that is, $\Delta \varphi \equiv 0$). Hence $u_{r,R}$ is the solution of the Dirichlet problem on $G_{r,R}$ with boundary value f .

By Theorem 25.38, for $x \in G_{r,R}$,

$$\mathbf{P}_x[\tau_{r,R} = \tau_r] = \mathbf{P}_x[\|W_{\tau_{r,R}}\| = r] = \mathbf{E}_x[f(W_{\tau_{r,R}})] = u_{r,R}(x). \quad (25.27)$$

Theorem 25.40 For $r > 0$ and $x, y \in \mathbb{R}^d$ with $\|x - y\| > r$, we have

$$\mathbf{P}_x[W_t \in B_r(y) \text{ for some } t > 0] = \begin{cases} 1, & \text{if } d \leq 2, \\ \left(\frac{\|x-y\|}{r}\right)^{2-d}, & \text{if } d > 2. \end{cases}$$

Proof Without loss of generality, assume $y = 0$. Then

$$\begin{aligned}\mathbf{P}_x[\tau_r < \infty] &= \lim_{R \rightarrow \infty} \mathbf{P}_x[\tau_{r,R} = \tau_r] = \lim_{R \rightarrow \infty} \frac{V(\|x\|) - V(R)}{V(r) - V(R)} \\ &= \begin{cases} 1, & \text{if } d = 2, \\ \frac{V_d(\|x\|)}{V_d(r)}, & \text{if } d > 2, \end{cases}\end{aligned}$$

since $\lim_{R \rightarrow \infty} V_d(R) = \infty$ if $d \leq 2$ and $= 0$ if $d > 2$.

Proof of Theorem 25.39 Using the strong Markov property of Brownian motion, we get for $r > 0$

$$\begin{aligned}\mathbf{P}_x \left[\liminf_{t \rightarrow \infty} \|W_t\| < s \right] &= \mathbf{P}_x \left[\bigcup_{s \in (0, r)} \bigcap_{R > \|x\|} \{ \|W_t\| < r \text{ for some } t > \tau_R \} \right] \\ &= \sup_{s \in (0, r)} \inf_{R > \|x\|} \mathbf{P}_x [\|W_t\| \leq s \text{ for some } t > \tau_R] \\ &= \sup_{s \in (0, r)} \inf_{R > \|x\|} \mathbf{P}_x [\mathbf{P}_{W_{\tau_R}}[\tau_s < \infty]]\end{aligned}$$

However, by Theorem 25.40 (since $\|W_{\tau_R}\| = R$ for $R > \|x\|$), we have

$$\mathbf{P}_{W_{\tau_R}}[\tau_s < \infty] = \begin{cases} 1, & \text{if } d \leq 2 \\ (s/R)^{d-2}, & \text{if } d > 2 \end{cases}$$

Therefore,

$$\mathbf{P} \left[\liminf_{t \rightarrow \infty} \|W_t\| < r \right] = \begin{cases} 1, & \text{if } d \leq 2 \\ 0, & \text{if } d > 2 \end{cases}$$

This implies the claim.

Definition 25.41 (Polar set) A set $A \subset \mathbb{R}^d$ is called polar if

$$\mathbf{P}_x[W_t \notin A \text{ for all } t > 0] = 1 \quad \text{for all } x \in \mathbb{R}^d.$$

Theorem 25.42 If $d = 1$, then only the empty set is polar. If $d \geq 2$, then $\{y\}$ is polar for every $y \in \mathbb{R}^d$.

Proof For $d = 1$, the statement is obvious since

$$\limsup_{t \rightarrow \infty} W_t = \infty \quad \text{and} \quad \liminf_{t \rightarrow \infty} W_t = -\infty \quad \text{a.s.}$$

Hence, due to the continuity of W , every point $y \in \mathbb{R}$ will be hit (infinitely often).

Now let $d \geq 2$. Without loss of generality, assume $y = 0$. If $x \neq 0$, then

$$\begin{aligned}\mathbf{P}_x[\tau_{\{0\}} < \infty] &= \lim_{R \rightarrow \infty} \mathbf{P}_x[\tau_{\{0\}} < \tau_R] \\ &= \lim_{R \rightarrow \infty} \inf_{r > 0} \mathbf{P}_x[\tau_{r,R} = \tau_r] \\ &= \lim_{R \rightarrow \infty} \inf_{r > 0} u_{r,R}(x) = 0\end{aligned} \tag{25.28}$$

since $V_d(r) \xrightarrow{r \rightarrow 0} -\infty$ if $d \geq 2$.

On the other hand, if $x = 0$, then the strong Markov property of Brownian motion (and the fact that $\mathbf{P}_0[W_t = 0] = 0$ for all $t > 0$) implies

$$\begin{aligned}
\mathbf{P}_0 [\tau_{\{0\}} < \infty] &= \sup_{t>0} \mathbf{P}_0 [W_s = 0 \text{ for some } s \geq t] \\
&= \sup_{t>0} \mathbf{P}_0 [\mathbf{P}_{W_t} [\tau_{\{0\}} < \infty]] = 0
\end{aligned}$$

Note that in the last step, we used (25.28).

5.26 26 Stochastic Differential Equations

Stochastic differential equations describe the time evolution of certain continuous Markov processes with values in $\mathbb{R}^n$. In contrast with classical differential equations, in addition to the derivative of the function, there is a term that describes the random fluctuations that are coded as an Itô integral with respect to a Brownian motion. Depending on how seriously we take the concrete Brownian motion as the driving force of the noise, we speak of strong and weak solutions. In the first section, we develop the theory of strong solutions under Lipschitz conditions for the coefficients. In the second section, we develop the so-called (local) martingale problem as a method of establishing weak solutions. In the third section, we present some examples in which the method of duality can be used to prove weak uniqueness.

As stochastic differential equations are a very broad subject, and since things quickly become very technical, we only excursively touch some of the most important results, partly without proofs, and illustrate them with examples.

5.26.1 26.1 Strong Solutions

Consider a stochastic differential equation (SDE) of the type

$$\begin{aligned}
X_0 &= \xi \\
dX_t &= \sigma(t, X_t) dW_t + b(t, X_t) dt
\end{aligned} \tag{26.1}$$

Here $W = (W^1, \dots, W^m)$ is an m -dimensional Brownian motion, ξ is an $\mathbb{R}^n$ valued random variable with distribution μ that is independent of W , $\sigma(t, x) = (\sigma_{ij}(t, x))_{\substack{i=1, \dots, n \\ j=1, \dots, m}}$ is a real $n \times m$ matrix and $b(t, x) = (b_i(t, x))_{i=1, \dots, n}$ is an n dimensional vector. Assume the maps $(t, x) \mapsto \sigma_{ij}(t, x)$ and $(t, x) \mapsto b_i(t, x)$ are measurable.

By a solution of (26.1) we understand a continuous adapted stochastic process X with values in $\mathbb{R}^n$ that satisfies the integral equation

$$X_t = \xi + \int_0^t \sigma(s, X_s) dW_s + \int_0^t b(s, X_s) ds \quad \mathbf{P}\text{-a.s. for all } t \geq 0 \tag{26.2}$$

Written in full, this is

$$X_t^i = \xi^i + \sum_{j=1}^m \int_0^t \sigma_{ij}(s, X_s) dW_s^j + \int_0^t b_i(s, X_s) ds \quad \text{for all } i = 1, \dots, n$$

Now the following problem arises: To which filtration $\mathbb{F}$ do we wish X to be adapted? Should it be the filtration that is generated by ξ and W , or do we allow $\mathbb{F}$ to be larger? Already for ordinary differential equations, depending on the equation, uniqueness of the solution may fail (although existence is usually not a problem); for example, for $f' = |f|^{1/3}$. If $\mathbb{F}$ is larger than the filtration generated by W , then we can define further random variables that select one out of a variety of possible solutions. We thus have more possibilities for solutions than if $\mathbb{F} = \sigma(W)$. Indeed, it will turn out that in some situations for the existence of a solution, it

is necessary to allow a larger filtration. Roughly speaking, X is a strong solution of (26.1) if (26.2) holds and if X is adapted to $\mathbb{F} = \sigma(W)$. On the other hand, X is a weak solution if X is adapted to a larger filtration $\mathbb{F}$ with respect to which W is still a martingale. Weak solutions will be dealt with in Section 26.2.

Definition 26.1 (Strong solution) We say that the stochastic differential equation (SDE) (26.1) has a strong solution X if there exists a map $F : \mathbb{R}^n \times C([0, \infty); \mathbb{R}^m) \rightarrow C([0, \infty); \mathbb{R}^n)$ with the following properties:

- (i) For every $t \geq 0$, the map $(x, w) \mapsto F(x, w)$ is measurable with respect to $\mathcal{B}(\mathbb{R}^n) \otimes \mathcal{G}_t^m - \mathcal{G}_t^n$, where (for $k = m$ or $k = n$) $\mathcal{G}_t^k := \sigma(\pi_s : s \in [0, t])$ is the σ algebra generated by the coordinate maps $\pi_s : C([0, \infty); \mathbb{R}^k) \rightarrow \mathbb{R}, w \mapsto w(s)$.
- (ii) The process $X = F(\xi, W)$ satisfies (26.2).

Condition (i) says that the path $(X_s)_{s \in [0, t]}$ depends only on ξ and $(W_s)_{s \in [0, t]}$ but not on further information. In particular, X is adapted to $\mathcal{F}_t = \sigma(\xi, W_s : s \in [0, t])$ and is progressively measurable; hence the Itô integral in (26.2) is well-defined if σ and b do not grow too quickly for large x .

Remark 26.2 Clearly, a strong solution of an SDE is a generalized n -dimensional diffusion. If the coefficients σ and b are independent of t , then the solution is an n -dimensional diffusion.

Remark 26.3 Let X be a strong solution and let F be as in Definition 26.1. If W' is an m -dimensional Brownian motion on a space $(\Omega', \mathcal{F}', \mathbf{P}')$ with filtration $\mathbb{F}'$, and if ξ' is independent of W' and is $\mathcal{F}'_0$ -measurable, then $X' = F(\xi', W')$ satisfies the integral equation (26.2). Hence, it is a strong solution of (26.1) with W replaced by W' . Thus the existence of a strong solution does not depend on the actual realization of the Brownian motion or on the filtration $\mathbb{F}$. $\square$

Definition 26.4 We say that the SDE (26.1) has a unique strong solution if there exists an F as in Definition 26.1 such that:

- (i) If W is an m -dimensional Brownian motion on some probability space $(\Omega, \mathcal{F}, \mathbf{P})$ with filtration $\mathbb{F}$ and if ξ is an $\mathcal{F}_0$ -measurable random variable that is independent of W and such that $\mathbf{P} \circ \xi^{-1} = \mu$, then $X := F(\xi, W)$ is a solution of (26.2).
- (ii) For every solution (X, W) of (26.2), we have $X = F(\xi, W)$.

Example 26.5 Let $m = n = 1, b \in \mathbb{R}$ and $\sigma > 0$. The Ornstein-Uhlenbeck process

$$X_t := e^{bt}\xi + \sigma \int_0^t e^{(t-s)b} dW_s, \quad t \geq 0 \quad (26.3)$$

is a strong solution of the SDE $dX_t = \sigma dW_t + bX_t dt$ and

$$dX_t = \sigma dW_t + bX_t dt$$

In the language of Definition 26.1, we have (in the sense of the pathwise Itô integral with respect to w)

$$F(x, w) = \left(t \mapsto e^{bt}x + \int_0^t e^{(t-s)b} dw(s) \right)$$

for all $w \in \mathcal{C}_{\text{qv}}$ (that is, with continuous square variation). Since $\mathbf{P}[W \in \mathcal{C}_{\text{qv}}] = 1$, we can define $F(x, w) = 0$ for $w \in C([0, \infty); \mathbb{R}) \setminus \mathcal{C}_{\text{qv}}$.

Indeed, by Fubini's theorem for Itô integrals, we have (Exercise 25.3.1)

$$\begin{aligned}
& \xi + \int_0^t \sigma dW_s + \int_0^t bX_s ds \\
&= \xi + \sigma W_t + \int_0^t b e^{bs} \xi ds + \int_0^t \sigma b \left(\int_0^s e^{b(s-r)} dW_r \right) ds \\
&= \xi + \sigma W_t + (e^{bt} - 1) \xi + \int_0^t \sigma \left(\int_r^t b e^{b(s-r)} ds \right) dW_r \\
&= e^{bt} \xi + \int_0^t (\sigma + (e^{b(t-r)} - 1) \sigma) dW_r \\
&= X_t
\end{aligned}$$

It can be shown (see Theorem 26.8) that the solution is also (strongly) unique. $\square$

Example 26.6 Let $\alpha, \beta \in \mathbb{R}$. The one-dimensional SDE $X_0 = \xi$ and

$$dX_t = \alpha X_t dW_t + \beta X_t dt \quad (26.4)$$

has the strong solution

$$X_t = \xi \exp \left(\alpha W_t + \left(\beta - \frac{\alpha^2}{2} \right) t \right)$$

In the language of Definition 26.1, we have $\sigma(t, x) = \alpha x$, $b(t, x) = \beta x$ and

$$F(x, w) = \left(t \mapsto x \exp \left(\alpha w(t) + \left(\beta - \frac{\alpha^2}{2} \right) t \right) \right)$$

for all $w \in C([0, \infty); \mathbb{R})$ and $x \in \mathbb{R}$. Indeed, by the time-dependent Itô formula (Corollary 25.35),

$$X_t = \xi + \int_0^t \alpha X_s dW_s + \int_0^t \left(\left(\beta - \frac{\alpha^2}{2} \right) + \frac{1}{2} \alpha^2 \right) X_s ds$$

Also in this case, we have strong uniqueness of the solution (see Theorem 26.8). The process X is called a geometric Brownian motion and, for example, serves in the so-called Black-Scholes model as the process of stock prices.

We give a simple criterion for existence and uniqueness of strong solutions. For an $n \times m$ matrix A , define the Hilbert-Schmidt norm

$$\|A\| = \sqrt{\text{trace}(AA^T)} = \sqrt{\sum_{i=1}^n \sum_{j=1}^m A_{i,j}^2}. \quad (26.5)$$

For $b \in \mathbb{R}^n$, we use the Euclidean norm $\|b\|$. Since all norms on finite-dimensional vector spaces are equivalent, it is not important exactly which norm we use. However, the Hilbert-Schmidt norm simplifies the computations, as the following lemma shows.

Lemma 26.7 Let $t \mapsto H(t) = (H_{ij}(t))_{i=1, \dots, n, j=1, \dots, m}$ be progressively measurable and $\mathbf{E} \left[\int_0^T H_{ij}^2(t) dt \right] < \infty$ for all i, j . Then

$$\mathbf{E} \left[\left\| \int_0^T H(t) dW_t \right\|^2 \right] = \mathbf{E} \left[\int_0^T \|H(t)\|^2 dt \right] \quad (26.6)$$

where $\|H\|$ is the Hilbert-Schmidt norm from (26.5).

Proof For $i = 1, \dots, n$, the process $I_i(t) := \sum_{j=1}^m \int_0^t H_{ij}(s) dW_s^j$ is a continuous martingale with square variation process $\langle I_i \rangle_t = \int_0^t \sum_{j=1}^m H_{ij}^2(s) ds$. Hence

$$\mathbf{E} [(I_i(T))^2] = \mathbf{E} \left[\int_0^T \sum_{j=1}^m H_{ij}^2(s) ds \right]$$

The left-hand side in (26.6) equals

$$\sum_{i=1}^n \mathbf{E} [(I_i(T))^2] = \mathbf{E} \left[\int_0^T \sum_{i=1}^n \sum_{j=1}^m H_{ij}^2(s) ds \right]$$

Hence the claim follows by the definition of $\|H(s)\|^2$.

Theorem 26.8 Let b and σ be Lipschitz continuous in the first coordinate. That is, we assume that there exists a $K > 0$ such that, for all $x, x' \in \mathbb{R}^n$ and $t \geq 0$,

$$\|\sigma(x, t) - \sigma(x', t)\| + \|b(x, t) - b(x', t)\| \leq K \|x - x'\|. \quad (26.7)$$

Further, assume the growth condition

$$\|\sigma(t, x)\|^2 + \|b(t, x)\|^2 \leq K^2 (1 + \|x\|^2) \quad \text{for all } x \in \mathbb{R}^n, t \geq 0. \quad (26.8)$$

Then, for every initial point $X_0 = x \in \mathbb{R}^n$, there exists a unique strong solution X of the SDE (26.1). This solution is a Markov process and in the case where σ and b do not depend on t , it is a strong Markov process.

As the main tool, we need the following lemma.

Lemma 26.9 (Gronwall) Let $f, g : [0, T] \rightarrow \mathbb{R}$ be integrable and let $C > 0$ such that

$$f(t) \leq g(t) + C \int_0^t f(s) ds \quad \text{for all } t \in [0, T] \quad (26.9)$$

Then

$$f(t) \leq g(t) + C \int_0^t e^{C(t-s)} g(s) ds \quad \text{for all } t \in [0, T]$$

In particular, if $g(t) \equiv G$ is constant, then $f(t) \leq Ge^{Ct}$ for all $t \in [0, T]$.

Proof Let $F(t) = \int_0^t f(s) ds$ and $h(t) = F(t)e^{-Ct}$. Then, by (26.9),

$$\frac{d}{dt} h(t) = f(t)e^{-Ct} - CF(t)e^{-Ct} \leq g(t)e^{-Ct}$$

Integration yields

$$F(t) = e^{Ct} h(t) \leq \int_0^t e^{C(t-s)} g(s) ds$$

Substituting this into (26.9) gives

$$f(t) \leq g(t) + CF(t) \leq g(t) + C \int_0^t g(s) e^{C(t-s)} ds$$

Proof of Theorem 26.8 It is enough to show that, for every $T < \infty$, there exists a unique strong solution up to time T .

Uniqueness. We first show uniqueness of the solution. Let X and X' be two solutions of (26.2). Then

$$X_t - X'_t = \int_0^t (b(s, X_s) - b(s, X'_s)) ds + \int_0^t (\sigma(s, X_s) - \sigma(s, X'_s)) dW_s$$

Hence

$$\begin{aligned} \|X_t - X'_t\|^2 &\leq 2 \left\| \int_0^t (b(s, X_s) - b(s, X'_s)) ds \right\|^2 \\ &\quad + 2 \left\| \int_0^t (\sigma(s, X_s) - \sigma(s, X'_s)) dW_s \right\|^2 \end{aligned} \quad (26.10)$$

For the first summand in (26.10), use the Cauchy-Schwarz inequality, and for the second one use Lemma 26.7 to obtain

$$\begin{aligned} \mathbf{E} [\|X_t - X'_t\|^2] &\leq 2t \int_0^t \mathbf{E} [\|b(s, X_s) - b(s, X'_s)\|^2] ds \\ &\quad + 2 \int_0^t \mathbf{E} [\|\sigma(s, X_s) - \sigma(s, X'_s)\|^2] ds \end{aligned}$$

Write $f(t) = \mathbf{E} [\|X_t - X'_t\|^2]$ and $C := 2(T+1)K^2$. Then $f(t) \leq C \int_0^t f(s) ds$. Hence Gronwall's lemma (with $g \equiv 0$) yields $f \equiv 0$.

Existence. We use a version of the Picard iteration scheme. For $N \in \mathbb{N}_0$, recursively define processes X^N by $X^0_t \equiv x$ and

$$X^N_t := x + \int_0^t b(s, X^{N-1}_s) ds + \int_0^t \sigma(s, X^{N-1}_s) dW_s \quad \text{for } N \in \mathbb{N} \quad (26.11)$$

Using the growth condition (26.8), it can be shown inductively that

$$\begin{aligned} \int_0^T \mathbf{E} [\|X^N_t\|^2] dt &\leq 2(T+1)K^2 \left(T + \int_0^T \mathbf{E} [\|X^{N-1}_t\|^2] dt \right) \\ &\leq (2T(T+1)K^2)^N (1 + \|x\|^2) < \infty, \quad N \in \mathbb{N} \end{aligned}$$

Hence, at each step, the Itô integral is well-defined.

Consider now the differences

$$X^{N+1}_t - X^N_t = I_t + J_t$$

where

$$I_t := \int_0^t (\sigma(s, X^N_s) - \sigma(s, X^{N-1}_s)) dW_s$$

and

$$J_t := \int_0^t (b(s, X^N_s) - b(s, X^{N-1}_s)) ds$$

By applying Doob's L^2 -inequality to the nonnegative submartingale $(\|I_t\|^2)_{t \geq 0}$, using Lemma 26.7 and (26.7), we obtain

$$\begin{aligned} \mathbf{E} \left[\sup_{s \leq t} \|I_s\|^2 \right] &\leq 4\mathbf{E} [\|I_t\|^2] \\ &= 4\mathbf{E} \left[\int_0^t \|\sigma(s, X^N_s) - \sigma(s, X^{N-1}_s)\|^2 ds \right] \\ &\leq 4K^2 \int_0^t \mathbf{E} [\|X^N_s - X^{N-1}_s\|^2] ds \end{aligned} \quad (26.12)$$

For J_t , by the Cauchy-Schwarz inequality, we get

$$\|J_t\|^2 \leq t \int_0^t \|b(s, X_s^N) - b(s, X_s^{N-1})\|^2 ds$$

Hence

$$\begin{aligned} \mathbf{E} \left[\sup_{s \leq t} \|J_s\|^2 \right] &\leq t \mathbf{E} \left[\int_0^t \|b(s, X_s^N) - b(s, X_s^{N-1})\|^2 ds \right] \\ &\leq t K^2 \int_0^t \mathbf{E} \left[\|X_s^N - X_s^{N-1}\|^2 \right] ds \end{aligned} \quad (26.13)$$

Defining

$$\Delta^N(t) := \mathbf{E} \left[\sup_{s \leq t} \|X_s^N - X_s^{N-1}\|^2 \right],$$

and $C := 2K^2(4+T) \vee 2(T+1)K^2(1 + \|x\|^2)$, we obtain (using the growth condition (26.8))

$$\Delta^{N+1}(t) \leq C \int_0^t \Delta^N(s) ds \quad \text{for } N \geq 1$$

and

$$\begin{aligned} \Delta^1(t) &\leq 2t \int_0^t \|b(s, x)\|^2 ds + 2 \int_0^t \|\sigma(s, x)\|^2 ds \\ &\leq 2(T+1)K^2(1 + \|x\|^2) \cdot t \leq Ct \end{aligned}$$

Inductively, we get $\Delta^N(t) \leq \frac{(Ct)^N}{N!}$. Thus, by Markov's inequality,

$$\begin{aligned} \sum_{N=1}^{\infty} \mathbf{P} \left[\sup_{s \leq t} \|X_s^N - X_s^{N-1}\|^2 > 2^{-N} \right] &\leq \sum_{N=1}^{\infty} 2^N \Delta^N(t) \\ &\leq \sum_{N=1}^{\infty} \frac{(2Ct)^N}{N!} \leq e^{2Ct} < \infty \end{aligned}$$

Using the Borel-Cantelli lemma, we infer $\sup_{s \leq t} \|X_s^N - X_s^{N-1}\|^2 \xrightarrow{N \rightarrow \infty} 0$ a.s. Hence a.s. $(X^N)_{N \in \mathbb{N}}$ is a Cauchy sequence in the Banach space $(C([0, T]), \|\cdot\|_{\infty})$. Therefore, X^N converges a.s. uniformly to some X . As uniform convergence implies convergence of the integrals, X is a strong solution of (26.2).

Markov property. The strong Markov property follows from the strong Markov property of the Brownian motion that drives the SDE.

We have already seen some important examples of this theorem. Many interesting problems, however, lead to stochastic differential equations with coefficients that are not Lipschitz continuous. In the one-dimensional case, using special comparison methods, one can show that it is sufficient that σ is Hölder-continuous of order $\frac{1}{2}$ in the space variable.

Theorem 26.10 (Yamada-Watanabe) Consider the one-dimensional situation where $m = n = 1$. Assume that there exist $K < \infty$ and $\alpha \in [\frac{1}{2}, 1]$ such that, for all $t \geq 0$ and $x, x' \in \mathbb{R}$, we have

$$|b(t, x) - b(t, x')| \leq K|x - x'| \quad \text{and} \quad |\sigma(t, x) - \sigma(t, x')| \leq |x - x'|^\alpha.$$

Then, for every $X_0 \in \mathbb{R}$, the SDE (26.1) has a unique strong solution X and X is a strong Markov process.

Proof See [172] or [85, Proposition 5.2.13] and [49, Theorem 5.3.11] for existence and uniqueness. The strong Markov property follows from Theorem 26.26.

Example 26.11 Consider the one-dimensional SDE

$$dX_t = \sqrt{\gamma X_t^+} dW_t + a(b - X_t^+) dt \quad (26.14)$$

with initial point $X_0 = x \geq 0$, where $\gamma > 0$ and $a, b \geq 0$ are parameters. The conditions of Theorem 26.10 are fulfilled with $\alpha = \frac{1}{2}$ and $K = \sqrt{\gamma} + a$. Obviously, the unique strong solution X remains nonnegative if $X_0 \geq 0$. (In fact, it can be shown that $X_t > 0$ for all $t > 0$ if $2ab/\gamma \geq 1$, and that X_t hits zero arbitrarily often with probability 1 if $2ab/\gamma < 1$. See, e.g., [78, Example IV.8.2, p. 237]. Compare Example 26.16. See Figs. 26.1 and 26.2 for computer simulations.)




Fig. 26.1 Cox-Ingersoll-Ross diffusion with parameters $\gamma = 1, b = 1$ and $a = 0.3$. The path hits zero again and again since $2ab/\gamma = 0.6 < 1$




Fig. 26.2 Cox-Ingersoll-Ross diffusion with parameters $\gamma = 1, b = 1$ and $a = 2$. The path never hits zero since $2ab/\gamma = 4 \geq 1$

Depending on the context, this process is sometimes called Feller's branching diffusion with immigration or the Cox-Ingersoll-Ross model for the time evolution of interest rates.

For the case $a = b = 0$, use the Itô formula to compute that

$$e^{-\lambda X_t} - e^{-\lambda x} - \gamma \frac{\lambda^2}{2} \int_0^t e^{-\lambda X_s} X_s ds = \lambda \int_0^t e^{-\lambda X_s} \sqrt{\gamma X_s} dW_s$$

is a martingale. Take expectations for the Laplace transform $\varphi(t, \lambda, x) = \mathbf{E}_x [e^{-\lambda X_t}]$ to get the differential equation

$$\frac{d}{dt} \varphi(t, \lambda, x) = \gamma \frac{\lambda^2}{2} \mathbf{E} [X_t e^{-\lambda X_t}] = -\frac{\gamma \lambda^2}{2} \frac{d}{d\lambda} \varphi(t, \lambda, x)$$

With initial value $\varphi(0, \lambda, x) = e^{-\lambda x}$, the unique solution is

$$\varphi(t, \lambda, x) = \exp\left(-\frac{\lambda}{(\gamma/2)\lambda t + 1}x\right)$$

However (for $\gamma = 2$), this is exactly the Laplace transform of the transition probabilities of the Markov process that we defined in Theorem 21.48 and that in Lindvall's theorem (Theorem 21.51) we encountered as the limit of rescaled Galton-Watson branching processes.

Exercise 26.1.1 Let $a, b \in \mathbb{R}$. Show that the stochastic differential equation

$$dX_t = \frac{b - X_t}{1 - t} dt + dW_t$$

with initial value $X_0 = a$ has a unique strong solution for $t \in [0, 1)$ and that $X_1 := \lim_{t \uparrow 1} X_t = b$ almost surely. Furthermore, show that the process $Y = (X_t - a - t(b - a))_{t \in [0, 1]}$ can be described by the Itô integral

$$Y_t = (1 - t) \int_0^t (1 - s)^{-1} dW_s, \quad t \in [0, 1)$$

and is hence a Brownian bridge (compare Exercise 21.5.3).

5.26.2 26.2 Weak Solutions and the Martingale Problem

In the last section, we studied strong solutions of the stochastic differential equation

$$dX_t = \sigma(t, X_t) dW_t + b(t, X_t) dt. \quad (26.15)$$

A strong solution is a solution where any path of the Brownian motion W gets mapped onto a path of the solution X . In this section, we will study the notion of a weak solution where additional information (or additional noise) can be used to construct the solution.

Definition 26.12 (Weak solution of an SDE) A weak solution of (26.15) with initial distribution $\mu \in \mathcal{M}_1(\mathbb{R}^n)$ is a triple

$$L = ((X, W), (\Omega, \mathcal{F}, \mathbf{P}), \mathbb{F}),$$

where

- $(\Omega, \mathcal{F}, \mathbf{P})$ is a probability space,
- $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ is a filtration on $(\Omega, \mathcal{F}, \mathbf{P})$ that satisfies the usual conditions,
- W is a Brownian motion on $(\Omega, \mathcal{F}, \mathbf{P})$ and is a martingale with respect to $\mathbb{F}$,
- X is continuous and adapted (hence progressively measurable),
- $\mathbf{P} \circ (X_0)^{-1} = \mu$, and
- (X, W) satisfies

$$X_t = X_0 + \int_0^t \sigma(s, X_s) dW_s + \int_0^t b(s, X_s) ds \quad \mathbf{P}\text{-a.s.} \quad (26.16)$$

A weak solution L is called (weakly) unique if, for any further solution L' with initial distribution μ , we have $\mathbf{P}' \circ (X')^{-1} = \mathbf{P} \circ X^{-1}$.

Remark 26.13 Clearly, a weak solution of an SDE is a generalized n -dimensional diffusion. If the coefficients σ and b do not depend on t , then the solution is an n -dimensional diffusion.

Remark 26.14 Clearly, every strong solution of (26.15) is a weak solution. The converse is false, as the following example shows.

Example 26.15 Consider the SDE (with initial value $X_0 = 0$)

$$dX_t = \text{sign}(X_t) dW_t, \quad (26.17)$$

where $\text{sign} = \mathbb{1}_{(0,\infty)} - \mathbb{1}_{(-\infty,0)}$ is the sign function. Then

$$X_t = X_0 + \int_0^t \text{sign}(X_s) dW_s \quad \text{for all } t \geq 0 \quad (26.18)$$

if and only if

$$W_t = \int_0^t dW_s = \int_0^t \text{sign}(X_s) dX_s \quad \text{for all } t \geq 0 \quad (26.19)$$

A weak solution of (26.17) is obtained as follows. Let X be a Brownian motion on a probability space $(\Omega, \mathcal{F}, \mathbf{P})$ and $\mathbb{F} = \sigma(X)$. If we define W by (26.19), then W is a continuous $\mathbb{F}$ -martingale with square variation

$$\langle W \rangle_t = \int_0^t (\text{sign}(X_s))^2 ds = t$$

Thus, by Lévy's characterization (Theorem 25.28), W is a Brownian motion. Hence $((X, W), (\Omega, \mathcal{F}, \mathbf{P}), \mathbb{F})$ is a weak solution of (26.17).

In order to show that a strong solution does not exist, take any weak solution and show that X is not adapted to $\sigma(W)$. Since, by (26.18), X is a continuous martingale with square variation $\langle X \rangle_t = t$, X is a Brownian motion.

Let $F_n \in C^2(\mathbb{R})$ be a convex even function with derivatives F'_n and F''_n such that

$$\sup_{x \in \mathbb{R}} |F_n(x) - |x|| \xrightarrow{n \rightarrow \infty} 0$$

$|F'_n(x)| \leq 1$ for all $x \in \mathbb{R}$ and $F'_n(x) = \text{sign}(x)$ for $|x| > \frac{1}{n}$. In particular, we have

$$\int_0^t (F'_n(X_s) - \text{sign}(X_s))^2 ds \xrightarrow{n \rightarrow \infty} 0 \quad \text{a.s.}$$

and thus

$$\int_0^t F'_n(X_s) dX_s \xrightarrow{n \rightarrow \infty} \int_0^t \text{sign}(X_s) dX_s \quad \text{in } L^2 \quad (26.20)$$

By passing to a subsequence, if necessary, we may assume that almost sure convergence holds in (26.20).

Since F''_n is even, we have

$$\begin{aligned} W_t &= \int_0^t \text{sign}(X_s) dX_s = \lim_{n \rightarrow \infty} \int_0^t F'_n(X_s) dX_s \\ &= \lim_{n \rightarrow \infty} \left(F_n(X_t) - F_n(0) - \frac{1}{2} \int_0^t F''_n(X_s) ds \right) \\ &= |X_t| - \lim_{n \rightarrow \infty} \frac{1}{2} \int_0^t F''_n(|X_s|) ds \end{aligned}$$

As the right-hand side depends only on $|X_s|$, $s \in [0, t]$, W is adapted to $\mathbb{G} := (\sigma(|X_s| : s \in [0, t]))$. Hence $\sigma(W) \subset \mathbb{G} \subsetneq \sigma(X)$, and thus X is not adapted to $\sigma(W)$.

Example 26.16 Let $B = (B^1, \dots, B^n)$ be an n -dimensional Brownian motion started at $y \in \mathbb{R}^n$. Let $x := \|y\|^2$, $X_t := \|B_t\|^2 = (B_t^1)^2 + \dots + (B_t^n)^2$ and

$$W_t := \sum_{i=1}^n \int_0^t \frac{1}{\sqrt{X_s}} B_s^i dB_s^i$$

Then W is a continuous local martingale with $\langle W \rangle_t = t$ for every $t \geq 0$ and

$$X_t = x + nt + \int_0^t \sqrt{X_s} dW_s$$

That is, (X, W) is a weak solution of the SDE $dX_t = \sqrt{2X_t} dW_t + n dt$. X is called an n -dimensional Bessel process. By Theorem 25.42, B (and thus X) hits the origin for some $t > 0$ if and only if $n = 1$. Clearly, we can define X also for noninteger $n \geq 0$. One can show that X hits zero if and only if $n \leq 1$. Compare Example 26.11. $\diamond$

For the connection between existence and uniqueness of weak solutions and strong solutions, we only quote here the theorem of Yamada and Watanabe.

Definition 26.17 (Pathwise uniqueness) A solution of the SDE (26.15) with initial distribution μ is said to be pathwise unique if, for every $\mu \in \mathcal{M}_1(\mathbb{R}^n)$ and for any two weak solutions (X, W) and (X', W') on the same space $(\Omega, \mathcal{F}, \mathbf{P})$ with the same filtration $\mathbb{F}$, we have $\mathbf{P}[X_t = X'_t \text{ for all } t \geq 0] = 1$.

Theorem 26.18 (Yamada and Watanabe) The following are equivalent.

- (i) The SDE (26.15) has a unique strong solution.
- (ii) For any $\mu \in \mathcal{M}_1(\mathbb{R}^n)$, (26.15) has a weak solution, and pathwise uniqueness holds.

If (i) and (ii) hold, then the solution is weakly unique.

Proof See [172], [147, pp. 151ff] or [78, pp. 163ff].

Example 26.19 Let X be a weak solution of (26.17). Then $-X$ is also a weak solution; that is, pathwise uniqueness does not hold (although it can be shown that the solution is weakly unique; see Theorem 26.25).

Consider the one-dimensional case $m = n = 1$. If X is a solution (strong or weak) of (26.15), then

$$M_t := X_t - \int_0^t b(s, X_s) ds$$

is a continuous local martingale with square variation

$$\langle M \rangle_t = \int_0^t \sigma^2(s, X_s) ds$$

We will see that this characterizes a weak solution of (26.15) (under some mild growth conditions on σ and b).

Now assume that, for all $t \geq 0$ and $x \in \mathbb{R}^n$, the $n \times n$ matrix $a(t, x)$ is symmetric and nonnegative definite, and let $(t, x) \mapsto a(t, x)$ be measurable.

Definition 26.20 An n -dimensional continuous process X is called a solution of the local martingale problem for a and b with initial condition $\mu \in \mathcal{M}_1(\mathbb{R}^n)$ (briefly, LMP(a, b, μ)) if $\mathbf{P} \circ X_0^{-1} = \mu$ and if, for every $i = 1, \dots, n$,

$$M_t^i := X_t^i - \int_0^t b_i(s, X_s) ds, \quad t \geq 0$$

is a continuous local martingale with quadratic covariation

$$\langle M^i, M^j \rangle_t = \int_0^t a_{ij}(s, X_s) ds \quad \text{for all } t \geq 0, i, j = 1, \dots, n$$

We say that the solution of $\text{LMP}(a, b, \mu)$ is unique if, for any two solutions X and X' , we have $\mathbf{P} \circ X^{-1} = \mathbf{P} \circ (X')^{-1}$.

Denote by σ^T the transposed matrix of σ . Clearly, $a = \sigma\sigma^T$ is a nonnegative semidefinite symmetric $n \times n$ matrix.

Theorem 26.21 X is a solution of $\text{LMP}(\sigma\sigma^T, b, \mu)$ if and only if (on a suitable extension of the probability space) there exists a Brownian motion W such that (X, W) is a weak solution of (26.15).

In particular, there exists a unique weak solution of the SDE (26.15) with initial distribution μ if $\text{LMP}(\sigma\sigma^T, b, \mu)$ is uniquely solvable.

Proof We show the statement only for the case $m = n = 1$. The general case needs some consideration on the roots of nonnegative semidefinite symmetric matrices, which, however, do not yield any further insight into the stochastics of the problem. For this we refer to [85, Proposition 5.4.6].

"." If (X, W) is a weak solution, then, by Corollary 25.19, X solves the local martingale problem.

" $\implies$ " Let X be a solution of $\text{LMP}(\sigma^2, b, \mu)$. By Theorem 25.29, on an extension of the probability space there exists a Brownian motion $\tilde{W}$ such that $M_t = \int_0^t |\sigma(s, X_s)| d\tilde{W}_s$. If we define

$$W_t := \int_0^t \text{sign}(\sigma(s, X_s)) d\tilde{W}_s$$

then $M_t = \int_0^t \sigma(s, X_s) dW_s$ and hence (X, W) is a weak solution of (26.15).

In some sense, a local martingale problem is a very natural way of writing a stochastic differential equation; that is:

X locally has derivative (drift) b and additionally has random normally distributed fluctuations of size σ .

Here, a concrete Brownian motion does not appear. In fact, in most problems its occurrence is rather artificial. Just as Markov chains are described by their transition probabilities and not by a concrete realization of the random transitions (as in Theorem 17.17), many continuous (space and time) processes are most naturally described by the drift and the size of the fluctuations but not by the concrete realization of the random fluctuations.

From a technical point of view, the formulation of a stochastic differential equation as a local martingale problem is very convenient since it makes SDEs accessible to techniques such as martingale inequalities and approximation theorems that can be used to establish existence and uniqueness of solutions. Here we simply quote two important results.

Theorem 26.22 (Existence of solutions) Let $(t, x) \mapsto b(t, x)$ and $(t, x) \mapsto a(t, x)$ be continuous and bounded. Then, for every $\mu \in \mathcal{M}_1(\mathbb{R}^n)$, there exists a solution X of the $\text{LMP}(a, b, \mu)$.

Proof See [147, Theorem V.23.5].

Definition 26.23 The $\text{LMP}(a, b)$ is said to be well-posed if, for every $x \in \mathbb{R}^n$, there exists a unique solution X of $\text{LMP}(a, b, \delta_x)$.

Remark 26.24 If σ and b satisfy the Lipschitz conditions of Theorem 26.8, then the $\text{LMP}(\sigma\sigma^T, b)$ is well-posed. This follows by Theorem 26.8, Theorem 26.18 and Theorem 26.21.

In the following, we assume

$$(t, x) \mapsto \sigma(t, x) \text{ resp. } (t, x) \mapsto a(t, x) \text{ is bounded on compact sets.} \quad (26.21)$$

This condition ensures the equivalence of the local martingale problems to the somewhat more common martingale problem (see [85, Proposition 5.4.11]).

Theorem 26.25 (Uniqueness in the martingale problem) Assume (26.21) and that, for any $x \in \mathbb{R}^n$, there exists a solution X^x of LMP (a, b, δ_x) . The distribution of X^x will be denoted by $\mathbf{P}_x := \mathbf{P} \circ (X^x)^{-1}$.

Assume that, for any two solutions X^x and Y^x of LMP (a, b, δ_x) , we have

$$\mathbf{P} \circ (X_T^x)^{-1} = \mathbf{P} \circ (Y_T^x)^{-1} \quad \text{for any } T \geq 0. \quad (26.22)$$

Then LMP (a, b) is well-posed, and the canonical process X is a strong Markov process with respect to $(\mathbf{P}_x, x \in \mathbb{R}^n)$. If $a = \sigma \sigma^T$, then under $\mathbf{P}_x$, the process X is the unique weak solution of the SDE (26.15).

Proof See [49, Theorem 4.4.2 and Problem 49] and [85, Proposition 5.4.11].

A fundamental strength of this theorem is that we do not need to check the uniqueness of the whole process but only have to check in (26.22) the onedimensional marginal distributions. We will use this in Section 26.3 in some examples.

The existence of solutions of a stochastic differential equation (or equivalently of a local martingale problem) is often easier to show than the uniqueness of solutions. We know already that Lipschitz conditions for the coefficients b and σ (not $\sigma \sigma^T$!) ensure uniqueness (Theorem 26.8 and Theorem 26.18), as here strong uniqueness of the solution holds.

At first glance, it might seem confusing that random fluctuations have a stabilising effect on the solution. That is, there are deterministic differential equations whose solution is unique only after adding random noise terms. For example, consider the following equation:

$$dX_t = \text{sign}(X_t) |X_t|^{1/3} dt + \sigma dW_t, \quad X_0 = 0 \quad (26.23)$$

If $\sigma = 0$, then the deterministic differential equation has a continuum of solutions that can be parameterized by $v \in \{-1, +1\}$ and $T \geq 0$, namely $X_t = v2\sqrt{2}(t - T)^{3/2} \mathbb{1}_{\{t > T\}}$. If $\sigma > 0$, then the noise eliminates the instability of (26.23) at $x = 0$. We quote the following theorem for the time-independent case from [147, Theorem V.24.1] (see also [161, Chapter 10]).

Theorem 26.26 (Stroock-Varadhan) Let $a_{ij} : \mathbb{R}^n \rightarrow \mathbb{R}$ be continuous and let $b_i : \mathbb{R}^n \rightarrow \mathbb{R}$ be measurable for $i, j = 1, \dots, n$. Assume

- (i) $a(x) = (a_{ij}(x))$ is symmetric and strictly positive definite for every $x \in \mathbb{R}^n$,
- (ii) there exists a $C < \infty$ such that, for all $x \in \mathbb{R}^n$ and $i, j = 1, \dots, n$, we have

$$|a_{ij}(x)| \leq C(1 + \|x\|^2) \quad \text{and} \quad |b_i(x)| \leq C(1 + \|x\|).$$

Then the LMP (a, b) is well-posed and the SDE (26.15) has a unique strong solution that is a strong Markov process. The solution X has the Feller property: For every $t > 0$ and every bounded measurable $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the map $x \mapsto \mathbf{E}_x[f(X_t)]$ is continuous.

We will present explicit examples in Section 26.3. Here we just remark that we have developed a particular method in order to construct Markov processes, namely as the solution of a stochastic differential equation or of a local martingale problem. In the framework of models in discrete time, in Section 17.2 and especially in Exercise 17.2.1, we characterized certain Markov chains as solutions of martingale problems. In order for drift and square variation to be sufficient for uniqueness of the Markov chain described by the martingale problem, it was essential that, for any step of the chain, we only allowed three possibilities. Here, however, the decisive restriction is the continuity of the processes.

Exercise 26.2.1 Consider the time-homogeneous one-dimensional case ($m = n = 1$). Let σ and b be such that, for every $X_0 \in \mathbb{R}$, there exists a unique weak solution of

$$dX_t = \sigma(X_t) dW_t + b(X_t) dt$$

that is a strong Markov process. Further, assume that there exists an $x_0 \in \mathbb{R}$ with

$$C := \int_{-\infty}^{\infty} \frac{1}{\sigma^2(x)} \exp \left(\int_{x_0}^x \frac{2b(r)}{\sigma^2(r)} dr \right) dx < \infty$$

(i) Show that the measure $\pi \in \mathcal{M}_1(\mathbb{R})$ with density

$$\frac{\pi(dx)}{dx} = C^{-1} \frac{1}{\sigma^2(x)} \exp \left(\int_{x_0}^x \frac{2b(r)}{\sigma^2(r)} dr \right)$$

is an invariant distribution for X .

(ii) For which values of b does the Ornstein-Uhlenbeck process $dX_t = \sigma dW_t + bX_t dt$ have an invariant distribution? Determine this distribution and compare the result with what could be expected by an explicit computation using the representation in (26.3).

(iii) Compute the invariant distribution of the Cox-Ingersoll-Ross SDE (26.14) (i.e., Feller's branching diffusion).

(iv) Let $\gamma, c > 0$ and $\theta \in (0, 1)$. Show that the invariant distribution of the solution X of the SDE on $[0, 1]$,

$$dX_t = \sqrt{\gamma X_t(1 - X_t)} dW_t + c(\theta - X_t) dt$$

is the Beta distribution $\beta_{2c\gamma/\theta, 2c\gamma/(1-\theta)}$.

Exercise 26.2.2 Let $\gamma > 0$. Let X^1 and X^2 be solutions of $dX_t^i = \sqrt{\gamma X_t^i} dW_t^i$, where W^1 and W^2 are two independent Brownian motions with initial values $X_0^1 = x_0^1 > 0$ and $X_0^2 = x_0^2 > 0$. Show that $Z := X^1 + X^2$ is a weak solution of $Z_0 = 0$ and $dZ_t = \sqrt{\gamma Z_t} dW_t$.

5.26.3 26.3 Weak Uniqueness via Duality

The Stroock-Varadhan theorem provides a strong criterion for existence and uniqueness of solutions of stochastic differential equations. However, in many cases, the condition of locally uniform ellipticity of a (Condition (i) in Theorem 26.26) is not fulfilled. This is the case, in particular, if the solutions are defined only on subsets of $\mathbb{R}^n$.

Here we will study a powerful tool that in many special cases can yield weak uniqueness of solutions.

Definition 26.27 (Duality) Let $X = (X^x, x \in E)$ and $Y = (Y^y, y \in E')$ be families of stochastic processes with values in the spaces E and E' , respectively, and such that $X_0^x = x$ a.s. and $Y_0^y = y$ a.s. for all $x \in E$ and $y \in E'$. We say that X and Y are dual to each other with duality function $H : E \times E' \rightarrow \mathbb{C}$ if, for all $x \in E, y \in E'$ and $t \geq 0$, the expectations $\mathbf{E}[H(X_t^x, y)]$ and $\mathbf{E}[H(x, Y_t^y)]$ exist and are equal:

$$\mathbf{E}[H(X_t^x, y)] = \mathbf{E}[H(x, Y_t^y)].$$

In the following, we assume that $\sigma_{ij} : \mathbb{R}^n \rightarrow \mathbb{R}$ and $b_i : \mathbb{R}^n \rightarrow \mathbb{R}$ are bounded on compact sets for all $i = 1, \dots, n, j = 1, \dots, m$. Consider the time-homogeneous stochastic differential equation

$$dX_t = \sigma(X_t) dW_t + b(X_t) dt \quad (26.24)$$

Theorem 26.28 (Uniqueness via duality) Assume that, for every $x \in \mathbb{R}^n$, there exists a solution of the local martingale problem for $(\sigma \sigma^T, b, \delta_x)$. Further, assume that there exists a family $(Y^y, y \in E')$ of Markov processes with values in the measurable space $(E', \mathcal{E}')$ and a measurable map $H : \mathbb{R}^n \times E' \rightarrow \mathbb{C}$ such that, for every $y \in E', x \in \mathbb{R}^n$ and $t \geq 0$, the expectation $\mathbf{E}[H(x, Y_t^y)]$ exists and is finite. Further, let $(H(\cdot, y), y \in E')$ be a separating class of functions for $\mathcal{M}_1(\mathbb{R}^n)$ (see Definition 13.9).

For every $x \in \mathbb{R}^n$ and every solution X^x of LMP $(\sigma\sigma^T, b, \delta_x)$, assume that the duality equation holds:

$$\mathbf{E}[H(X_t^x, y)] = \mathbf{E}[H(x, Y_t^y)] \quad \text{for all } y \in \mathbf{E}', t \geq 0. \quad (26.25)$$

Then the local martingale problem of $(\sigma\sigma^T, b)$ is well-posed and hence (26.24) has a unique weak solution that is a strong Markov process.

Proof By Theorem 26.25, it is enough to check that, for every $x \in \mathbb{R}^n$, every solution X^x of LMP $(\sigma\sigma^T, b, \delta_x)$ and every $t \geq 0$, the distribution $\mathbf{P} \circ (X_t^x)^{-1}$ is unique. Since $(H(\cdot, y), y \in E')$ is a separating class of functions, this follows from (26.16).

Example 26.29 (Wright-Fisher diffusion) Consider the Wright-Fisher SDE

$$dX_t = \mathbb{1}_{[0,1]}(X_t) \sqrt{\gamma X_t(1-X_t)} dW_t, \quad (26.26)$$

where $\gamma > 0$ is a parameter. See Fig. 26.3 for a computer simulation. By Theorem 26.22, for every $x \in \mathbb{R}$, there exists a weak solution $(\tilde{X}, W)$ of (26.26). $\tilde{X}$ is a continuous local martingale with square variation

$$\langle \tilde{X} \rangle_t = \int_0^t \gamma \tilde{X}_s (1 - \tilde{X}_s) \mathbb{1}_{[0,1]}(\tilde{X}_s) ds$$

Let $\tau := \inf \{t > 0 : \tilde{X}_t \notin [0, 1]\}$ and let $X := \tilde{X}^\tau$ be the process stopped at τ . Then X is a continuous bounded martingale with

$$\langle X \rangle_t = \int_0^t \gamma X_s (1 - X_s) \mathbb{1}_{[0,1]}(X_s) ds$$

Hence, (X, W) is a solution of (26.26). By construction, $X_t \in [0, 1]$ for all $t \geq 0$ if $X_0 = \tilde{X}_0 \in [0, 1]$.


Let $\tau' := \inf \{t > 0 : \tilde{X}_t \in [0, 1]\}$. If $\tilde{X}_0 \notin [0, 1]$, then $\tau' > 0$ since $\tilde{X}$ is continuous. Since $\tilde{X}^{\tau'}$ is a continuous local martingale with $\langle \tilde{X}^{\tau'} \rangle \equiv 0$, we have $\tilde{X}_t^{\tau'} = \tilde{X}_0$ for all $t \geq 0$. However, this implies $\tilde{X}_t = \tilde{X}_0$ for all $t < \tau'$. Again, by continuity of $\tilde{X}$, we get $\tau' = \infty$ and $\tilde{X}_t = \tilde{X}_0$ for all $t \geq 0$.

Hence, it is enough to show uniqueness of the solution for $\tilde{X}_0 = x \in [0, 1]$. To this end, let $Y = (Y_t)_{t \geq 0}$ be the Markov process on $\mathbb{N}$ with Q -matrix

$$q(m, n) = \begin{cases} \gamma \binom{m}{2}, & \text{if } n = m - 1, \\ -\gamma \binom{m}{2}, & \text{if } n = m, \\ 0, & \text{else.} \end{cases}$$

We show duality of X and Y with respect to $H(x, n) = x^n$:

$$\mathbf{E}_x[X_t^n] = \mathbf{E}_n[x^{N_t}] \quad \text{for all } t \geq 0, x \in [0, 1], n \in \mathbb{N}. \quad (26.27)$$

Fig. 26.3 Simulation of a Wright-Fisher diffusion with parameter $\gamma = 1$

Define $m^{x,n}(t) = \mathbf{E}_x[X_t^n]$ and $g^{x,n}(t) = \mathbf{E}_n[x^{N_t}]$. By the Itô formula,

$$X_t^n - x^n - \int_0^t \gamma \binom{n}{2} X_s^{n-1} (1 - X_s) ds = \int_0^t n X_s^{n-1} \sqrt{\gamma X_s (1 - X_s)} dW_s$$

is a martingale.

Taking expectations, we obtain the following recursive equations for the moments of X :

$$\begin{aligned} m^{x,1}(t) &= x \\ m^{x,n}(t) &= x^n + \gamma \binom{n}{2} \int_0^t (m^{x,n-1}(s) - m^{x,n}(s)) ds \quad \text{for } n \geq 2 \end{aligned} \quad (26.28)$$

Clearly, this system of linear differential equations can be uniquely solved recursively in n . Due to the Markov property of Y , for $h > 0$ and $t \geq 0$, we have

$$\begin{aligned} g^{x,n}(t+h) &= \mathbf{E}_n[x^{Y_{t+h}}] = \mathbf{E}_n[\mathbf{E}_{Y_h}[x^{Y_t}]] \\ &= \sum_{m=1}^n \mathbf{P}_n[Y_h = m] \mathbf{E}_m[x^{Y_t}] \\ &= \sum_{m=1}^n \mathbf{P}_n[Y_h = m] g^{x,m}(t) \end{aligned}$$

This implies

$$\begin{aligned} \frac{d}{dt} g^{x,n}(t) &= \lim_{h \downarrow 0} h^{-1} [g^{x,n}(t+h) - g^{x,n}(t)] \\ &= \lim_{h \downarrow 0} h^{-1} \sum_{m=1}^n \mathbf{P}_n[Y_h = m] (g^{x,m}(t) - g^{x,n}(t)) \\ &= \sum_{m=1}^n q(n, m) g^{x,m}(t) \\ &= \gamma \binom{n}{2} (g^{x,n-1}(t) - g^{x,n}(t)). \end{aligned} \quad (26.29)$$

Evidently, $g^{x,1}(t) = x$ for all $x \in [0, 1]$ and $t \geq 0$ and $g^{x,n}(0) = x^n$. That is, $g^{x,n}$ solves (26.28), and thus (26.27) holds.

By Theorem 15.4, the family $(H(\cdot, n), n \in \mathbb{N}) \subset C([0, 1])$ is separating for $\mathcal{M}_1([0, 1])$; hence the conditions of Theorem 26.28 are fulfilled. Therefore, X is the unique weak solution of (26.26) and is a strong Markov process.

Remark 26.30 The martingale problem for the Wright-Fisher diffusion is almost identical to the martingale problem for the Moran model (see Example 17.22) $M^N = (M_n^N)_{n \in \mathbb{N}_0}$ with population size N : M^N is a martingale with values in the set $\{0, 1/N, \dots, (N-1)/N, 1\}$ and with square variation process

$$\langle M^N \rangle_n = \frac{2}{N^2} \sum_{k=0}^{n-1} M_k^N (1 - M_k^N)$$

At each step, M^N can either stay put or increase or decrease by $1/N$. In Exercise 17.2.1, we saw that this determines the process M^N uniquely. Similarly as in Theorem 21.51 for branching processes, it can be shown that the time-rescaled Moran processes $\tilde{M}_t^N = M_{\lfloor N^2 t \rfloor}^N$ converge to the Wright-Fisher diffusion with $\gamma = 2$. The Wright-Fisher diffusion thus occurs as the limiting model of a genealogical model and describes the gene frequency (that is, the fraction) of a certain allele in a population that fluctuates randomly due to resampling.

Example 26.31 (Feller's branching diffusion) Let $(Z_n^N)_{n \in \mathbb{N}_0}$ be a Galton-Watson branching process with critical geometric offspring distribution $p_k = 2^{-k-1}$, $k \in \mathbb{N}_0$ and $Z_0^N = N$ for any $N \in \mathbb{N}$. Then Z^N is a discrete martingale and we have

$$\mathbf{E} \left[(Z_n^N - Z_{n-1}^N)^2 \mid Z_{n-1}^N \right] = Z_{n-1}^N \left(\sum_{k=0}^{\infty} p_k k^2 - 1 \right) = 2Z_{n-1}^N$$

Hence Z^N has square variation

$$\langle Z^N \rangle_n = \sum_{k=0}^{n-1} 2Z_k^N$$

Define the linearly interpolated version

$$Z_t^N := (t - N^{-1} \lfloor tN \rfloor) (Z_{\lfloor tN \rfloor + 1}^N - Z_{\lfloor tN \rfloor}^N) + \frac{1}{n} Z_{\lfloor tN \rfloor}^N$$




Fig. 26.4 Simulation of Feller's branching diffusion with parameter $\gamma = 1$

of $N^{-1}Z_{\lfloor tN \rfloor}^N$. By Lindvall's theorem (Theorem 21.51), there is a continuous Markov process Z such that $Z^N \xrightarrow{N \rightarrow \infty} Z$ in distribution. See Fig. 26.4 for a computer simulation of Z . Since it can be shown that the moments also converge, we have that Z is a continuous martingale with square variation

$$\langle Z \rangle_t = \int_0^t 2Z_s ds$$

In fact, in Example 26.11, we have already shown that Z is the unique solution of the SDE

$$dZ_t = \sqrt{2Z_t} dW_t \quad (26.30)$$

with initial value $Z_0 = 1$. There we also showed that Z is dual to $Y_t^y = \left(\frac{t\gamma}{2} + \frac{1}{y}\right)^{-1}$ with $H(x, y) = e^{-xy}$. This implies uniqueness of the solution of (26.30) and the strong Markov property of Z . $\square$

It could be objected that in Examples 26.29 and 26.31, we considered only onedimensional problems for which the Yamada-Watanabe theorem (Theorem 26.10) yields uniqueness (indeed of a strong solution) anyway. The full strength of the method of duality is displayed only in higher-dimensional problems. As an example, we consider an extension of Example 26.29.

Example 26.32 (Interacting Wright-Fisher diffusions) The Wright-Fisher diffusion from Example 26.29 describes the fluctuations of the gene frequency of an allele in one large population. Now we consider more populations, which live at the points $i \in S := \{1, \dots, N\}$ and interact with each other by a migration that is quantified by migration rates $r(i, j) \geq 0$. As a model for the gene frequencies $X_t(i)$ at site i at time t we use the following N -dimensional SDE for $X = (X(1), \dots, X(N))$:

$$dX_t(i) = \sqrt{\gamma X_t(i)(1 - X_t(i))} dW_t^i + \sum_{j=1}^N r(i, j) (X_t(j) - X_t(i)) dt \quad (26.31)$$

Here $W = (W^1, \dots, W^N)$ is an N -dimensional Brownian motion. By Theorem 26.22, this SDE has weak solutions; however, none of our general criteria for weak uniqueness apply. We will thus show weak uniqueness by virtue of duality.

As in Example 26.29, it is not hard to show that solutions of (26.31), started at $X_0 = x \in E := [0, 1]^S$, remain in $[0, 1]^S$. The diagonal terms $r(i, i)$ do not appear in (26.31). We use our freedom and define these terms as $r(i, i) = -\sum_{j \neq i} r(i, j)$. Let $Y = (Y_t)_{t \geq 0}$ be the Markov process on $E' := (\mathbb{N}_0)^S$ with the following Q -matrix:

$$q(\varphi, \eta) = \begin{cases} \varphi(i)r(i, j), & \text{if } \eta = \varphi - \mathbb{1}_{\{i\}} + \mathbb{1}_{\{j\}} \text{ for} \\ \gamma \binom{\varphi(i)}{2}, & \text{some } i, j \in S, i \neq j, \\ \sum_{i \in S} \left(\varphi(i)r(i, i) - \gamma \binom{\varphi(i)}{2} \right), & \text{if } \eta = \varphi - \mathbb{1}_{\{i\}} \text{ for some } i \in S, \\ 0, & \text{if } \eta = \varphi, \\ \text{else.} & \end{cases}$$

Here $\varphi \in E'$ denotes a generic state with $\varphi(i)$ particles at site $i \in S$, and $\mathbb{1}_{\{i\}} \in E'$ denotes the state with exactly one particle at site i . The process Y describes a system of particles that independently with rate $r(i, j)$ jump from site i to site j . If there is more than one particle at the same site i , then any of the $\binom{\varphi(i)}{2}$ pairs of particles coalesce with the same rate γ to one particle. The common genealogical interpretation of this process is that (in reversed time) it describes the lines of descent of samples of $Y_0(i)$ individuals at each site $i \in S$. By migration, the lines change sites. If two individuals have the same common ancestor, then their lines coalesce. Clearly, for two particles to have the same ancestor at a given time, it is necessary but not sufficient for them to be at the same site.

For $x \in \mathbb{R}^n$ and $\varphi \in E'$, we denote $x^\varphi := \prod_{i \in S} x(i)^{\varphi(i)}$. We show that X and Y are dual to each other with the duality function $H(x, \varphi) = x^\varphi$:

$$\mathbf{E}_x [X_t^\varphi] = \mathbf{E}_\varphi [x^{Y_t}] \quad \text{for all } \varphi \in S^{\mathbb{N}_0}, x \in [0, 1]^S, t \geq 0 \quad (26.32)$$

Let $m^{x,\varphi}(t) := \mathbf{E}_x[X_t^\varphi]$ and $g^{x,\varphi}(t) := \mathbf{E}_\varphi[x^{Y_t}]$. Clearly, H has the derivatives $\partial_i H(\cdot, \varphi)(x) = \varphi(i)x^{\varphi - \mathbb{1}_{\{i\}}}$ and $\partial_i \partial_i H(\cdot, \varphi)(x) = 2 \binom{\varphi(i)}{2} x^{\varphi - 2\mathbb{1}_{\{i\}}}$.

By the Itô formula,

$$\begin{aligned} X_t^\varphi - X_0^\varphi - \int_0^t \sum_{i,j \in S} \varphi(i) r(i, j) (X_s(j) - X_s(i)) X_s^{\varphi - \mathbb{1}_{\{i\}}} ds \\ - \sum_{i \in S} \int_0^t \gamma \binom{\varphi(i)}{2} (X_s(i) (1 - X_s(i))) X_s^{\varphi - 2\mathbb{1}_{\{i\}}} ds \end{aligned}$$

is a martingale. Taking expectations, we get a system of linear integral equations

$$\begin{aligned} m^{x,0}(t) &= 1 \\ m^{x,\varphi}(t) &= x^\varphi + \int_0^t \sum_{i,j \in S} \varphi(i) r(i, j) (m^{x,\varphi + \mathbb{1}_{\{j\}} - \mathbb{1}_{\{i\}}}(s) - m^{x,\varphi}(s)) ds \\ &\quad + \int_0^t \gamma \sum_{i \in S} \binom{\varphi(i)}{2} (m^{x,\varphi - \mathbb{1}_{\{i\}}}(s) - m^{x,\varphi}(s)) ds \end{aligned} \quad (26.33)$$

This system of equations can be solved uniquely by induction on $n = \sum_{i \in I} \varphi(i)$. However, we do not intend to compute this solution explicitly. We show only that it coincides with $g^{x,\varphi}(t)$ by showing that g solves an equivalent system of differential equations.

For g as in (26.29), we obtain

$$\begin{aligned} \frac{d}{dt} g^{x,\varphi}(t) &= \sum_{\eta \in E'} q(\varphi, \eta) g^{x,\varphi}(t) \\ &= \sum_{i,j \in S} r(i, j) (g^{x,\varphi + \mathbb{1}_{\{j\}} - \mathbb{1}_{\{i\}}}(t) - g^{x,\varphi}(t)) \\ &\quad + \sum_{i \in S} \gamma \binom{\varphi(i)}{2} (g^{x,\varphi - \mathbb{1}_{\{i\}}}(t) - g^{x,\varphi}(t)) \end{aligned} \quad (26.34)$$

Together with the initial values $g^{x,0}(t) = 1$ and $g^{x,\varphi}(0) = x^\varphi$, the system (26.34) of differential equations is equivalent to (26.33). Hence the duality (26.32) holds, and thus the SDE (26.31) has a unique weak solution. (In fact, it can be shown that there exists a unique strong solution, even if S is countably infinite, as long as r then satisfies certain regularity conditions such as if it is the Q -matrix of a random walk on $S = \mathbb{Z}^d$; see [153].)

Exercise 26.3.1 (Extinction probability of Feller's branching diffusion) Let $\gamma > 0$ and let Z be the solution of $dZ_t := \sqrt{\gamma Z_t} dW_t$ with initial value $Z_0 = z > 0$. Use the duality to show

$$\mathbf{P}_z[Z_t = 0] = \exp\left(-\frac{2z}{\gamma t}\right) \quad (26.35)$$

Use Lemma 21.44 to compute the probability that a Galton-Watson branching process X with critical geometric offspring distribution and with $X_0 = N \in \mathbb{N}$ is extinct by time $n \in \mathbb{N}$. Compare the result with (26.35).

Part III

Problems

6 Typical tasks

(У Чистякова есть некоторые ответы и решения!!!! тренироваться много еще нужно!!)

6.1 Problems About the Simplest Probabilistic Schemes

6.1.1 Problems About Combinatorial Probability Models (?)

(тут задачи, которые напрямую опираются на комбинаторику)

Ч-1.16

По схеме случайного выбора с возвращением из множества целых чисел $\{0, 1, 2, \dots, 10^n - 1\}$ выбираются числа ξ и η . Обозначим p_m вероятность того, что сумма $\xi + \eta$ будет m -значным натуральным числом в десятичной записи. Найти вероятности $p_{n-k+1}, k = \Rightarrow 0, 1, \dots, n$, и $q_h = \lim_{n \rightarrow \infty} p_{n-k+1}, k = 0, 1, 2, \dots$

Ч-1.17

По схеме случайного выбора с возвратением из множества целых чисел $\{0, 1, 2, \dots, 10^n - 1\}$ выбираются числа ξ и η . Обозначим p_m вероятность того, что произведение $\xi\eta$ будет m -значным натуральным числом в десятичной записи. Найти $q_k = \lim_{n \rightarrow \infty} p_{2n-k} \quad k = 0, 1, 2, \dots$

Количество упаковок для сбора коллекции игрушек в них (????)

В каждую жевачку вложен один из n вкладышей, причем каждый вкладыш встречается с вероятностью $1/n$. Сколько в среднем надо купить жевачек, чтобы собрать полную коллекцию вкладышей?

Пусть у нас уже собрано k вкладышей. Посчитаем, сколько ждать следующего вкладыша. Назовем жевачку новой, если вложенный в нее вкладыш не встречается среди k уже собранных вкладышей. Обозначим M_k среднее количество жевачек, которое нужно купить, чтобы последняя купленная жевачка была новой. Рассмотрим следующую купленную жевачку. Она с вероятностью $\frac{n-k}{n}$ новая (и при этом ожидание $(k+1)$ -ой жевачки в коллекции заканчивается), и с вероятностью $\frac{k}{n}$ она не является новой (в этом случае мы, купив одну жевачку, снова оказываемся в состоянии, когда среднее число жевачек, которые нужно купить до покупки новой, равно M_k). Тем самым, мы приходим к уравнению

$$M_k = \frac{n-k}{n} \cdot 1 + \frac{k}{n} \cdot (M_k + 1),$$

откуда $M_k = \frac{n}{n-k}$. Отсюда получаем, что среднее число жевачек, которое необходимо купить для полной коллекции вкладышей, равно

$$M_0 + M_1 + \dots + M_{n-1} = \frac{n}{n} + \frac{n}{n-1} + \dots + \frac{n}{1} = n \left(1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n} \right).$$

(??? укажу, почему на самом деле с 0-ля суммируем, важный момент.)

(!! очень крутая задача!! укажу в 1й части суть этих методов!!! очень крутой, очень нетривиальный метод.)

Вероятность совпадения нумераций кубика (??)

Двое человек взяли два кубика, в каждом 6 граней, и произвольного пронумеровали стороны. Какова вероятность того, что с точностью до вращений у них получился один и тот же кубик.

Посчитаем число различных способов, которыми можно пронумеровать грани. 1-ю можно выбрать 6-ю способами, занумеровать 6-ю разными цифрами. 2-ю грань - 5-ю способами, занумеровать 5-ю числами. И т.д. В итоге число различных кубиков будет:

$$N = 6 \cdot 6 \cdot 5 \cdot 5 \cdot 4 \cdot 4 \cdot 3 \cdot 3 \cdot 2 \cdot 2 = (6!)^2$$

Поэтому вероятность совпадения $P = \frac{1}{(6!)^2}$

(?? потом проверю, так ли решил как нужно? пока не уверен, 1я часть слабая.)

Ч-1.18

*. Из множества $\{1, \dots, N\}$ по схеме случайного равновероятного выбора с возвращением выбираются элементы $X'_1, \dots, X'_k$, а по схеме равновероятного выбора для любой совокупности $\mathcal{A} := \{A_1, \dots, A_m\}$, состоящей из попарно различных k -элементных подмножеств $A_i = \{a_{i,1}, \dots, a_{i,k}\} \subset \{1, \dots, N\}$, $0 \leq \mathbf{P}\{\{X''_1, \dots, X''_k\} \in \mathcal{A}\} - \mathbf{P}\{\{X'_1, \dots, X'_k\} \in \mathcal{A}\} \leq \frac{1}{N} C_k^2$.

Ч-1.19

*. По схеме случайного выбора с возвращением из множества натуральных чисел $\{1, 2, \dots, N\}$, $N \geq 4$, выбираются числа X и Y . Что больше: и ли

$$p_2 = \mathbf{P}\{X^2 - Y^2 \text{ делится на } 2\}$$

$$p_3 = \mathbf{P}\{X^2 - Y^2 \text{ делится на } 3\}?$$

Ч-1.22

По схеме случайного выбора с возвращением из множества $\{1, 2, \dots, N\}$ выбираются числа X и Y . Показать, что при $N \geq 4$

$$P\{X^3 + Y^3 = 0 \pmod{3}\} < P\{X^3 + Y^3 = 0 \pmod{7}\}.$$

Ч-1.23

Из совокупности всех подмножеств множества $S = \{1, 2, \dots, N\}$ по схеме выбора с возвращением выбираются множества A_1, A_2 . Найти вероятность того, что $A_1 \cap A_2 = \emptyset$.

Ч-1.24

Из совокупности всех подмножеств множества $S = \{1, 2, \dots, N\}$ по схеме выбора с возвращением выбираются подмножества $A_1, A_2, \dots, A_r$. Найти вероятность того, что множества $A_1, A_2, \dots, A_r$ попарно не пересекаются.

Ч-1.25

*. В урне содержится $(2n+1)^2$ карточек, на каждой из которых написана упорядоченная пара целых чисел (x, y) (x и y принимают значения от $-n$ до n , каждая пара чисел написана ровно на одной карточке). Из урны по схеме выбора без возвращения извлекаются три карточки: $(\xi_1, \eta_1), (\xi_2, \eta_2), (\xi_3, \eta_3)$. Рассмотрим эти пары как координаты случайных точек Ξ_1, Ξ_2, Ξ_3 плоскости в декартовой системе координат. Найти вероят-

Ч-1.26

Брошено 10 игральных костей. Предполагается, что все комбинации выпавших очков равновероятны. Найти вероятности событий:

- а) не выпало ни одной «6»;
- б) выпало ровно три «6»;
- в) выпала хотя бы одна «6»;
- г) выпало хотя бы две «6».

Ч-1.27

Из множества $\{0, 1, \dots, N\}$ по схеме равновероятного выбора с возвращением извлекаются числа

Ч-1.28

Найти вероятность того, что в помете случайно выбранного в большом городе автомобиля сумма первых двух цифр равна сумме двух последних.

Ч-1.29

Некоторые москвичи считают трамвайный, троллейбусный или автобусный билет "счастливым", если сумма первых трех цифр его шестизначного номера совпадает с суммой последних трех цифр (?). Найти вероятность получить «счастливые» билеты.

Ч-1.30

Из карточек разрезной азбуки составлено слово "СТАТИСТИКА". Затем из этих 10 карточек по схеме случайного выбора без возвращения отобрано 5 карточек. Найти вероятность того, что из отобранных карточек можно составить слово «ТАИСИ».

Ч-1.31

Из 30 чисел $(1, 2, \dots, 29, 30)$ случайно отбирается 10 различных чисел. Найти вероятности событий: $A = \{ \text{все числа нечетные} \}$, $B = \{ \text{ровно 5 чисел делится на 3} \}$, $C = \{ \text{5 чисел четных и 5 нечетных, причем ровно одно число делится на 10} \}$.

Ч-1.32

Из урны, содержащей M_1 шаров с номером 1, M_2 шаров с номером 2, ..., M_N шаров с номером N , случайно без возвращения выбирается n шаров. Найти вероятности событий: 1) появилось m_1 шаров с номером 1, m_2 шаров с номером 2, ..., m_N шаров с номером N ; 2) каждый из N номеров появлялся хотя бы один раз.

Ч-1.33

Из множества чисел $\{1, 2, \dots, N\}$ по схеме выбора без возвращения выбираются числа ξ_1 и ξ_2 . Найти $P\{\xi_2 > \xi_1\}$. При выборе трех чисел найти вероятность того, что второе число лежит между первым и третьим.

Ч-1.36

Десять рукописей разложены по 30 папкам (на одну рукопись 3 папки). Найти вероятность того, что в случайно выбранных 6 папках не содержится целиком ни одной рукописи.

Ч-1.37

За круглый стол рассаживаются в случайном порядке $2n$ гостей. Какова вероятность того, что гостей можно разбить на n непересекающихся пар так, чтобы каждая пара состояла из сидящих рядом мужчины и женщины?

Ч-1.38

Участник лотереи «6 из 49» на первой карточке отметил номера $(4, 12, 20, 31, 32, 33)$, а на второй $(4, 12, 20, 41, 42, 43)$. Найти вероятность того, что участник получит ровно два минимальных выигрыша, т. е. что каждый из этих наборов имеет ровно 3 общих элемента с набором номеров $(\alpha_1, \dots, \alpha_6) \subset \{1, 2, \dots, 49\}$, появившихся при розыгрыше тиража.

Ч-1.39

Найти вероятность того, что при случайной расстановке двух ладей на шахматной доске они не будут угрожать друг другу.

Ч-1.40

*. Найти вероятность P_k того, что при случайной расстановке k ($2 \leq k \leq 8$) ладей на шахматной доске никакие две ладьи не будут угрожать друг другу. При каких k эта вероятность меньше $1/2$? Менее $1/100$?

Ч-1.41

Собираясь в путешествие на воздушном шаре, Пончик положил в каждый из 20 карманов своего костюма по прянику. Через каждые 10 минут полета у Пончика возникает желание подкрепиться, и он начинает в случайном порядке просматривать свои карманы до тех пор, пока не найдет очередной пряник. Найти вероятность того, что поиск k -го пряника наталкивается на пустой карман.

Ч-1.50

Равновероятной схемой размещения частиц по ячейкам называют схему размещения, в которой по мере ячейек, последовательно занимаемых частицами, получают посредством случайного выбора с возвращением.

Обозначим $\mu_r = \mu_r(n, N)$ число ячеек, содержащих ровно по r частиц после размещения n частиц по N ячейкам. Найти вероятности следующих событий: 1) $\mu_0(n, N) > 0$ (при $n = N$); 2) $\mu_0(n, N) = 0$ (при $n = N + 1$); 3) $\mu_0(n, N) = 1$ (при $n = N + 1$); 4) найдется ячейка, содержащая хотя бы две частицы (при любых соотношениях между n и N).

Ч-1.51

(см. п.1.50). Найти $P\{\mu_0(n, N) = 0\}$ при произвольных n, N .

Ч-1.52

По N различным ячейкам размещается случайно n неразличимых частиц. (Элементарными событиями являются наборы чисел $(r_1, r_2, \dots, r_N)$, где r_k число частиц в k -й ячейке, $k = 1, 2, \dots, N$.) Найти вероятности событий: 1) $\mu_0(n, N) > 0$; 2) $\mu_0(n, N) = 1$.

Ч-1.53

В первом ряду кинотеатра, состоящем из N кресел, сидит n человек. Предполагая, что все возможные размещения этих n человек в первом ряду равновероятны, найти вероятности следующих событий:

- а) $A_{n,N} = \{\text{никакие 2 человека не сидят рядом}\}$;
- б) $B_{n,N} = \{\text{каждый из } n \text{ человек имеет ровно одного соседа}\}$;
- в) $C_{n,N} = \{\text{из любых двух кресел, расположенных симметрично относительно середины ряда, хотя бы одно свободно}\}$.

Ч-1.54

В зале кинотеатра в первых двух рядах, каждый из которых состоит из N кресел, сидит n человек. Найти вероятности следующих событий:

- а) в первом ряду никакие 2 человека не сидят рядом;
- б) во втором ряду каждый человек имеет ровно одного соседа;
- в) в первом ряду из любых двух кресел, расположенных симметрично относительно середины ряда, хотя бы одно свободно.

Ч-1.55

Из всех отображений множества $\{1, 2, \dots, n\}$ в себя случайно выбирается отображение. Найти вероятности событий: 20

- а) выбранное отображение каждый из n элементов переводит в 1;
- б) элемент i имеет ровно k прообразов;
- в) элемент i переводится в j ;
- г) выбранное отображение элементы $i_1, i_2, \dots, i_k$ ($1 \leq i_1 < i_2 < \dots < i_k \leq n$) переводит в элементы $j_1, j_2, \dots, j_k$ соответственно.

Вероятность попадания при одном выстреле при вероятности попадания хотя бы раз при трех (??)

Вероятность хотя бы одного попадания в цель при трех выстрелах равна 0,973. Найти вероятность попадания в цель при одном выстреле.

Решение: Пусть p - вероятность попадания в цель при одном выстреле. Введем событие $X = \{\text{при трех выстрелах есть хотя бы одно попадание}\}$ и противоположное ему событие $\bar{X} = \{\text{при трех выстрелах нет ни одного попадания}\}$.

Вероятность события $\bar{X}$ равна $P(\bar{X}) = (1-p)^3$, тогда вероятность события X равна $P(X) = 1 - P(\bar{X}) = 1 - (1-p)^3$. По условию эта вероятность равна 0,973, откуда получаем уравнение относительно p :

$$\begin{aligned} 1 - (1-p)^3 &= 0,973, \\ (1-p)^3 &= 0,027, \\ (1-p) &= 0,3, \\ p &= 0,7. \end{aligned}$$

Таким образом, вероятность попадания в цель при одном выстреле равна 0,7.

Ответ: 0,7.

Вероятность, что ладьи на шахматной доске не будут бить друг друга

На шахматную доску 8×8 ставится наудачу две ладьи белого и черного цвета. С какой вероятностью они не будут «бить» друг друга. Решение: Испытание состоит в расположении двух ладей на шахматной доске. Найдем общее число исходов эксперимента. Первую ладью можно поставить 64 способами, после чего вторую ладью можно поставить 63 способами. Общее число исходов равно $64 \cdot 63$.

Событие A - ладьи не будут «бить» друг друга. Очевидно, что проще вычислить вероятность противоположного события $\bar{A}$ - ладьи будут «бить» друг друга.

Найдем благоприятные исходы для события $\bar{A}$. Первую ладью также можно поставить 64 способами, а вторую - только 14-ю (число свободных клеток горизонтали и вертикали, на пересечении которых стоит первая ладья). Общее число благоприятных исходов равно $64 \cdot 14$. Тогда $P(\bar{A}) = \frac{64 \cdot 14}{64 \cdot 63} = \frac{2}{9}$, а $P(A) = 1 - P(\bar{A}) = \frac{7}{9}$.

Тест на заболевание

Заболееваемость некоторым заболеванием X составляет 0,1% (то есть им страдает 0,1% населения). Для выявления этого заболевания был разработан тест. Тест не дает ложноотрицательных результатов (то есть любой, у кого есть заболевание, даст положительный результат на него), но частота ложноположительных результатов составляет 5% (то есть около 5% людей, прошедших тест, получают положительный результат, даже если у них нет этой болезни). Предположим, что случайно выбранный человек проходит тест и у него положительный результат на X . Какова вероятность того, что этот человек действительно болен X ?

Вероятность получить сообщения за меньшее время, чем время, за которое наверняка они придут

(натянута задача, нет такого, хотя отдаленно похожее может быть.)

Володя получает 3 сообщения в час.

- Какова вероятность того, что он получит 5 сообщений в течение двух часов?
- Какова вероятность того, что он получит более двух сообщений в течение двух часов?

$$\lambda = 3 \quad P(X = k) = \frac{(\lambda t)^k}{k!} e^{-\lambda t}$$

X - количество получаемых сообщений

- $P(A) = \frac{(3 \cdot 2)^5}{5!} e^{-6} \approx 0,1606$
- $k > 2 \quad t = 2$

$$P(A) = P(k > 2) = 1 - P(0) - P(1) - P(2) = 1 - \left(\frac{6^2}{2!} + \frac{6^1}{1!} + \frac{6^0}{0!} \right) e^{-6} \approx 0,938$$

(??? причем тут распределение Пуассона!!!!? насколько я плохо это знаю, что вообще не готов к такому??)

Кондратенко-1.

Есть n палок, каждую из которых разломали на 2 части. После этого получившиеся части соединили в пары произвольным образом. Какова вероятность того, что получились в точности исходные палки?

Solution.

Точно исходные палки получатся в единственном из всех случаев сборки. Посчитаем количество вариантов собрать n пар из $2n$ частей. Представим, что мы положили части в определенном порядке и потом соединяем 1 и 2, 3 и 4 и так далее. Тогда способов разложить части будет $(2n)!$. Нам неважно, в каком порядке будут лежать части «внутри палки» - палка получится одинаковая. То есть из одного упорядоченного разложения просто перекалывая части «внутри палки» мы можем получить 2^n эквивалентных сборок. Поэтому делим на 2^n . Кроме того, нас не интересует, в каком порядке будут в итоге лежать получившиеся палки. Способов их переложить $n!$. Значит, искомая вероятность $\frac{n! 2^n}{(2n)!}$.

Кондратенко-2.

У страховой компании 10^4 клиентов, вероятность смерти каждого равна $6 \cdot 10^{-3}$, страховой взнос 12 у.е., выплата в случае смерти - 10^3 \$. Найти вероятность того, что доход компании превысит $4 \cdot 10^4$ и вероятность разорения.

Решение. Подробно рассмотрим случай небывалого дохода, разорение - аналогично. Годовой доход равен $12 \cdot 10^4$, значит, должно произойти не более 80 смертей. Применим теорему Муавра-Лапласа:

$$P\left(a \leq \frac{\mu_n - np}{\sqrt{npq}} \leq b\right) \approx \frac{1}{\sqrt{2\pi}} \int_a^b e^{-\frac{t^2}{2}} dt$$

Подгоним под формулу:

$$\mu_n \leq 80 \Leftrightarrow \frac{\mu_n - 10000 \cdot 0.006}{\sqrt{10000 \cdot 0.006(1 - 0.006)}} \leq \frac{80 - 10000 \cdot 0.006}{\sqrt{10000 \cdot 0.006(1 - 0.006)}} = b$$

В нашем случае $a = -\infty$, поэтому искомая вероятность равна

$$P(\mu_n \leq 80) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^b e^{-\frac{t^2}{2}} dt = \frac{1}{2} + \frac{1}{\sqrt{2\pi}} \int_0^b e^{-\frac{t^2}{2}} dt = \frac{1}{2} + \Phi_0(b)$$

Кондратенко-

3. Вероятность попадания одной пули в бочку с бензином равна p . При одном попадании бочка взрывается с вероятностью p_1 , при двух и более - взрывается наверняка. Найти вероятность того, что бочка рванет при n выстрелах.

Решение. Рассмотрим три случая: не попали ни разу, одно попадание и больше двух попаданий. Обозначим эти события A_0, A_1 и A_2 соответственно. Обозначим событие A - бочка взорвалась. Тогда $P(A) = P(A|A_0)P(A_0) + P(A|A_1)P(A_1) + P(A|A_2)P(A_2)$.

Имеем $P(A_0) = (1-p)^n, P(A_1) = C_n^1 p(1-p)^{n-1}, P(A_2) = 1 - P(A_0) - P(A_1)$.

Из условия $P(A|A_0) = 0, P(A|A_1) = p_1, P(A|A_2) = 1$

Подставляем, $P(A) = np p_1 (1-p)^{n-1} + 1 - np(1-p)^{n-1} - (1-p)^n = 1 + (1-p)^{n-1}(np(p_1 - 1)) - (1-p)^n = 1 + (1-p)^{n-1}(np(p_1 - 1) - (1-p))$

Кондратенко-

4. Завод выпускает изделия с вероятностью брака 0.04. Первый контролер находит брак из брака с вероятностью 0.92, второй - 0.98. Найти вероятность, с которой признанное годным изделие будет бракованным

Решение. Пусть события A_1 - деталь изготовлена с браком, $P(A_1) = 0.04, A_2$ - без брака, $P(A_2) = 0.96, B$ - деталь признана годной. Тогда искомая вероятность

$$P(A_1|B) = \frac{P(B|A_1)P(A_1)}{P(B|A_1)P(A_1) + P(B|A_2)P(A_2)} = \frac{0.04 \cdot P(B|A_1)}{0.04 \cdot P(B|A_1) + 0.96 \cdot 1} = \frac{1}{481}$$

так как $P(B|A_1) = \frac{0.08+0.02}{2} = 0.05$

Кондратенко-

5. Вероятность прихода в бюро k человек равна $e^{-\lambda} \frac{\lambda^k}{k!}$. Вероятность получения отказа p . Найти вероятность ровно m отказов.

Решение. Введем дополнительное обозначение: η - количество деталей.

$$\begin{aligned} P(\xi = m) &= P(\xi = m, (\eta = m) + (\eta = m+1) + \dots) = \sum_{k=m}^{\infty} P(\xi = m|\eta = k)P(\eta = k) = \\ &= \sum_{k=m}^{\infty} e^{-\lambda} \frac{\lambda^k}{k!} p^m q^{k-m} C_k^m = \sum_{k=m}^{\infty} e^{-\lambda} \frac{\lambda^k}{k!} p^m q^{k-m} \frac{k!}{m!(k-m)!} = e^{-\lambda} \frac{p^m}{m!} \sum_{k=m}^{\infty} \frac{\lambda^k}{(k-m)!} q^{k-m} = \\ &= e^{-\lambda} \lambda^m \frac{p^m}{m!} \sum_{k=m}^{\infty} \frac{\lambda^{k-m}}{(k-m)!} q^{k-m} = e^{-\lambda} e^{\lambda(1-p)} \frac{p^m \lambda^m}{m!} = e^{-p\lambda} \frac{(p\lambda)^m}{m!} \end{aligned}$$

Кондратенко-

6. На отрезок $[0, L]$ бросают три точки. Найти вероятность того, что третья окажется между первыми двумя

Решение. По сути требуется найти объем множества $\{(x, y, z) \mid (x < z < y) \vee (y < z < x)\}$. Так как все множества вида $\{(x, y, z) \mid (x < y < z)\}$ $(x, y, z$ идут в определенном порядке) получаются друг из друга движениями, не пересекаются и в сумме покрывают весь куб кроме множества нулевой меры, а всего их 6, то объем каждого из них $1/6$. Значит, объем нашего множества $1/3$.

Кондратенко-

7. Найти мат.ожидание и дисперсию случайной величины:

$$p(x) = \begin{cases} \frac{4x^2}{\alpha^3\sqrt{\pi}} e^{-\frac{x^2}{\alpha^2}}, & x > 0 \\ 0, & x \leq 0 \end{cases}$$

Ответ: $M\xi = \frac{\alpha}{2\sqrt{\pi}}, M\xi^2 = \frac{3}{8}\alpha^2$

Кондратенко-

8. На отрезок бросаются две точки. Найти мат. ожидание и дисперсию расстояния между ними.
Решение. Мат. ожидание:

$$M|\xi - \eta| = \iint_{\substack{0 \leq x \leq 1 \\ 0 \leq y \leq 1}} |x - y| dx dy = 2 \iint_{0 \leq x \leq y \leq 1} (y - x) dx dy = 2 \left(\iint_{0 \leq x \leq y \leq 1} y dx dy - \iint_{0 \leq x \leq y \leq 1} x dx dy \right) = \frac{1}{3}$$

Дисперсия:

$$\begin{aligned} D|\xi - \eta| &= M(\xi - \eta)^2 - \frac{1}{9} = \\ &= \iint_{\substack{0 \leq x \leq 1 \\ 0 \leq y \leq 1}} (x - y)^2 dx dy - \frac{1}{9} = \int_0^1 \left(\frac{1}{3} - y + y^2 \right) dy - \frac{1}{9} = \frac{1}{3} - \frac{1}{2} + \frac{1}{3} - \frac{1}{9} = \frac{1}{18} \end{aligned}$$

Кондратенко-

9. Найти мат. ожидание и дисперсию величины ξ с плотностью $p(x) = \frac{1}{2\alpha} e^{-\frac{|x-a|}{\alpha}}$

Решение. Плотность симметрична относительно точки $x = a$, поэтому $M\xi = a$. Найдём дисперсию:

$$\begin{aligned} D\xi &= \frac{1}{2\alpha} \int_{-\infty}^{\infty} (x - a)^2 e^{-\frac{|x-a|}{\alpha}} dx = \frac{1}{2\alpha} \int_{-\infty}^{\infty} x^2 e^{|x|} \alpha dx = \frac{\alpha^3}{2\alpha} \int_{-\infty}^{\infty} x^2 e^{-|x|} dx = \\ &= \frac{\alpha}{2} \int_{-\infty}^{\infty} x^2 e^{-|x|} dx = \alpha^2 \int_0^{\infty} x^2 e^{-x} dx = \alpha^2 \Gamma(3) = 2\alpha^2 \end{aligned}$$

Кондратенко-

10. Найти мат. ожидание и дисперсию гипергеометрического распределения.

Решение. Гипергеометрическое распределение задается $p_m = \frac{C_M^m C_{N-M}^{n-m}}{C_N^n}$

Его смысл: в урне находится N шаров, из них M белых. Из урны достают без возвращения n шаров, выписана вероятность того, что m из них будут белыми.

Решаем через индикаторы: наша с.в. - количество вынутых белых шаров $\xi = I_1 + \dots + I_n$, где I_i с.в. равная 1, если i -тый вытащенный шар белый, 0 иначе. Индикаторы зависимы, $P(I_i = 1) = \frac{M}{N}$, $P(I_i I_j = 1) = \frac{M(M-1)}{N(N-1)}$ при $i \neq j$. $M I_i = M I_i^2 = \frac{M}{N}$. Поэтому

$$M\xi = \sum_{i=1}^n M I_i = \frac{nM}{N}$$

$$\begin{aligned} M\xi^2 &= M(I_1 + \dots + I_n)^2 = M I_1^2 + \dots + M I_n^2 + 2 \sum_{1 \leq i < j \leq n} M I_i I_j = \\ &= M\xi + 2 \cdot \frac{n!}{(n-2)!2!} \cdot \frac{M}{N} \cdot \frac{M-1}{N-1} = \frac{nM}{N} \left(\frac{(n-1)(M-1)}{N-1} + 1 \right) \end{aligned}$$

Отсюда

$$D\xi = M\xi^2 - M^2\xi = \dots = \frac{nM(N-n)(N-M)}{N^2(N-1)}$$

Кондратенко-

11. Найти мат. ожидание и дисперсию числа смен успеха на неуспех и неуспеха на успех в схеме Бернулли.

Решение. Решим задачу для смен успеха на неуспех. Вероятность смены успеха на неуспех при переходе с i -той позиции на j -тую равна pq . Следовательно, мат. ожидание числа смен успеха на неуспех при этом переходе равно pq . Но мат. ожидание суммы равно сумме мат. ожиданий. Следовательно, $M\xi = (n-1)pq$. Дисперсию сами.

Кондратенко-

12. Пусть ξ_1, ξ_2 - случайные пуассоновские величины с параметрами λ_1 и λ_2 соответственно, причём $\lambda_1 \leq \lambda_2$. Доказать, что $\forall t > 0$ выполняется $P(\xi_1 \leq t) \geq P(\xi_2 \leq t)$.

Решение. Заметим, что

$$P(\xi \leq t) = \sum_{k=0}^{[t]} P(\xi = k) = \sum_{k=0}^{[t]} \frac{\lambda^k}{k!} = e^{-\lambda} Q_{[t]}(\lambda)$$

Найдём производную этой вероятности по λ :

$$P'_\lambda(\xi \leq t) = (e^{-\lambda} Q_{[t]}(\lambda))'_\lambda = -e^{-\lambda} Q_{[t]}(\lambda) + e^{-\lambda} Q'_{[t]}(\lambda) = -e^{-\lambda} (Q_{[t]}(\lambda) - Q_{[t]-1}(\lambda)) = -e^{-\lambda} \frac{\lambda^{[t]}}{[t]!} < 0$$

Из чего следует утверждение задачи.

Кондратенко-

13. Даны ξ_1, ξ_2 - независимые, имеют геометрическое распределение. Найти вероятность того, что $\xi_1 = k$ при условии что $\xi_1 + \xi_2 = n$

Решение. По определению, $P(A|B) = \frac{P(AB)}{P(B)}$

Найдём вероятность того, что $\xi_1 = k$ и $\xi_1 + \xi_2 = n$:

$$P(\xi_1 = k, \xi_1 + \xi_2 = n) = P(\xi_1 = k, \xi_2 = n - k) = p_1 q_1^k p_2 q_2^{n-k}$$

Теперь найдём вероятность того, что $\xi_1 + \xi_2 = n$:

$$P(\xi_1 + \xi_2 = n) = \sum_{k=0}^n P(\xi_1 = k, \xi_1 + \xi_2 = n) = \sum_{k=0}^n p_1 q_1^k p_2 q_2^{n-k} = p_1 p_2 \frac{q_1^{n+1} - q_2^{n+1}}{q_1 - q_2}$$

Осталось поделить одно на другое.

Кондратенко-

14. Пусть ξ - геометрически распределённая случайная величина. Найти распределение величины $\eta = \xi^{1 + \frac{(-1)^\xi}{2}}$.

Решение. Заметим, что $\eta = \xi$, если ξ принимает чётное значение и $\eta = 0$ в противном случае. Значит, $P(\eta = 0) = \sum_{k=0}^{\infty} P(\xi = 2k + 1) = \sum_{k=0}^{\infty} pq^{2k+1} = \frac{pq}{1-q^2}$, $P(\eta = 2k + 1) = 0$, $P(\eta = 2k) = pq^{2k}$

Кондратенко-

15. ξ, η - независимые случайные величины с распределением $\mathcal{N}(0, 1)$. Найти распределение (или плотность) величины $\chi = \xi^2 + \eta^2$

Решение. Найдём функцию распределения. $F_\chi(t) = 0$ при $t \leq 0$, так как $\chi \geq 0$. При $t > 0$:

$$F_\chi(t) = P(\chi \leq t) = \iint_{x^2 + y^2 \leq t} p_{\xi, \eta}(x, y) dx dy = \iint_{x^2 + y^2 \leq t} p_\xi(x) p_\eta(y) dx dy = \frac{1}{2\pi} \iint_{x^2 + y^2 \leq t} e^{-\frac{x^2 + y^2}{2}} dx dy$$

Перейдём к полярным координатам, получим:

$$F_\chi(t) = \frac{1}{2\pi} \int_0^{2\pi} d\varphi \int_0^{\sqrt{t}} e^{-\frac{r^2}{2}} r dr = \int_0^{\sqrt{t}} e^{-\frac{r^2}{2}} d\frac{r^2}{2} = \int_0^{t/2} e^{-u} du = -e^{-u} \Big|_0^{t/2} = 1 - e^{-\frac{t}{2}}$$

Плотность будет

$$p_X(t) = \begin{cases} \frac{1}{2}e^{-\frac{t}{2}}, & t \geq 0 \\ 0, & t < 0 \end{cases}$$

Кондратенко-

16. Пусть ξ, η - нормальные распределения с параметрами $(0, 1)$. Найти распределение $\frac{\xi}{\eta}$

Решение. Зная, что $p_\xi(x) = p_\eta(x) = \frac{1}{\sqrt{2\pi}}e^{-\frac{x^2}{2}}$, найдем распределение

$$F_{\frac{\xi}{\eta}}(t) = \iint_{\frac{y}{x} \leq t} \frac{e^{-\frac{x^2+y^2}{2}}}{2\pi} dx dy = \iint_{\text{tg } \varphi \leq t} \frac{e^{-\frac{r^2}{2}}}{2\pi} r dr d\varphi = 2 \left(\frac{\pi}{2} + \text{arctg } t \right) \cdot \frac{1}{2\pi} \int_0^\infty e^{-u} du = \frac{1}{2} + \frac{\text{arctg } t}{\pi}$$

Получили распределение Коши.

Кондратенко-

17. Даны ξ, η - независимые нормально распределенные величины с параметрами $(0, 1)$. Найти $F_{\frac{\xi}{\eta}, \xi^2 + \eta^2}$

Кондратенко-

18. Пусть η - распределение Коши, $\xi = b\eta + a$, где $b \neq 0$. Найти $p_{\frac{1}{\xi}}(x)$

Кондратенко-

19. Выполняется ли ЗБЧ для такой последовательности случайных величин:

$$P(\xi_n = 2^n) = P(\xi_n = -2^n) = 2^{-2n+1}, P(\xi_n = 0) = 1 - 2^{-2n}?$$

Решение. Да, так как дисперсии $D_{\xi_n} = M\xi_n^2 = 1$ ограничены в совокупности.

Кондратенко-

20. Пусть $M\xi = 0$. Доказать, что $M|\xi| \leq \frac{1}{2}(D\xi + 1)$

Решение. $M\xi = 0$, значит, $D\xi = M\xi^2$. Надо доказать, что $2M|\xi| \leq M(\xi^2 + 1)$, то есть $M(\xi^2 - 2|\xi| + 1) \geq 0$ - правда, так как мат. ожидание неотрицательной величины неотрицательно.

Кондратенко-

21. Дана $\{\xi_n\}_{i=1}^\infty$ - последовательность независимых случайных величин, $M\xi_i = 0, D\xi_i < K \quad \forall i \in \mathbb{N}$

Доказать, что $A_n = \frac{1}{\sqrt{n}} M \left| \sum_{i=1}^n \xi_i \right|$ ограничены в совокупности.

Решение. Обозначим $\eta_n = \frac{\xi_1 + \dots + \xi_n}{\sqrt{n}}$. Так как $M\eta_n = 0$, то $M|\eta_n| \leq \frac{D\eta_n + 1}{2}$. Кроме того, $D\eta_n = \frac{D\xi_1 + \dots + D\xi_n}{n} \leq \frac{Kn}{n} = K$. Значит, $A_n = M|\eta_n| \leq \frac{K+1}{2}$

Кондратенко-

22. Пусть $\xi_1, \dots, \xi_n$ - независимые одинаково распределенные случайные величины, $\eta_n = \xi_1 + \dots + \xi_n$. При этом $M\xi_i = 0$ и $\exists D\xi_i$. Также дано, что $\lim_{n \rightarrow \infty} P(\eta_n > \sqrt{n}) = \frac{1}{3}$. Найти $D\xi_i$ и доказать, что $\forall a, b < \infty \lim_{n \rightarrow \infty} P(\eta_n \in [a, b]) = 0$

Решение. Обозначим $\sigma^2 = D\xi_i$. Применим ЦПТ:

$$\frac{1}{3} = \lim_{n \rightarrow \infty} P(\eta_n > \sqrt{n}) = \lim_{n \rightarrow \infty} P\left(\frac{\eta_n}{\sqrt{n}\sigma} > \frac{1}{\sigma}\right) = 1 - \lim_{n \rightarrow \infty} P\left(\frac{\eta_n}{\sqrt{n}\sigma} < \frac{1}{\sigma}\right) = 1 - \int_{-\infty}^{1/\sigma} e^{-\frac{t^2}{2}} dt = \frac{1}{2} -$$

$$\int_0^{1/\sigma} e^{-\frac{t^2}{2}} dt = \frac{1}{2} - \Phi_0(1/\sigma)$$

Таким образом, $D\xi_i = \sigma^2 = \left(\frac{1}{\Phi_0^{-1}(1/6)}\right)^2$. Посмотрим на второе утверждение:

$$a \leq \eta_n \leq b \Leftrightarrow \frac{a}{\sigma\sqrt{n}} \leq \frac{\eta_n}{\sigma\sqrt{n}} \leq \frac{b}{\sigma\sqrt{n}}$$

Применим ЦПТ:

$$P\left(\frac{a}{\sigma\sqrt{n}} \leq \frac{\eta_n}{\sigma\sqrt{n}} \leq \frac{b}{\sigma\sqrt{n}}\right) \rightarrow \int_{a/(\sigma\sqrt{n})}^{b/(\sigma\sqrt{n})} e^{-\frac{t^2}{2}} dt \rightarrow 0, n \rightarrow \infty$$

Кондратенко-

23. Дана производящая функция Φ . Найти характеристическую.

Решение. По определению,

$$f(t) = Me^{it\xi} = \sum_{k=0}^{\infty} P(\xi = k)e^{itk} = \sum_{k=0}^{\infty} (e^{it})^k = \Phi(e^{it})$$

Кондратенко-

24. Доказать, что $\int_{-\infty}^{\infty} p^2(x)dx = \frac{1}{2\pi} \int_{-\infty}^{\infty} |f(t)|^2 dt$

Решение. Рассмотрим независимые случайные величины ξ_1, ξ_2 с плотностью p . Тогда характеристическая функция их разности будет в точности $f_{\xi_1 - \xi_2}(t) = |f_{\xi}(t)|^2$.

По формуле свёртки $p_{\xi_1 - \xi_2}(x) = \int_{-\infty}^{\infty} p(t-x)p(t)dt$, так как $p_{-\xi}(x) = p_{\xi}(-x)$

Запишем формулу обращения:

$$p_{\xi_1 - \xi_2}(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-itx} |f(t)|^2 dt$$

Приравниваем правые части из обеих формул, подставляем $x = 0$, получаем требуемое утверждение.

Кондратенко-

25. Пусть $(\xi_n - \xi)^2 \xrightarrow{P} 0$. Доказать ¹, что $\xi_n^2 \xrightarrow{P} \xi^2$

Решение. Докажем, что $\xi_n \xrightarrow{P} \xi$. Действительно, $P(|\xi_n - \xi| > \varepsilon) = P((\xi_n - \xi)^2 > \varepsilon^2) \rightarrow 0$ при $n \rightarrow \infty$.

Теперь докажем, что из $\xi_n \xrightarrow{P} \xi$ следует $\xi_n^2 \xrightarrow{P} \xi^2$:

$$P(|\xi_n^2 - \xi^2| > \varepsilon) =$$

$$P(|\xi_n^2 - \xi^2| > \varepsilon, |\xi| > C - 1) + P(|\xi_n^2 - \xi^2| > \varepsilon, |\xi| \leq C - 1, |\xi_n| \leq C) +$$

$$P(|\xi_n^2 - \xi^2| > \varepsilon, |\xi| \leq C - 1, |\xi_n| > C) \leq P(|\xi| > C - 1) + P(|\xi - \xi_n| > \frac{\varepsilon}{2C}) + P(|\xi - \xi_n| > 1)$$

Сначала выберем C так, чтобы первое слагаемое стало маленьким. Потом выберем N так, что при $n > N$ второе и третье слагаемое тоже стали малы.

Кондратенко-

26. Пусть $b > a > 0$, Φ - производящая функция. Доказать, что

$$M \frac{1}{(\xi + a)(\xi + b)} = \int_0^1 z^{b-a-1} \int_0^z u^{a-1} \Phi(u) du dz$$

Решение. По определению,

$$\begin{aligned}
\int_0^1 z^{b-a-1} \int_0^z u^{a-1} \Phi(u) du dz &= \int_0^1 z^{b-a-1} \int_0^z u^{a-1} \sum_{k=0}^{\infty} p_k u^k du dz = \\
\sum_{k=0}^{\infty} p_k \int_0^1 z^{b-a-1} \int_0^z u^{k+a-1} du dz &= \sum_{k=0}^{\infty} p_k \int_0^1 z^{b-a-1} \frac{z^{k+a}}{k+a} dz = \\
\sum_{k=0}^{\infty} p_k \int_0^1 \frac{z^{b+k-1}}{k+a} dz &= \sum_{k=0}^{\infty} p_k \frac{1}{(k+a)(k+b)} = M_{\frac{1}{(k+a)(k+b)}}
\end{aligned}$$

Кондратенко-

27. Пусть ξ_i - независимые одинаково распределенные случайные величины, $M\xi_i = a, D\xi_i = \sigma^2$. Доказать, что последовательность $\frac{\xi_1 + \dots + \xi_n}{\xi_1^2 + \dots + \xi_n^2}$ сходится и найти чему.

Решение. Эта последовательность сходится к $\frac{a}{a^2 + \sigma^2}$. Действительно, $M\xi_i^2 = M^2\xi_i + D\xi_i = a^2 + \sigma^2$. Еще заметим, что $\frac{\xi_1 + \dots + \xi_n}{\xi_1^2 + \dots + \xi_n^2} = \frac{\xi_1 + \dots + \xi_n}{n} \cdot \frac{n}{\xi_1^2 + \dots + \xi_n^2}$. Но первая из дробей, согласно ЗБЧ, сходится к a , а последовательность, обратная второй - к $a^2 + \sigma^2$. Осталось доказать, что предел частного есть частное пределов для сходимости по вероятности.

Кондратенко-

28. Пусть ξ_i - н.о.р.с.в., у которых $M\xi_i = 0, D\xi_i = 1$. Доказать, что величина $\eta_n = \frac{\sqrt{n}(\xi_1 + \dots + \xi_n)}{\xi_1^2 + \dots + \xi_n^2}$ асимптотически нормально распределенная

Решение. Заметим, что по ЦПТ $\frac{\xi_1 + \dots + \xi_n}{\sqrt{n}}$ сходится к нормальному распределению, а $\frac{\xi_1^2 + \dots + \xi_n^2}{n}$ по ЗБЧ сходится к своему мат. ожиданию, в нашем случае к 1. Теперь надо правильно сказать, что частное последовательностей сходится к частному пределов.

Кондратенко-

29. Доказать, что ξ - геометрически распределенная случайная величина тогда и только тогда когда $P(\xi = n + k | \xi \geq k) = P(\xi = n)$.

Решение. Посчитаем для геометрического распределения $P(\xi = n + k | \xi \geq k)$:

$$P(\xi = n + k | \xi \geq k) = \frac{P(\xi = n + k, \xi \geq k)}{P(\xi \geq k)}$$

Если ξ геометрически распределена, то $P(\xi = k) = (1 - q)q^k$, откуда

$$P(\xi \geq k) = \sum_{i=k}^{\infty} P(\xi = i) = (1 - q)q^k (1 + q + q^2 + \dots) = (1 - q) \frac{q^k}{1 - q} = q^k$$

Числитель:

$$P(\xi = n + k, \xi \geq k) = P(\xi = n + k, n + k \geq k) = P(\xi = n + k, n \geq 0) = P(\xi = n + k) = (1 - q)q^{n+k}$$

Значит,

$$P(\xi = n + k | \xi \geq k) = \frac{(1 - q)q^{n+k}}{q^k} = (1 - q)q^n = P(\xi = n)$$

Заметим, что для $n < 0$ числитель будет 0, как и $P(\xi = n)$.

В одну сторону доказали. Теперь докажем, что это свойство и число $p = 1 - q = P(\xi = 0)$ однозначно задает распределение ξ . Действительно, при $n < 0$ числитель нулевой, значит, $P(\xi < 0) = 0$. Подставляя $n = 0$ в свойство, получим:

$$p = P(\xi = 0) = P(\xi = k | \xi \geq k) = \frac{P(\xi = k, \xi \geq k)}{P(\xi \geq k)} = \frac{P(\xi = k)}{P(\xi \geq k)}$$

значит,

⁰¹ Верен и более общий факт, который, кстати, и доказывается более изящно. Именно, если $\xi_i \xrightarrow{P} \xi$, а f непрерывна, то $f(\xi_i) \xrightarrow{P} f(\xi)$. Доказательство смотри М.И.Дьяченко, П.Л.Ульянов. Мера и интеграл, стр. 52-53.

$$P(\xi = k) = p \cdot P(\xi \geq k) = p(1 - P(\xi < k))$$

Зная, что $P(\xi < 0) = 0, P(\xi = 0) = p$ и подставляя в это соотношение $k = 1, 2, \dots$ можно найти все вероятности, то есть распределение определено однозначно. Одно такое мы знаем - геометрическое, значит, это оно и есть.

Кондратенко-

30. Найти свёртку двух нормальных распределений с параметрами (a_1, σ_1^2) и (a_2, σ_2^2)

Решение. Нормальное распределение с произвольными параметрами $\eta \sim \mathcal{N}(a, \sigma^2)$ может быть представлено в виде $\eta = \sigma\xi + a, \xi \sim \mathcal{N}(0, 1)$. Характеристическая функция нормального распределения $f_\xi(t) = e^{-\frac{t^2}{2}}$. Кроме того, $f_{b\xi+a} = e^{ita} f_\xi(bt)$ Значит, $f_\eta = f_{\sigma\xi+a} = e^{ita} e^{-\frac{\sigma^2 t^2}{2}}$. Нашли характеристическую функцию нормального распределения с параметрами (a, σ^2) .

Теперь задача. Пусть $\xi_1 \sim \mathcal{N}(a_1, \sigma_1^2), \xi_2 \sim \mathcal{N}(a_2, \sigma_2^2)$. Тогда

$$f_{\xi_1+\xi_2}(t) = f_{\xi_1}(t)f_{\xi_2}(t) = e^{ita_1} e^{-\frac{\sigma_1^2 t^2}{2}} e^{ita_2} e^{-\frac{\sigma_2^2 t^2}{2}} = e^{it(a_1+a_2)} e^{-\frac{(\sigma_1^2+\sigma_2^2)t^2}{2}}$$

Значит, $\xi_1 + \xi_2 \sim \mathcal{N}(a_1 + a_2, \sigma_1^2 + \sigma_2^2)$

Кондратенко-

31. Найти свёртку двух равномерно распределённых на отрезках $[A, B]$ и $[C, D]$ случайных величин.

Решение. Надо найти распределение суммы. Другими словами, надо найти, какую часть площади прямоугольника с углами (A, C) и (B, D) лежит в полуплоскости $x + y \leq \alpha$. Из геометрии очевидно, что при $\alpha < A + C$ эта площадь равна нулю, при $\alpha \in (A + C, A + D)$ происходит квадратичный рост, при $\alpha \in (A + D, C + B)$ - линейный рост, при $\alpha \in (C + B, B + D)$ - опять квадратичный рост, а при $\alpha > B + D$ доля площади равна единице. Таким образом, плотность сначала равна нулю, потом линейно растёт, потом константа, потом линейно убывает, потом опять ноль.

Кондратенко-

32. Являются ли $\sin x, \cos x$ и $|\cos x|$ характеристическими функциями?

Решение. $\sin x$ не является, так как $\sin 0 = 0 \neq 1$.

$\cos x$ выражается через e^{ix} и e^{-ix} через формулы Эйлера:

$$\begin{cases} e^{ix} = \cos x + i \sin x \\ e^{-ix} = \cos x - i \sin x \end{cases}$$

То есть, $\cos t = 1/2(e^{it} + e^{-it})$, что соответствует характеристической функции дискретного распределения $\xi: P(\xi = 1) = P(\xi = -1) = 1/2$

$|\cos x|$ в окрестности нуля совпадает с просто косинусом и имеет в нуле вторую производную, а значит, если это характеристическая функция какой-то случайной величины, то у этой с.в. должен быть момент второго порядка, но если есть момент k -того порядка, то у характеристической функции всюду существует k -тая производная, но у $|\cos x|$ нет производных в точках $\frac{\pi}{2} + \pi n$. Противоречие.

Кондратенко-

33. Является ли функция

$$f(t) = \begin{cases} 1 - t^2, & |t| \leq 1 \\ 0, & \text{иначе} \end{cases}$$

характеристической, если да, то для какого распределения?

Решение. Не является, так как имеет вторую производную в нуле, но не имеет производной в точках ± 1

Кондратенко-

34. Является ли e^{-t^4} характеристической функцией?

Решение. Не является.

Кратко - по теореме Марцинкевича (если х.ф. представляется в виде экспонента в степени многочлен, то такой многочлен не может быть степени больше 2).

Строго:

При разложении в ряд e^{-t^4} первый член, содержащий t будет иметь степень 4. Это значит, что в нуле у нашей функции первая и вторая (и третья) производные равны 0. Но мы знаем, что $f^{(k)}(0) = i^k \alpha_k$, то

есть, у нашей предполагаемой с.в. первый и второй моменты нулевые. Это значит, что мат. ожидание $M\xi = \alpha_1$ и дисперсия $D\xi = \alpha_2 - \alpha_1^2$ равны 0. Значит, $P(\xi = \text{const}) = 1$, но у константы х.ф. e^{itc} , противоречие.

Кондратенко-

35. Пусть $\sum_{k=1}^n a_k = 1$, будет ли $\sum_{k=1}^n a_k \cos kt$ характеристической функцией?

Решение. Рассмотрим ξ :

$$\begin{pmatrix} -n & -n+1 & \dots & -1 & 1 & \dots & n \\ \frac{a_n}{2} & \frac{a_{n-1}}{2} & \dots & \frac{a_1}{2} & \frac{a_1}{2} & \dots & \frac{a_n}{2} \end{pmatrix}$$

$$\text{Тогда } f_\xi = \sum_{k=1}^n a_k \frac{e^{itk} + e^{-itk}}{2} = \sum_{k=1}^n a_k \cos kt.$$

Кондратенко-

36. Пусть $f(t)$ - характеристическая функция. Является ли $\text{Re } f(t)$ характеристической функцией (а если да, то какого распределения)?

Ответ: Да, является. (Указание: $\text{Re } f(t) = \frac{f(t) + f(-t)}{2}$).

Кондратенко-

37. Найти распределение суммы пуассоновских случайных величин через производящие функции

Решение. Производящая функция для пуассоновского распределения с параметром λ :

$$\xi \sim \Pi(\lambda), \Phi_\xi(x) = \sum_{k=0}^{\infty} e^{-\lambda} \frac{x^k \lambda^k}{k!} = e^{-\lambda} \sum_{k=0}^{\infty} \frac{(\lambda x)^k}{k!} = e^{-\lambda} e^{\lambda x} = e^{\lambda(x-1)}$$

Таким образом, если $\xi_1 \sim \Pi(\lambda_1), \xi_2 \sim \Pi(\lambda_2)$, то

$$\Phi_{\xi_1 + \xi_2}(x) = \Phi_{\xi_1}(x) \Phi_{\xi_2}(x) = e^{\lambda_1(x-1)} e^{\lambda_2(x-1)} = e^{(x-1)(\lambda_1 + \lambda_2)} = \Phi_{\Pi(\lambda_1 + \lambda_2)}$$

То есть $\Pi(\lambda_1) + \Pi(\lambda_2) \sim \Pi(\lambda_1 + \lambda_2)$

Кондратенко-

38. Пусть $\xi_n \xrightarrow{P} 0$. Доказать, что $\forall \varepsilon$, такого, что ξ независимо с любым ξ_n , выполнено $\xi \xi_n \xrightarrow{P} 0$.

Решение. Пусть $\varepsilon > 0$, тогда $P(|\xi_n \xi| \geq \varepsilon) = P\left(|\xi_n| \geq \frac{\varepsilon}{|\xi|}\right) P(|\xi| > 0) + P(0 \geq \varepsilon) P(|\xi| = 0) = P\left(|\xi_n| \geq \frac{\varepsilon}{C}\right) P(0 \leq |\xi| < C) + P\left(|\xi_n| \geq \frac{\varepsilon}{C}\right) P(|\xi| \geq C)$. Тогда $\forall \delta > 0 \exists C > 0 : P\left(|\xi_n| \geq \frac{\varepsilon}{C}\right) P(|\xi| \geq C) < 1 \cdot \delta = \delta$, тогда $\exists N > 0 : \forall n > N$ выполнено $P\left(|\xi_n| \geq \frac{\varepsilon}{C}\right) P(0 \leq |\xi| < C) < \delta \cdot 1 = \delta$. Значит, $\xi_n \xrightarrow{P} 0$.

Кондратенко-

39. Пусть $\xi_n \xrightarrow{d} \xi$, верно ли, что $\xi_n - \xi \xrightarrow{d} 0$?

Ответ: Не верно.

Кондратенко-

40. Пусть $x i_n$ - н.о.р.с.в., $\eta_n = \frac{\xi_1 + \dots + \xi_n}{n}$ и $p_{\xi_n} = \frac{1 - \cos x}{\pi x^2}$ $\xi_n \xrightarrow{d} ?$

Кондратенко-

41. Пусть $x i_n$ - н.о.р.с.в. Доказать, что $\forall k \leq n : M\left(\frac{\xi_1 + \dots + \xi_k}{\xi_1 + \dots + \xi_n} = \frac{k}{n}\right)$

Кондратенко-

42. Доказать формулу Байеса.

Кондратенко-

43. Доказать интегральную теорему Муавра-Лапласа (без использования ЦПТ).

Кондратенко-

44. Доказать эквивалентность сходимости $\xi_i \xrightarrow{d} \xi$ и сходимости их производящих функций.

Кондратенко-

45. Доказать ЗБЧ (в простой форме).

Кондратенко-

46. Доказать ЦПТ (в простой форме).
Решения смотри учебник Севастьянова.

Кондратенко-

47. Доказать, что мощность σ -алгебры не может равняться 130.

Решение. У всякой конечной σ -алгебры есть порождающее множество, т.е. множество таких элементов a_i , что $\forall \omega$ из σ -алгебры пересечение a_i и ω равно либо a_i , либо $\emptyset$, и при этом $\bigcup a_i = \Omega$. Такое множество можно построить: взять элемент A из σ -алгебры, потом взять B и, если A и B пересекаются, то оставить вместо них $A \setminus B, B \setminus A$ и $A \cap B$. Затем берём элемент C и продолжаем данную процедуру. Она оборвётся на некотором шаге в силу конечности алгебры. В итоге получится порождающее множество мощности n . Но тогда можно установить биекцию между n -битными словами и элементами σ -алгебры по следующему принципу: на i -ая буква есть 1, если $a_i \in \omega$, и 0 в противном случае. Так как таких слов 2^n , то и мощность σ -алгебры есть 2^n .

Кондратенко-

48. Доказать, что мощность σ -алгебры не может быть счётной.

Решение. Пусть σ -алгебра счётна, тогда можно построить счётную последовательность непустых непересекающихся её элементов таким образом: возьмём $A \neq \Omega$ за первый элемент и $\bar{A}$ за второй. Очевидно, что они не пересекаются, и хотя бы один из них можно опять поделить (иначе σ алгебра была бы конечной). Будем продолжать этот процесс до тех пор, пока не получим $\aleph_0$ непустых непересекающихся элементов. Но любое счётное объединение элементов из неё будет принадлежать σ -алгебре а если брать объединения различных элементов, то получать будем так же разные. Т.о. в σ -алгебре присутствуют как минимум $2^{\aleph_0} = \mathfrak{c}$ элементов.

Кондратенко-

49. $\xi > 0, \exists M\xi$. Доказать, что $1 \leq M\xi M_{\xi}^{\frac{1}{\xi}}$.

Решение. Нужно доказать, что $1 - M\xi M_{\xi}^{\frac{1}{\xi}} \leq 0$. Посмотрим на это выражение, как на дискриминант некоторого квадратного уравнения:

$$D \leq 0 \Leftrightarrow M\xi x^2 - x + \frac{1}{4}M_{\xi}^{\frac{1}{\xi}} \geq 0 \Leftrightarrow M \left(\frac{1}{\xi} (4x^2\xi^2 - 4x\xi + 1) \right) \geq 0 \Leftrightarrow M \frac{(2x\xi + 1)^2}{\xi} \geq 0$$

что верно при любых x , так как это математическое ожидание неотрицательной функции.

Кондратенко-

50. ξ_i - последовательность независимых неотрицательных целочисленных одинаково распределённых случайных величин, ν - независимая с ними неотрицательная целочисленная случайная величина. Найти $\Phi_{\xi_1 + \dots + \xi_{\nu}}$.

$$\begin{aligned} \text{Решение. } \Phi_{\xi_1+\dots+\xi_\nu} &= Mx^{\xi_1+\dots+\xi_\nu} = \sum_{k=0}^{\infty} x^k P(\xi_1+\dots+\xi_\nu=k) = \sum_{k=0}^{\infty} \sum_{m=0}^{\infty} x^k P(\xi_1+\dots+\xi_m=k|\nu=m) \\ &= \sum_{m=0}^{\infty} \sum_{k=0}^{\infty} x^k P(\xi_1+\dots+\xi_m=k) P(\nu=m) = \sum_{m=0}^{\infty} (\Phi_{\xi_1}(x))^m P(\nu=m) = \Phi_\nu(\Phi_{\xi_1}(x)) \end{aligned}$$

Кондратенко-

51. Доказать, что если ξ и η независимы, то ρ (коэффициент корреляции) равен нулю, но обратное не всегда верно.

Решение. $\Rightarrow$ Если ξ и η независимы, то $\text{cov}(\xi, \eta) = M\xi\eta - M\xi M\eta = 0$, а значит и $\rho = 0$.
 $\Leftarrow$ Рассмотрим величину ξ , принимающие значения $0, \frac{\pi}{2}$ и π с равными вероятностями, и величины $\sin \xi$ и $\cos \xi$. $0 = P(\cos \xi = 1, \sin \xi = 1) \neq P(\cos \xi = 1)P(\sin \xi = 1) = \frac{1}{9}$, значит величины зависимы. При этом $\text{cov}(\cos \xi, \sin \xi) = M \cos \xi \sin \xi - M \cos \xi M \sin \xi = \frac{1}{2} M \sin 2\xi = 0$.

Кондратенко-

52. Доказать, что $|\rho(\xi, \eta)| = 1 \Leftrightarrow \eta = a\xi + b$.

Решение. То, что $|\rho(\xi, \eta)| \leq 1$ следует из неравенства Коши-Буняковского $M^2\xi\eta \leq M\xi^2 M\eta^2$. Вспомним, как мы его доказывали: $M(\xi + \lambda\eta)^2$ неотрицательно как мат. ожидание неотрицательной случайной величины, притом равенство нулю достигается тогда, когда эта случайная величина с вероятностью 1 является константой 0. С другой стороны,

$$M(\xi + \lambda\eta)^2 = M\xi^2 + 2\lambda M\xi\eta + \lambda^2 M\eta^2 \geq 0$$

Это означает, что дискриминант квадратного уравнения на λ меньше либо равен нулю, что и даёт нам неравенство Коши-Буняковского. Пусть он равен нулю (как раз в этом случае $|\rho(\xi, \eta)| = 1$), это значит, что существует такое λ , что $M(\xi + \lambda\eta)^2 = 0$, откуда заключаем, что с вероятностью 1 выполнено $\xi = -\lambda\eta$

Кондратенко-

53. Про производящую функцию $\Phi_\xi(x)$ известно, что $\forall n \geq 0 \Phi_\xi\left(\frac{1}{2^n}\right) = \frac{1}{2^n}$. Найти распределение ξ .

Решение. $\Phi_\xi(0) = \Phi_\xi\left(\lim_{n \rightarrow \infty} \frac{1}{2^n}\right) = \lim_{n \rightarrow \infty} \Phi_\xi\left(\frac{1}{2^n}\right) = 0$, значит $p_0 = 0$. Выпишем два равенства:

$$\begin{aligned} \sum_{k=0}^{\infty} \frac{1}{2^k} p_k &= p_0 + \frac{1}{2} p_1 + \frac{1}{4} p_2 + \dots = \frac{1}{2} \\ \sum_{k=0}^{\infty} \frac{1}{2^{2k}} p_k &= p_0 + \frac{1}{4} p_1 + \frac{1}{16} p_2 + \dots = \frac{1}{4} \end{aligned}$$

Домножим второе равенство на 2 и вычтем из первого, помня, что $p_0 = 0$:

$$0 + 0 + \frac{1}{8} p_2 + \dots = 0$$

Все p_i и коэффициенты перед ними неотрицательны, значит все p_i , кроме p_1 , равны 0. Отсюда $p_1 = 1$.

Кондратенко-

54. Является ли функция $\Phi(x) = \frac{2e}{e^2+1} \text{sh } x$ производящей, и если да, то какого распределения?

Решение. Нужно проверить, чтобы в точке $x = 1$ функция равнялась единице и чтобы все её производные в нуле были неотрицательны. Первое условие выполнено.

$$\begin{aligned} \Phi^{2n+1}(x) &= -\frac{2e}{e^2+1} \text{sh } x \\ \Phi^{2n}(x) &= \Phi(x) \end{aligned}$$

В нуле чётные производные равны $\frac{2e}{e^2+1}$, а нечётные равны 0. Таким образом,

$$p_{2n+1} = 0, p_{2n} = \frac{2e}{(e^2+1)(2n)!}$$

Кондратенко-

55. Верно ли, что сходимость почти всюду равносильна сходимости по вероятности? Решение. Из сходимости почти всюду следует сходимость по вероятности, обратное неверно.

Кондратенко-

56. Известно, что $\xi_n \xrightarrow{P} \xi, \xi_n \xrightarrow{P} \eta$. Доказать, что $P(\xi = \eta) = 1$.

Кондратенко-

57. Известно, что $\xi_n \xrightarrow{P} 0, \eta_n \xrightarrow{P} \eta$. Доказать, что $\xi_n \eta_n \xrightarrow{P} 0$.

Кондратенко-

58. Пусть ξ_i - независимые одинаково распределённые случайные величины, с характеристической функцией

$$f_{\xi_i}(t) = \begin{cases} 1 - |t|, & |t| < 1 \\ 0, & |t| \geq 1 \end{cases}$$

найти $\frac{\xi_1 + \dots + \xi_n}{n} \xrightarrow{d} ?$

Решение. Обозначим $\eta_n = \frac{\xi_1 + \dots + \xi_n}{n}$. Найдём $f_{\eta_n}(t)$:

$$f_{\eta_n}(t) = \left(f_{\xi_1}\left(\frac{t}{n}\right)\right)^n = \left(f_{\xi_1}\left(\frac{t}{n}\right)\right)^n = \left(1 - \frac{|t|}{n}\right)^n$$

отсюда получаем:

$$\lim_{n \rightarrow \infty} f_{\eta_n}(t) = e^{-|t|}$$

Кондратенко-

59. Про события A и B известно, что $P(A) = P(B) = \frac{1}{2}$. Доказать, что $P(AB) = P(\bar{A}\bar{B})$

Решение. Напомню обозначения для событий-множеств: $A \cup B = A + B, A \cap B = AB$. Теперь, имея на руках тождество $P(A + B) = P(A) + P(B) - P(AB)$ и правило де Моргана $\overline{A + B} = \bar{A}\bar{B}$, подгоним искомое равенство:

$$P(AB) = P(A) + P(B) - P(A + B) = 1 - P(\overline{A + B}) = 1 - P(\bar{A}\bar{B}) = 1 - (1 - P(\bar{A}\bar{B})) = P(\bar{A}\bar{B})$$

Кондратенко-

60. Известно, что $\xi_n \xrightarrow{d} \xi$.

Верно ли:

- а) $\forall i \exists M \xi_i \Rightarrow \exists M \xi$
- б) $\exists M \xi \Rightarrow \exists M \xi_i$, начиная с некоторого номера
- в) $\exists M \xi_i, \exists M \xi \Rightarrow M \xi_i \rightarrow M \xi$

Решение. Ничего не верно. Приводим контрпримеры:

- а) Возьмем η - последовательность Коши, а

$$\xi_n = \begin{cases} \eta, & \eta \in [-n, n] \\ 0 & \text{иначе} \end{cases}$$

Вроде это называется обрезкой. Получаем $\xi_i \xrightarrow{d} \eta = \xi$. Из-за того, что плотность у распределения Коши симметрична относительно 0, а интеграл по конечной области можно взять, то $\forall i M \xi_i = 0$, но, как мы знаем, $\nexists M \xi$

- б) Опять с Коши. Возьмем

$$\xi_n = \begin{cases} 0, & \eta \in [-n, n] \\ \eta & \text{иначе} \end{cases}$$

Тогда предельная последовательность будет всюду 0 и мат. ожидание у нее 0, но у любой ξ_i мат. ожидания нет, так как у нее остались от Коши «хвосты» на бесконечности, которые, домноженные на x , дают расходящийся интеграл.

- в) Было на лекции скорее всего. Возьмем дискретную с.в.:

$$\xi_n \left(\begin{array}{cc} 0 & n \\ \frac{n-1}{n} & \frac{1}{n} \end{array} \right)$$

Она опять же стремится к константе 0 с мат. ожиданием 0. Но $M \xi_n = 0 \cdot \frac{n-1}{n} + n \cdot \frac{1}{n} = 1$

Кондратенко-

61. Найти коэффициент корреляции между числом выпадения единицы, и числом выпадения шестёрок при бросании кубика n раз

Решение. Рассмотрим случайные величины $\xi_i =$ числу выпадения грани i . Очевидно, что они распределены одинаково и по сути нет разницы между ξ_1 и ξ_6 . Так же ясно, что при бросании кубика n раз выпадет n граней (n - константа). Таким образом получаем:

$$0 = D(n) = D(\xi_1 + \dots + \xi_6) = D\xi_1 + \dots + D\xi_6 + 2 \sum_{1 \leq i < j \leq 6} \text{cov}(\xi_i, \xi_j) = 6D\xi_1 + 2C_6^2 \text{cov}(\xi_1, \xi_6)$$

$$\text{отсюда следует, что } \rho(\xi_1, \xi_6) = \frac{\text{cov}(\xi_1, \xi_6)}{\sqrt{D\xi_1} \sqrt{D\xi_6}} = \frac{\text{cov}(\xi_1, \xi_6)}{D\xi_1} = -\frac{1}{5}.$$

Кондратенко-

62. Доказать, что если ξ, η имеют совместное нормальное распределение (вектор $(\xi, \eta)^T$ нормально распределен) и их ковариация равна нулю, то ξ и η независимы.

Кондратенко-

63. Двумерная случайная величина (ξ, η) равномерно распределена на прямоугольнике. Что можно сказать о зависимости ξ и η ?

Ответ: Независимы.

Кондратенко-

64. Двумерная случайная величина (ξ, η) равномерно распределена на круге. Что можно сказать о зависимости ξ и η ?

Ответ: Зависимы.

6.1.2 Problems About Geometric Probabilities**МФТИ-Т.14 Бросание иглы**

На плоскость, разлинованную параллельными линиями, расстояние между которыми $L = 1$, бросают иглу длины $l \leq L$. Какова вероятность того, что игла пересечет одну из линий?

Solution

Введем параметр x - расстояние от середины иглы до ближайшей линии. Этим параметром, а также углом α , который составляет игла с линией, определяется однозначно положение иглы. (Мы помещаем центр координат плоскости в центр между линиями, в области, где есть центр иглы). Таким образом, пространство событий это множество из всех x, α : $\Omega = [0, L/2] \times [0, \pi]$. Его можно представить как прямоугольник.

В Ω нам нужны события, соответствующие пересечениям иглы с линией, то есть когда проекция иглы на перпендикуляр $\frac{l}{2} \sin \alpha$ больше расстояния от центра иглы до линии: $x \leq \frac{l}{2} \sin \alpha$. В итоге в прямоугольнике Ω появляется нужная нам область площадью $S = \frac{l}{2} \int_0^\pi \sin \alpha d\alpha = \frac{l}{2} (-\cos(\pi) - (-\cos 0)) = l$.

В итоге по формуле геометрической вероятности

$$P = \frac{l}{\pi L/2} = \frac{2l}{\pi L}$$

Ответ: $\frac{2l}{\pi L}$

Комментарий про решение: Можно было ввести и другой угол, тогда проекция была бы $\sim |\cos \phi_2|$, но так не удобнее.

Комментарий про задачу: это известная задача Бюффона. В прошлом были эксперименты по проверке этого ответа и про вычисление числа π , что сейчас не особо уже интересно и важно.

ФР-2023. Парадокс Бертрана о хорде

В круге радиуса R случайно проводится хорда. Обозначим через ξ ее длину. Найти вероятность $Q_x = P(\xi > x)$. Вычислить вероятности Q_R и $Q_{R\sqrt{3}}$ того, что длина хорды больше стороны правильного вписанного шестиугольника и треугольника соответственно. Рассмотреть следующие случаи:

1. Считать, что слово “случайно” означает, что середина хорды равномерно распределена в круге.
2. Считать, что слово “случайно” означает, что направление хорды фиксировано, а середина равномерно распределена на диаметре окружности.
3. Считать, что слово “случайно” означает, что один из концов хорды закреплён, а другой равномерно распределён на окружности.


Указание: В качестве ответа на каждый из пунктов принимается формула для Q_x , выражающая искомую вероятность через R и параметр x , а также два значения, вычисленные по этой формуле: при $x = R$ и при $x = R\sqrt{3}$.

1. Случай «случайного центра»

Хорда длиннее стороны равностороннего треугольника, если выбранная точка находится внутри круга, вписанного в треугольник. Это можно легко увидеть, представим, что мы вращаем вписанный треугольник, так что его сторона может стать параллельно хорде и очевидно, кто из них больше. Поэтому

$$Q_x = \frac{r_{\text{int}}}{R^2} = \frac{R^2 \cos^2 \frac{\pi}{n}}{R^2} = \cos^2 \frac{\pi}{n}.$$

Рассмотрим случай с треугольником. Площадь вписанного круга есть $1/4$ от площади большего, значит, исходная вероятность равна $\frac{1}{4}$.




Рассмотрим случай с шестиугольником. Аналогично, только площадь вписанного круга есть $\pi R^2 \left(\frac{\sqrt{3}}{2}\right)^2$ от площади большего, значит, исходная вероятность равна $\frac{3}{4}$.

2. Случай «случайного диаметра»

Так как все на диаметр в 2-е больше радиуса и все симметрично, то вероятность выбора случайной точки на диаметре такая же, что и на радиусе. Представим, что n -угольник повернут так, что одна из его сторон перпендикулярна зафиксированному радиусу. Хорда длиннее стороны n -угольника, если её центр ближе к центру, чем точка пересечения треугольника с зафиксированным радиусом. Угол между расстоянием от центра до ребра равен (исходя из равного деления дуги) $\frac{2\pi}{2n}$, поэтому

$$Q_x = \frac{R \cos(\pi/n)}{R}$$

Рассмотрим случай с треугольником. Сторона треугольника делит пополам радиус, следовательно вероятность выбрать хорду длиннее стороны треугольника равна $\frac{1}{2}$.




Рассмотрим случай с шестиугольником. Аналогично, но теперь нужно заметить из геометрии, что от центра шестиугольника до его стороны расстояние $R\frac{\sqrt{3}}{2}$. Ответ: $\frac{\sqrt{3}}{2}$.

3. Случай «случайных концов»

Представим, что n -угольник повернут так, что одна из его вершин совпадает с концом хорды. Нам нужно, чтобы конец хорды лежал на большей дуге между двумя ближайшими вершинами шестиугольника от первого зафиксированного конца.

$$Q_x = \frac{2\pi R(n-2)/n}{2\pi R}$$

Рассмотрим случай с треугольником. Длина рассмотренной дуги равна $1/3$ длины окружности, поэтому искомая вероятность равна $\frac{1}{3}$.



Рассмотрим случай с шестиугольником. Аналогично, но теперь Длина этой дуги равна $2/3$ длины окружности, поэтому искомая вероятность равна $\frac{2}{3}$.

МФТИ-Т.15 Три точки на отрезке

На отрезок наудачу последовательно одну за другой бросают три точки. Какова вероятность того, что третья по счету точка попадет между двумя первыми?

Solution

Пусть три точки уже находятся на отрезке. Между двумя крайними точками находится одна из них. Она могла бы попасть первой, могла бы второй, а могла бы и третьей. Таким образом, из всех способов размещения трех точек только один нужен нам, когда средняя точка попадает последней.

Так как какой бы ни был расклад точек, для каждого расклада только ровно $1/3$ из всех возможных событий успешный, то и в итоге вероятность нужного нам события равна $1/3$.

Ответ $1/3$.

МФТИ-Т.16 Вероятность выполнения неравенства на координаты случайной точки в квадрате

Пусть в единичном квадрате $[0, 1] \times [0, 1]$ выбирается случайно точка $A = (\xi, \eta)$. Докажите, что вероятность, что её координаты удовлетворяют $|\eta - \xi| < \frac{1}{5}$, равна:

$$P\left(|\eta - \xi| < \frac{1}{5}\right) = \frac{9}{25}.$$

Solution

Построим вероятностную модель следующим образом. Пространство элементарных событий будет единичный квадрат $\Omega = [0, 1] \times [0, 1]$. σ -алгеброй F будет совокупность борелевских подмножеств Ω , вероятностной мерой P — мера Лебега на Ω .

На $\Omega = [0, 1] \times [0, 1]$ событию $B = \{|\xi - \eta| < \frac{1}{5}\}$ соответствует полоса S , ограниченной отрезками, которые соединяют точку $(0, \frac{1}{5})$ с точкой $(\frac{4}{5}, 1)$ и точку $(\frac{1}{5}, 0)$ с точкой $(1, \frac{4}{5})$. Дополнение к этой области до квадрата Ω состоит из двух равнобедренных прямоугольных треугольников, с длинами катетов равными $\frac{4}{5}$.

(мб рисунок добавлю)

В итоге вероятность — отношение площади полосы S к площади Ω . Поэтому осталось просто посчитать площадь S , что легко сделать, разбив полосу на прямоугольник и два треугольника.

В итоге $P(|\eta - \xi| < \frac{1}{5}) = \frac{|S|}{|\Omega|} = \frac{9}{25}$.

МФТИ-Т.17 Вероятность составления треугольника из двух точек на отрезке

На отрезке наудачу выбираются две точки. Какова вероятность того, что из получившихся трех отрезков можно составить треугольник?

Solution

Треугольник можно составить если для любых сторон его a, b, c выполняется $a < c + b$. То есть $2c < a + b + c = l$, поэтому любая сторона его должна быть меньше $\frac{l}{2}$. Введем координаты $a = \min(x_1, x_2)$, $b = |x_2 - x_1|$, $c = l - \max(x_1, x_2)$

У нас должно выполняться при $x_1 < x_2$:

$$x_1 < \frac{l}{2}; x_2 > \frac{l}{2}; x_2 - x_1 > \frac{l}{2}$$

А при $x_1 > x_2$:

$$x_1 > \frac{l}{2}; x_2 < \frac{l}{2}; x_2 - x_1 < \frac{l}{2}$$

Изобразим это на рисунке, посчитаем долю площади от всего квадрата, запишем ответ.




Figure 6.1: геометрическая вероятность для двух точек на отрезке

В итоге нужная нам область составляет $\frac{1}{4}$ от площади всех возможных вариантов.

Ответ: $\frac{1}{4}$

Ч-1.61

Случайная точка A равномерно распределена на отрезке $[0, 1]$ и делит этот отрезок на две части. Пусть η_1 - длина большей части и η_2 - длина меньшей части. Найти $P\{\eta_1 \leq x\}$, $P\{\eta_2 \leq x\}$ при любом x .

Solution

Ч-1.62

Случайная точка A имеет равномерное распределение в квадрате со стороной 1. Найти вероятности следующих событий:

- расстояние от точки A до фиксированной стороны квадрата не превосходит x ;
- расстояние от точки A до ближайшей стороны квадрата не превосходит x ;
- расстояние от точки A до центра квадрата не превосходит x ;
- расстояние от точки A до фиксированной вершины квадрата не превосходит x .

Solution

Ч-1.63

Случайная точка A имеет равномерное распределение в прямоугольнике со сторонами 1 и 2. Найти вероятностей следующих событий:

- расстояние от A до ближайшей стороны прямоугольника не превосходит x ;
- расстояние от A до любой стороны прямоугольника не превосходит x ;
- расстояние от A до диагоналей прямоугольника не превосходит x .

Solution

Ч-1.64 Свойства точки в квадрате

Случайная точка A имеет равномерное распределение в квадрате со стороной a . Найти вероятность того, что расстояние от A до ближайшей стороны квадрата меньше, чем расстояние от A до ближайшей диагонали квадрата.

Solution

Ч-1.65

Случайная точка X имеет равномерное распределение в квадрате $A = \{(x, y) : |x| \leq a, |y| \leq a\}$. Найти вероятность того, что квадрат с центром X и сторонами длины b , параллельными осям координат, целиком содержится в квадрате A .

Ч-1.67

Случайная точка X равномерно распределена в правильном треугольнике с вершинами $(a, 0)$, $(-a, 0)$, $(0, a\sqrt{3})$. Найти вероятность того, что квадрат с центром X а сторонами длины b , параллельными осям координат, целиком содержится в этом треугольнике.

Ч-1.68

Случайная точка X равномерно распределена в круге $S = \{(x, y) : x^2 + y^2 \leq R^2\}$. Найти вероятность того, что параллельный оси абсцисс отрезок длины R с серединой в точке X целиком содержится в круге S .

Ч-1.69

Случайная точка A имеет равномерное распределение в правильном n -угольнике. Найти вероятность P_n того, что A находится ближе к границе многоугольника, чем к его диагонали. Найти такие числа C и a , что

$$P_n = Cn^2(1 + o(1)), \quad n \rightarrow \infty.$$

Ч-1.70

Случайная точка (ξ_1, ξ_2) равномерно распределена в единичном квадрате $K = \{(u, v) : 0 \leq u \leq 1, 0 \leq v \leq 1\}$. Обозначим η число действительных корней многочлена

$$f_{\xi_1, \xi_2}(x) = \frac{1}{3}x^3 - \xi_1^2 x + \xi_2.$$

Найти вероятности

$$P\{\eta = k\}, \quad k = 1, 3.$$

Ч-1.71

На паркет, составленный из правильных k -угольников со стороной a , случайно бросается монета радиуса r . Найти вероятность того, что упавшая монета не заденет границу ни одного из k -угольников паркета для:

а) $k = 3$ б) $k = 4$; в) $k = 6$.

Ч-1.73

Парадокс Бертра. В круге радиуса R случайно проводится хорда. Обозначим ξ ее длину. Найти вероятность $Q_x = P\{\xi > x\}$, если середина хорды равномерно распределена в круге. Вычислить вероятности Q_R и $Q_{R\sqrt{3}}$ того, что длина хорды больше стороны правильного вписанного шестиугольника и треугольника соответственно.

Результат зависит от того, как понимать слово «случайно». См. задачи п1.74 и п1.75.

Ч-1.74

Решить задачу п1.73, если направление хорды задано, а ее середина равномерно распределена на диаметре, перпендикулярном ее направлению.

Ч-1.75

Решить задачу п1.73, если один конец хорды закреплен, а другой равномерно распределен на окружности.

Ч-1.76

На плоскость, разную параллельными прямыми (расстояние между соседними прямыми равно $2a$), орошена полуокружность радиуса $r < a$; точка (φ, x) (x - расстояние от центра окружности до ближайшей прямой, $0 \leq x \leq a$; φ - угол между этой прямой и диаметром, соединяющим концы дуги) равномерно распределена в прямоугольнике $[0, a] \times [-\pi/2, \pi/2]$. Найти вероятность пересечений с полуокружностью.

Ч-1.77

В интервале времени $[0, T]$ в случайный момент u появляется сигнал длительности Δ . Приемник включается в случайный момент $v \in [0, T]$ на время t . Предположив, что точка (u, v) равномерно распределена в квадрате $[0, T] \times [0, T]$, найти вероятность обнаружения сигнала.

Ч-1.78

Пассажир может воспользоваться трамваями двух маршрутов, следующих с интервалами T_1, T_2 . Момент прихода пассажира определяет на отрезках $[0, T_1], [0, T_2]$ числа u и v , равные временам, оставшимся до прихода трамваев соответствующего маршрута. Предполагая, что точка (u, v) равномерно распределена на $\Omega = \{(u, v) : 0 \leq u \leq T_1, 0 \leq v \leq T_2\}$, найти вероятность того, что пассажир, пришедший на остановку, будет ждать не дольше t ($0 < t < \min(T_1, T_2)$).

Ч-1.79

Однородный прямой круговой цилиндр случайно бросается на горизонтальную плоскость. Найти вероятность того, что цилиндр упадет на боковую поверхность, если его высота h , а радиус основания r . Вычислить эту вероятность при $h = 2r$. При каких h и r вероятности упасть на основание и на боковую поверхность одинаковы?

Ч-1.80

Неоднородный прямой круговой цилиндр случайно бросается на горизонтальную плоскость. Радиус основания цилиндра r , центр тяжести расположен на оси симметрии цилиндра на расстоянии a от одного основания и $b > a$ от другого основания цилиндра. Найти вероятность того, что цилиндр упадет: а) на основание, расположенное ближе к центру; б) на основание, более удаленное от центра тяжести; в) на боковую поверхность.

Ч-1.81

Однородный прямой круговой конус с высотой h и радиусом основания r случайно бросается на горизонтальную плоскость. а) Найти вероятность того, что он упадет на основание; б) вычислить эту вероятность при $r = h$; в) при каком отношении r/h эта вероятность равна $1/4$?

Ч-1.82

Однородное тело, ограниченное сферой и плоскостью, проходящей через центр сферы (полушар), случайно бросается на горизонтальную плоскость. Найти вероятность того, что полушар упадет на плоскую часть своей границы.

6.1.3 Set Property Tasks

Обсудим другие имеющиеся свойства множеств, использующиеся в теории вероятностей.

(!!! убежусь, чтобы в 1 и 2 части была нужная теория про это!! пока теории нет и я не разобрался нормально.)

Вероятность $p(AB + C|AC + B)$ при условиях на вероятность каждого (!?!???)

Вычислить $P(AB + C|AC + B)$, если известно, что $P(B) = 0.4, P(C) = 0.3, P(A|B) = P(A|C) = 0.7$, B и C - несовместные события.

Как всегда:

$$P(AB + C|AC + B) = \frac{(AB + C)(AC + B)}{AC + B} =$$

Также $P(AB) = P(A|B)P(B) = 0.7 \cdot 0.4 = 0.28$, $P(AC) = P(A|C)P(C) = 0.7 \cdot 0.3 = 0.21$.
(?!?!?! а дальше какое-то преобразование неизвестное мне.)

МФТИ-? Стрелки, начинающие стрелять по жребью (?????)

Два стрелка соревнуются, по очереди стреляя в мишень до первого попадания. Вероятность попадания в мишень стрелка А равна p_1 , а стрелка Б - p_2 . Право первого выстрела определяется жребием, в котором стрелок А побеждает с вероятностью α . Вычислить значение α , при котором шансы на победу у стрелков А и Б равны, если $p_1 = 0.84, p_2 = 0.8$.

Solution

(??? честно, я хз, не отработал еще это.)

МФТИ-Т.1

Пусть A, B - два события. Найти все события X такие, что $\overline{(X \cup A)} \cup \overline{(X \cup \bar{A})} = B$
По правилу де Моргана:

$$\bar{B} = (X \cup A) \cap (X \cup \bar{A})$$

По правилу де Моргана $(X \cup \bar{A}) = \overline{(\bar{X} \cap A)}$, таким образом:

$$(X \cup A) \cap \overline{(\bar{X} \cap A)} = \overline{(\bar{X} \cup \bar{A})} \cup \overline{(\bar{X} \cap A)} = \overline{(\bar{X} \cap \bar{A})} \cup \overline{(\bar{X} \cap A)} = X$$

Ответ $X = \bar{B}$.

МФТИ-Т.2

Пусть A, B - два события. Найти все события X такие, что $AX = AB \iff A \cup X = A \cup B$.

Очевидно, что необходимо, чтобы $X \in A \cap B$. Но также X может содержать любое множество $Y \in X$, которое при $A \cup X$ дало бы ноль, таким образом задача решена.

Ответ $X = Y \cup (A \cap B) \forall Y \in \bar{A}$.

МФТИ-Т.3

Пусть $A_1, \dots, A_n$ - события. Покажите, что

$$\mathbf{P} \left(\bigcup_{k=1}^n A_k \right) + \mathbf{P} \left(\bigcap_{k=1}^n \bar{A}_k \right) = 1$$

Воспользуемся свойству вероятностной меры: $\forall B \quad \mathbf{P}(B) + \mathbf{P}(\bar{B}) = 1$. Возьмем в качестве B выражение: $B = \bigcup_{k=1}^n A_k$, тогда $\bar{B} = \bigcap_{k=1}^n \bar{A}_k$. Подставляя, получаем:

$$\mathbf{P} \left(\bigcup_{k=1}^n A_k \right) + \mathbf{P} \left(\bigcap_{k=1}^n \bar{A}_k \right) = 1$$

МФТИ-Т.4

Пусть $A_1, A_2, \dots$ - последовательность событий и $\mathbf{P}(A_n) = 1$ для всех $n = 1, 2, \dots$. Покажите, что

$$\mathbf{P} \left(\bigcap_{n=1}^{\infty} A_n \right) = 1$$

Первый способ: Пусть $B_n = \bigcap_{i=1}^n A_i$. Заметим, что $B_1 \in B_2 \in \dots$, тогда

$$P(\bar{B}_n) = P \left(\bigcap_{i=1}^n \bar{A}_i \right) \leq \sum_{i=1}^n P(\bar{A}_i) = 0$$

Поэтому в силу положительной определенности $P(\bar{B}_n) = 0$, поэтому $P(B_n) = 1$

Окончательно, заменяя A на B и используя непрерывность вероятностной меры получаем:

$$\mathbf{P} \left(\bigcap_{n=1}^{\infty} A_n \right) = \mathbf{P} \left(\lim_{n \rightarrow \infty} B_n \right) = \lim_{n \rightarrow \infty} \mathbf{P}(B_n)$$

Второй способ: знаем, что на всем вероятностном пространстве Ω выполняется $\mathbf{P}(\Omega) = 1$. Таким образом $A_i = \Omega - A_{i0}$, где $A_{i0} : \mathbf{P}(A_{i0}) = 0$. Таким образом,

$$\mathbf{P} \left(\bigcap_{n=1}^{\infty} A_n \right) = \mathbf{P} \left(\Omega - \bigcap_{n=1}^{\infty} A_{i0} \right) = 1 - \sum_{n=1}^{\infty} 0 = 1$$

МФТИ-Т.5

Пусть $A_1, A_2, \dots$ - последовательность событий. Покажите, что

$$P\left(\lim_{n \rightarrow \infty} A_n\right) \leq \lim_{n \rightarrow \infty} P(A_n) \leq \overline{\lim}_{n \rightarrow \infty} P(A_n) \leq P\left(\overline{\lim}_{n \rightarrow \infty} A_n\right)$$

Вспомним определение:

$$A^* = \overline{\lim}_{n \rightarrow \infty} A_n = \bigcap_{n=1}^{\infty} \bigcup_{k \geq n} A_k, \quad A_* = \lim_{n \rightarrow \infty} A_n = \bigcup_{n=1}^{\infty} \bigcap_{k \geq n} A_k$$

Первое неравенство доказывается с помощью непрерывности вероятностной меры:

$$P\left(\lim_{n \rightarrow \infty} A_n\right) = P\left(\bigcup_{n=1}^{\infty} \bigcap_{k \geq n} A_k\right) \equiv P\left(\bigcup_{n=1}^{\infty} B_n\right) \equiv P(\lim B_n) \stackrel{\substack{\text{т.к. } B_1 \subseteq B_2 \subseteq \dots \\ \text{по непрер. меры}}}{=} \lim P(B_n) \stackrel{\substack{\leq \\ B_n = \bigcup_{N=n}^{\infty} A_n \subset A_n}}{=} \lim_{n \rightarrow \infty} P(A_n)$$

Выражение $\lim_{n \rightarrow \infty} P(A_n) \leq \overline{\lim}_{n \rightarrow \infty} P(A_n)$ очевидно.

И последнее неравенство:

$$P\left(\overline{\lim}_{n \rightarrow \infty} A_n\right) = P\left(\bigcup_{n=1}^{\infty} \bigcap_{k \geq n} A_k\right) \equiv P\left(\bigcup_{n=1}^{\infty} B_n\right) \equiv P(\lim B_n) \stackrel{\substack{\text{т.к. } B_1 \supseteq B_2 \supseteq \dots \\ \text{по непрер. меры}}}{=} \lim P(B_n) \stackrel{\substack{\geq \\ B_n = \bigcup_{N=n}^{\infty} A_n \supset A_n}}{=} P(A_n)$$

МФТИ-Т.6 Неравенство $|P(AB) - P(A)P(B)| \leq \frac{1}{4}$

Покажите, что для любых двух событий A и B выполняется неравенство

$$|P(AB) - P(A)P(B)| \leq \frac{1}{4}$$

Обозначим $AB \equiv c$, $A \setminus B \equiv a$, $B \setminus A \equiv b$, тогда доказать нужно нам, что

$$|c - (a + c)(b + c)| \leq \frac{1}{4}$$

Причем по условию, так как мы работаем с событиями, должно выполняться: $a + b + c \leq 1$. Докажем, что $c - (a + c)(b + c) \geq -\frac{1}{4}$. Действительно, $2\sqrt{ab} \leq a + b \leq 1 \rightarrow ab \leq \frac{1}{4}$.

$$c - (a + c)(b + c) = (1 - \underbrace{a + b + c}_{\leq 0}) - ab \geq -\frac{1}{4}$$

Докажем, что $c - (a + c)(b + c) \leq \frac{1}{4}$. Преобразуем выражение уже в другом виде:

$$c - (a + c)(b + c) \leq c - c^2 = c(1 - c) = (\sqrt{c}\sqrt{1 - c})^2 \stackrel{\substack{\leq \\ \text{неравенство Коши}}}{=} \left(\left(\frac{1}{2}\right)c + (1 - c)\right)^2 \leq \frac{1}{4}$$

Таким образом, уж точно

$$|c - (a + c)(b + c)| \leq \frac{1}{4}$$

6.1.4 Bone and Coin Problems (!!!)

(пока по актуальности собираю)

МФТИ-Т.16 Вероятность, что при подбрасывании двух костей сумма в 5 очков появится раньше, чем сумма в 7 очков

Случайный эксперимент заключается в последовательном подбрасывании двух игральных костей. Найти вероятность того, что сумма в 5 очков появится раньше, чем сумма в 7 очков.

Найдем, чему равна вероятность, что при подбрасывании двух игральных костей выпадет 5: всего таких исходов 4: если выпадает 1,4; 4,1; 2,3; 3,2. Таким образом, $p_5 = 4/36 = 1/9$

Вероятность, что при подбрасывании двух игральных костей выпадет 7: всего таких исходов 4: если выпадает 2,5; 5,2; 2,4; 4,2; 1,6; 6,1. Таким образом, $p_7 = 6/36 = 1/6$

Пусть A - событие, состоящее в выпадении 5 раньше 7. Оно состоит из исхода, при котором на i -м шаге выпадает 5, а на $(i-1)$ -х шагах не выпадало ни 5, ни 7, обозначим его B_i . Запишем событие A в виде

$$A = \bigcup_{i=1}^{\infty} B_i$$

тогда его вероятность:

$$P(A) = \sum_{i=1}^{\infty} P(B_i)$$

Также $P(B_i) = (1 - p_5 - p_7)^{i-1} p_5$, потому что отдельные броски независимы. Таким образом, осталось посчитать ответ:

$$P(A) = \sum_{i=1}^{\infty} (1 - p_5 - p_7)^{i-1} p_5 = \frac{p_5}{1 - (1 - p_5 - p_7)} = \frac{p_5}{p_5 + p_7} = 4/(6 + 4) = 2/5$$

Ответ: $\frac{2}{5}$

МФТИ-Т.? Вероятность первого выпадения герба у каждого из троих подбрасывающих (?!!!!)

Трое игроков по очереди подбрасывают монету. Выигрывает тот, у кого раньше появится герб. Найти вероятности выигрыша каждого игрока

(!! оч крутая задача)

Пусть A_i - вероятность того, что у первого игрока при i -м бросании будет герб. Аналогично, B_i - вероятность того, что у второго игрока при i -м бросании будет герб, C_i - у третьего.

Тогда запишем $A = \bigcup A_i$, $B = \bigcup B_i$, $C = \bigcup C_i$.

И также..

По формуле полной вероятности $P(A) = P(A_1)P(A|A_1) + P(\overline{A_1})P(A|\overline{A_1})$.

И понятно, что $P(A|A_1) = 1$, $P(A_1) = P(\overline{A_1}) = \frac{1}{2}$, поэтому

$$P(A) = \frac{1}{2}(1 + P(A|\overline{A_1}))$$

Дальше заметим, что $P(A|\overline{A_1})$ это вероятность успеха после двух последовательных неудач второго и третьего игрока или после 5 неудач или после 8 неудач и т.п. Это же в точности $P(C)$.

По этой же логике $P(B|\overline{A_1}) = P(A)$

Формула полной вероятности для $P(B)$:

$$P(B) = P(B|A_1)P(A_1) + P(\overline{A_1})P(B|\overline{A_1}) = 0 + \frac{1}{2}P(A)$$

В итоге получается система

$$P(A) + P(B) + P(C) = 1 \quad (6.1)$$

$$P(C) = 2P(A) - 1 \quad (6.2)$$

$$P(B) = \frac{1}{2}P(A) \quad (6.3)$$

Она легко решается.

Ответ: $P(A) = \frac{4}{7}$, $P(B) = \frac{2}{7}$, $P(C) = \frac{1}{7}$.

Одна единица при бросании четырех игральных костей или хотя бы одна пару единиц за один бросок при 24 бросаниях двух костей?

Что вероятнее, получить хотя бы одну единицу при бросании четырех игральных костей или хотя бы одну пару единиц за один бросок при 24 бросаниях двух костей?

Посмотрим на случай четырех костей.

Можно считать количество успешных исходов, но это сделать сложно, так что пойдем от обратного, посчитаем неудачные исходы.

Вероятность невыпадения единицы в одном испытании: $1 - \frac{1}{6} = \frac{5}{6}$, в четырех экспериментах вероятность того, что среди всех выпавших костей нет ни одной единицы равна $(\frac{5}{6})^4$.

Тогда вероятность выпадения единицы равна

$$P_1 = 1 - \left(\frac{5}{6}\right)^4$$

Посмотрим на случай двух костей. При одном бросании двух костей у нас 36 может быть разных исходов, при этом подходит только один. Поэтому аналогично первому пункту, вероятность хотя бы за какой-то бросок получить пару единиц, равна:

$$P_2 = 1 - \left(1 - \frac{1}{36}\right)^{24} = 1 - \left(\frac{35}{36}\right)^{24}$$

Осталось только сравнить эти P_2 и P_1 . Просто посчитаем их на калькуляторе. В итоге

$$P_1 > P_2$$

Ответ: Успех в случае бросания одной кости имеет большую вероятность.

При подбрасывании трёх игральных костей 11 очков выпадают чаще, чем 12

(потренируюсь, сверну скорее всего потом в пару строк, тут всё очевидно, просто подсчетом событий. пока потренироваться еще нужно, научиться на самом деле быстро решать такие задачи.)

Объяснить, почему при подбрасывании трёх игральных костей 11 очков выпадают чаще, чем 12.

Посчитаем количество исходов, которые дадут 11 очков в сумме. По сути это количество целочисленных положительных решений уравнения $x_1 + x_2 + x_3 = 11$ при условии $x_1, x_2, x_3 \leq 6$.

$$P(x_1 + x_2 + x_3 = 11) = \sum_{i=1}^6 P(a = i)P(x_2 + x_3 = 11 - i) = \frac{1}{6} \sum_{i=5}^{10} P(x_2 + x_3 = i)$$

$$P(x_1 + x_2 + x_3 = 12) = \sum_{i=1}^6 P(a = i)P(x_2 + x_3 = 12 - i) = \frac{1}{6} \sum_{i=6}^{11} P(x_2 + x_3 = i)$$

$$P(x_1 + x_2 + x_3 = 11) - P(x_1 + x_2 + x_3 = 12) = P(x_2 + x_3 = 11) - P(x_2 + x_3 = 5)$$

считаем благоприятные $\frac{2}{36} - \frac{4}{36} < 0$
исходы

Поэтому при подбрасывании трёх игральных костей 11 очков выпадают чаще, чем 12.

6.1.5 Map Tasks (!!!)

(потренирую позже!!)

вероятность того, что среди выбранных карт из колоды будут представительницы всех мастей

Из колоды в 52 карты наудачу берется 6 карт. Какова вероятность того, что среди них будут представительницы всех четырех мастей?

Назовем наши масти A, B, C, D. Назовем A_A, A_B, A_C, A_D - исходы, в которых есть масть A, B, C, D соответственно.

Нам нужно событие $X = A_A \cap A_B \cap A_C \cap A_D$ - что есть представители всех мастей.

Посчитаем вероятность события $\bar{X} = \bar{A_A} \cup \bar{A_B} \cup \bar{A_C} \cup \bar{A_D}$ - что нет представителя какой-то из мастей. Для этого воспользуемся формулой включения-исключения:

$$P\left\{\bigcup_{k=1}^n A_k\right\} = \sum_{m=1}^n (-1)^{m-1} \sum_{1 \leq i_1 < \dots < i_m \leq n} P\{A_{i_1} A_{i_2} \dots A_{i_m}\}$$

И так как одной из четырех масти не может быть - 4мя способами, двух мастей - $C_4^2 = 6$ способами, трех - $C_4^3 = 4$ способами, то вторая сумма, учитывая что вероятности наличия разных мастей не зависят от конкретного вида масти, будет иметь вид:

$$P\left\{\bigcup_{k=1}^n A_k\right\} = 4P(\overline{A_1}) - 6P(\overline{A_1 A_2}) + 3P(\overline{A_1 A_2 A_3})$$

Посчитаем вероятности членов в этой сумме, найдя количество благоприятствующих исходов:

$$P(\overline{A_i}) = \frac{C_{52-13}^6}{C_{52}^6}; \quad P(\overline{A_i A_j}) = \frac{C_{52-13-13}^6}{C_{52}^6}; \quad P(\overline{A_i A_j A_k}) = \frac{C_{52-13-13-13}^6}{C_{52}^6};$$

Подставим все в ответ: $P_{\text{искомое}} = 1 - \left(4 \frac{C_{39}^6}{C_{52}^6} - 6 \frac{C_{26}^6}{C_{52}^6} + 4 \frac{C_{13}^6}{C_{52}^6}\right)$
(?? число???)

6.2 Test Sequence Tasks (!!!)

(нужно проработать этот раздел, типичные задачи!)

6.2.1 Problems About Conditional Probabilities

Ч-2.пр.2

Из урны, содержащей 3 белых и 2 черных шара, по схеме случайного выбора без возвращения последовательно извлекаются пары. Найти вероятность p_k того, что черный шар впервые появится при k -м испытании ($k = 1, 2, 3, 4$).

Обозначим C_i событие, состоящее в том, что в i -м испытании появился черный шар. Тогда события $B_k = \{\text{впервые черный шар появился при } k\text{-м испытании}\}$, $k = 1, 2, 3, 4$ можно выразить через C_i и $\bar{C}_i$:

$$B_1 = C_1, \quad B_2 = \bar{C}_1 C_2, \quad B_3 = C_1 C_2 C_3, \quad B_4 = \bar{C}_1 C_2 \bar{C}_3 C_4$$

По формуле (n2.3)

$$\begin{aligned} P(B_1) &= P(C_1), \quad P(B_2) = P(\bar{C}_1) P(C_2|\bar{C}_1), \\ P(B_3) &= P(\bar{C}_1) P(\bar{C}_2|\bar{C}_1) P(C_3|\bar{C}_1 \bar{C}_2) \\ P(B_4) &= P(\bar{C}_1) P(\bar{C}_2|\bar{C}_1) P(\bar{C}_3|\bar{C}_1 \bar{C}_2) P(C_4|\bar{C}_1 \bar{C}_2 \bar{C}_3) \end{aligned}$$

По классическому определению вероятности

$$\begin{aligned} P(C_1) &= \frac{2}{5}, \quad P(\bar{C}_1) = \frac{3}{5}, \quad P(\bar{C}_{i+1}|\bar{C}_1 \dots \bar{C}_i) = \frac{3-i}{5-i} \\ P(C_{i+1}|\bar{C}_1 \dots \bar{C}_i) &= \frac{2}{5-i}, \quad i = 1, 2, 3. \end{aligned}$$

В результате получим

$$\begin{aligned} p_1 &= P(B_1) = \frac{2}{5} = 0,4, \quad p_2 = P(B_2) = \frac{3}{5} \cdot \frac{2}{4} = 0,3, \\ p_3 &= P(B_3) = \frac{3}{5} \cdot \frac{2}{4} \cdot \frac{2}{3} = 0,2 \\ p_4 &= P(B_4) = \frac{3}{5} \cdot \frac{2}{4} \cdot \frac{1}{3} \cdot \frac{2}{2} = 0,1. \end{aligned}$$

Ч-2.пр.3

Найти вероятность того, что среди 10 однозначных случайных чисел ровно 4 четных числа и 2 нечетных числа, кратных 3.

Решен и е. Однозначное случайное число четно с вероятностью $p_1 = \frac{5}{10} = \frac{1}{2}$, нечетно и кратно 3 с вероятностью $p_2 = \frac{2}{10} = \frac{1}{5}$. Вероятность остальных исходов $p_3 = 1 - p_1 - p_2 = \frac{3}{10}$. Если ξ_1 - число четных чисел среди 10 случайных чисел, ξ_2 - число нечетных чисел, кратных 3, то число остальных чисел $\xi_3 = 10 - \xi_1 - \xi_2$, и по формуле (n2.10) с $N = n = 10$ находим $P\{\xi_1 = 4, \xi_2 = 2\} = P\{\xi_1 = 4, \xi_2 = 2, \xi_3 = 5\} =$

$$= \frac{10!}{4!2!4!} \left(\frac{1}{2}\right)^4 \left(\frac{1}{5}\right)^2 \left(\frac{3}{10}\right)^4 = 0,063787...$$

Частный случай полиномиальной схемы с $N = 2$ называют схемой Бернулли. Ниже два исхода каждого испытания в схеме Бернулли будем обозначать символами 1 и 0 или называть успехом и неудачей, а соответствующие им вероятности - буквами p и $q = 1 - p$. Если $\mu_n \rightarrow$ число успехов (или число единиц) в n испытаниях Бернулли, то

$$P\{\mu_n = m\} = P_n(m, p) = C_n^m p^m q^{n-m}, \quad m = 0, 1, \dots, n. \quad (\text{n2.11})$$

Ч-2.пр.4

Проводится n независимых опытов, состоящих в одновременном подбрасывании k монет. Вычислить вероятности событий: $A = \{\text{хотя бы один раз все } k \text{ монет выпали гербами}\}$, $B = \{\text{ровно } m \text{ раз все } k \text{ монет выпали гербами}\}$.

Появление k гербов в одном опыте будем называть успехом. Вероятность успеха $p = \frac{1}{2^k}$. Проще вычислить вероятность противоположного события не выпали гербами. Положим $C_l = \{\text{в } l\text{-м испытании произошел успех}\}$. Тогда $\bar{A} = \bar{C}_1 \bar{C}_2 \dots \bar{C}_n$, и по формуле (n2.6) находим

$$P(\bar{A}) = \prod_{l=1}^n P(\bar{C}_l) = \left(1 - \frac{1}{2^k}\right)^n.$$

Для вычисления $P(\bar{A})$ можно было также воспользоваться формулой (n2.11), положив в ней $m = 0, p = \frac{1}{2^k}, q = 1 - \frac{1}{2^k}$. Таким образом,

$$P(A) = 1 - P(\bar{A}) = 1 - \left(1 - \frac{1}{2^k}\right)^n.$$

Вероятность события B определяем по формуле (n2.11):

$$P(B) = C_n^m \left(\frac{1}{2^k}\right)^m \left(1 - \frac{1}{2^k}\right)^{n-m}$$

Ч-2.пр.5

В партии из $n = 200$ изделий каждое изделие независимо от остальных может быть бракованным с вероятностью $p = 0,01$. Определить вероятность того, что число μ_n бракованных изделий в этой партии равно трем.

Решение. Величина $np^2 = 0,02$ мала. В качестве приближенного значения искомой вероятности можно использовать предельное значение в (n2.12) (теорема Пуассона) с $\lambda = np = 2$: $P\{\mu_n = 3\} \approx \frac{2^3}{3!} e^{-2}$. По табл. 3 (с. 308) находим числовое значение: $P\{\mu_n = 3\} \approx 0,18045\dots$ Точное значение вероятности можно найти по формуле (n2.11):

$$P\{\mu_n = 3\} = C_{200}^3 (0,01)^3 (0,99)^{197} = 0,181355\dots$$

Ч-2.пр.6

В партии из $n = 22500$ изделий, каждое изделие независимо от других может быть бракованным с вероятностью $p = 1/5$. Найти вероятность того, что число μ_n бракованных изделий заключено между 4380 и 4560.

Значение $npq = 3600$ велико, поэтому можно воспользоваться нормальным приближением (n2.14). Вычтем $np = 4500$ из трех частей неравенства $4380 < \mu_n < 4560$ и получившееся неравенство поделим почленно на $-2 < \frac{\mu_n - np}{\sqrt{npq}} < 1$. $P\{4380 < \mu_n < 4560\} = P\left\{-2 < \frac{\mu_n - np}{\sqrt{npq}} < 1\right\} \approx \frac{1}{\sqrt{2\pi}} \int_{-2}^1 e^{-\frac{x^2}{2}} dx$ Так как

$$\frac{1}{\sqrt{2\pi}} \int_{-2}^1 e^{-\frac{x^2}{2}} dx = \frac{1}{\sqrt{2\pi}} \int_0^1 e^{-\frac{x^2}{2}} dx + \frac{1}{\sqrt{2\pi}} \int_0^2 e^{-\frac{x^2}{2}} dx = \Phi_0(1) + \Phi_0(2)_2$$

то (см. табл. 1, с. n306-n307).

$$P\{4380 < \mu_n < 4560\} \approx 0,3413 + 0,4772 = 0,8185.$$

Ч-2.пр.7

Найти вероятность того, что в схеме Бернулли первый успех появится в k -м испытании, если вероятность успеха в отдельном испытании равна p . Реше. $A_k = \{ \text{первый успех появился в } k \text{ - испытании} \}$ определяется исходами в k первых испытаниях. Пусть событие $C_i = \{ \text{в } i \text{ - М испытании наступил успех} \}$. Тогда

$$A_k = \bar{C}_1 \bar{C}_2 \dots \bar{C}_{k-1} C_k,$$

и по формуле (п2.17) находим

$$P(A_k) = P(\bar{C}_1) P(\bar{C}_2) \dots P(\bar{C}_{k-1}) P(C_k) = q^{k-1} p.$$

Приведем еще две формулы, полезные в тех случаях, когда заданы или могут быть легко вычислены условные вероятности. Если события $B_1, B_2, \dots, B_n$ попарно несовместны и $B_1 \cup B_2 \cup \dots \cup B_n = \Omega$, то для любого события A имеем

$$P(A) = \sum_{k=1}^n P(B_k) P(A|B_k)$$

(формула полной вероятности) и

$$P(B_m|A) = \frac{P(B_m) P(A|B_m)}{\sum_{k=1}^n P(B_k) P(A|B_k)}, \quad m = 1, \dots, n_s$$

Ч-2.пр.8

В ящик, содержащий 8 исправных изделий, добавлено 2 изделия, взятых со склада; известно, что доля бракованных изделий на складе равна 5%. Найти вероятность того, что взятое наудачу из пополненного ящика изделие не будет бракованным. Решение. Определим события A, B_0, B_1, B_2 ? A = изделие, взятое из пополненного ящика, $B_k = \{ \text{из двух изделий, добавленных в ящик, } k \text{ бракованных, } k = 0, 1, 2. \}$ Очевидно, что $B_0 \cup B_1 \cup B_2 = \Omega$ и $B_k B_l = \emptyset$ ($k \neq l$). Можно воспользоваться формулой полной вероятности (п2.18):

$$P(A) = P(B_0) P(A|B_0) + P(B_1) P(A|B_1) + P(B_2) P(A|B_2)$$

Если произошло событие B_k , то в ящике из 10 изделий k бракованных, следовательно,

$$P(A|B_k) = \frac{10-k}{10}, \quad k = 0, 1, 2.$$

При выборе малого числа изделий из большой партии схемы выбора без возвращения и с возвращением имеют близкие вероятности исходов (см. задачу п18). Считая поэтому, что каждое из добавленных изделий независимо от другого может быть бракованным с вероятностью 0,05, находим

$$P(B_0) = 0,95^2, \quad P(B_1) = 2 \cdot 0,95 \cdot 0,05, \quad P(B_2) = 0,05^2.$$

Таким образом,

$$P(A) = 0,95^2 \cdot 1 + 2 \cdot 0,95 \cdot 0,05 \cdot \frac{9}{10} + 0,05^2 \cdot \frac{8}{10} = 0,99.$$

Пример

Ч-2.пр.9

Из урны, содержащей 3 белых и 2 черных шара, по схеме выбора без возвращения отобраны 2 шара. Шар, взятый наудачу из этих двух, оказался белым. Какова вероятность того, что второй шар тоже белый? Р. $B_k = \{ \text{среди двух отобранных из урны шаров } k \text{ белых} \}$, $k = 0, 1, 2$, $A = \{ \text{шар, взятый наудачу из двух отобранных, белый} \}$. Условную вероятность $P(B_2|A)$, которую требуется найти, можно вычислить по формуле Байеса (Ч-2.19): $P(B_2|A) = \frac{P(B_2) P(A|B_2)}{P(B_0) P(A|B_0) + P(B_1) P(A|B_1) + P(B_2) P(A|B_2)}$. Вероятности, связанные с извлечением двух шаров из урны, найдем с помощью классического определения вероятности:

$$P(B_0) = \frac{C_2^2}{C_5^2} = 0,1, \quad P(B_1) = \frac{C_3^1 C_2^1}{C_5^2} = 0,6, \quad P(B_2) = \frac{C_3^2}{C_5^2} = 0,3.$$

Так же определяются условные вероятности, связанные с извлечением одного шара из двух отобранных:

$$\text{Значит, } P(A|B_0) = 0, \quad P(A|B_1) = 0,5, \quad P(A|B_2) = 1. \\ P(B_2|A) = \frac{0,3}{0,1 \cdot 0 + 0,6 \cdot 0,5 + 0,3 \cdot 1} = \frac{1}{2}.$$

В задачах п2.1-п2.6 вероятностное пространство считается заданным; Для вычисления условных вероятностей нужно использовать формулу (п2.1).

Ч-2.1

Из множества чисел $\{1, 2, \dots, N\}$ по схеме случайного выбора без возвращения выбираются три числа. Найти условную вероятность того, что третье число попадет в интервал, образованный первыми двумя, если известно, что первое число меньше второго.

Ч-2.2

Брошено две игральные кости. Предполагается, что все комбинации выпавших очков равновероятны. Найти условную вероятность того, что выпали две пятерки, если известно, что сумма выпавших очков делится на пять.

Ч-2.3

Из 100 карточек с числами 00, 01, ..., 98, 99 случайно выбирается одна. Пусть η_1 и η_2 - соответственно сумма и произведение цифр на выбранной карточке. Найти $P\{\eta_1 = i | \eta_2 = 0\}, i = 0, 1, \dots, 18$.

Ч-2.4

Из урны, содержащей M белых и $N - M$ черных шаров, последовательно без возвращения извлекают n шаров. Пусть $A_0^{(i)}(A_1^{(i)})$ - событие, состоящее в том, что i -й шар был черным (белым). Используя классическое определение случайного выбора (гл. 1), найти

$$P\{A_1^{(s+1)} | A_{\varepsilon_1}^{(1)} A_{\varepsilon_2}^{(2)} \dots A_{\varepsilon_s}^{(s)}\}, \quad \varepsilon_i = 0, 1.$$

Ч-2.5

Решить задачу п2.4 в случае выбора с возвращением (см. введение к гл. 1, (п3)).

Ч-2.6

Случайный выбор двух подмножеств A_1 и A_2 из множества $S = \{1, 2, \dots, N\}$ производится так же, как в задаче п2.3. Найти условную вероятность $P\{|A_1| = l_1, |A_2| = l_2 | A_1 \cap A_2 = \emptyset\}$ того, что множества A_1 и A_2 состоят из l_1 и l_2 элементов соответственно при условии, что A_1 и A_2 не пересекаются.

В задачах п2.7-п2.11 предполагаются заданными условные вероятности; при решении используются формулы (п2.2), (п2.3).

Ч-2.7

Среди 25 экзаменационных билетов 5 «хороших». Два студента по очереди берут по одному билету. Найти вероятность того, что:

- первый студент взял «хороший» билет,
- второй студент взял «хороший» билет,
- оба студента взяли «хорошие» билеты.

Ч-2.8

Два игрока поочередно извлекают шары (без возвращения) из урны, содержащей M белых и $N - M$ черных шаров. Выигрывает тот, кто первым вынет белый шар. Используя формулу (п2.3), найти вероятность выигрыша первого участника, если: а) $N = 4, M = 1$; б) $N = 5, M = 1$; в) $N = 7, M = 2$.

Ч-2.9

Из урны, содержащей M белых и $N - M$ черных шаров, по одному без возвращения извлекаются все шары. Используя определение случайного выбора в терминах условных вероятностей, найти вероятности событий $A_k = \{k\text{-й шар белый}\}, B_{h,l} = \{k\text{-й и } l\text{-й шары белые}\}, C_{A_1 1} = \{k\text{-й шар черный, а } l\text{-й белый}\}.$

Ч-2.10

Из урны, содержащей 3 белых шара, 5 черных и 2 красных, два игрока поочередно извлекают по одному шару без возвращения. Выигрывает тот, кто первым вынет белый шар. Если появляется красный шар, то объявляется ничья. Пусть $A_1 = \{\text{выигрывает игрок, начавший игру}\}, A_2 = \{\text{выигрывает второй участник}\}, B = \{\text{игра закончилась вничью}\}.$ Найти $P(A_1), P(A_2), P(B).$

Ч-2.13

Из урны, содержащей a белых и b черных шаров, три игрока извлекают шары по очереди. Выигрывает тот, кому раньше попадается белый шар. Найти вероятности P_1, P_2, P_3 выигрыша 1-го, 2-го, 3-го игроков в случаях, когда шары извлекаются:


а) по схеме равновероятного выбора с возвращением, б) по схеме равновероятного выбора без возвращения.

Ч-2.14

На остановку прибывают автобусы маршрутов $1, 2, \dots, k$. Номера последовательно прибывающих автобусов получаются по схеме равновероятного выбора с возвращением из урны, содержащей шары с номерами $1, 2, \dots, k$. Найти вероятность P_k того, что до появления автобуса маршрута 1 ни на одном из остальных маршрутов не придет более одного автобуса. При каком минимальном $k \geq 2$ эта вероятность меньше $1/2$?

О дорогах (?)

На рисунке изображена схема дорог. Туристы выходят из пункта O , выбирая наугад на разветвлении дорог один из возможных путей. Какова вероятность того, что туристы попадут в пункт M ?



Решение: Путь туристов обязательно проходит через один из промежуточных пунктов B_1, B_2, B_3 .

Пусть H_i - событие, состоящее в том, что при своем движении туристы попадут в пункт B_i ($i = 1, 2, 3$). События H_1, H_2, H_3 по условию несовместны и равновероятны, т. к. путь из точки O выбирается наугад.

Поэтому события H_1, H_2, H_3 образуют полную группу, т. к. $P(H_1) = P(H_2) = P(H_3) = \frac{1}{3}$, и $\sum_{i=1}^3 P(H_i) = 1$.

Если событие A - туристы попадут в пункт M , то $P(A/H_1) = 1/3$, $P(A/H_2) = 1/2$, $P(A/H_3) = 1/4$. Искомую вероятность находим по формуле полной вероятности:

$$\begin{aligned} P(A) &= P(H_1) \cdot P(A/H_1) + P(H_2) \cdot P(A/H_2) + P(H_3) \cdot P(A/H_3) = \\ &= \frac{1}{3} \cdot \frac{1}{3} + \frac{1}{3} \cdot \frac{1}{2} + \frac{1}{3} \cdot \frac{1}{4} = \frac{13}{36}. \end{aligned}$$

Билеты экзамена

Из n -экзаменационных билетов студент знает m («хорошие билеты»). Что лучше: брать на экзамене билет первым или вторым? Решение: Введем событие A - студент взял «хороший» билет. 1. Студент берет билет первым. В этом случае $P(A) = \frac{m}{n}$. 2. Студент берет билет вторым. Введем две гипотезы: H_1 - первый студент взял «хороший» билет, $H_2 = \overline{H_1}$.

События H_1, H_2 образуют полную группу, т. к. $\sum_{i=1}^2 P(H_i) = 1$. $P(A) = P(H_1) P(A/H_1) + P(H_2) P(A/H_2) = \frac{m}{n} \cdot \frac{m-1}{n-1} + (1 - \frac{m}{n}) \cdot \frac{m}{n-1} = \frac{m}{n}$. Вывод: безразлично, брать билет первым или вторым. Пример 3. В первой урне находится 6 белых и 4 черных шара, во второй - 5 белых и 5 черных шаров. Из первой урны наугад берут шар и перекладывают во вторую. Затем из второй урны достают один шар. Какова вероятность того, что он белый? Решение:

$$\begin{aligned} P(A) &= P(H_1) \cdot P(A/H_1) + P(H_2) \cdot P(A/H_2) + P(H_3) \cdot P(A/H_3) = \\ &= \frac{1}{3} \cdot \frac{3}{8} + \frac{1}{3} \cdot \frac{1}{2} + \frac{1}{3} \cdot 0 = \frac{1}{8} + \frac{1}{6} = \frac{7}{24}. \end{aligned}$$

Пример 5.

На фабрике первая машина производит 25%, вторая 35, а третья 40 деталей. Брак в их продукции равен, соответственно, 5, 4 и 3%. Какова вероятность, что случайно выбранная деталь дефектна? Решение: Событие A - выбранная деталь имеет дефект. Обозначим через H_i событие, состоящее в том, что случайно выбранная деталь изготовлена на i -й машине ($i = 1, 2, 3$).

Тогда $P(H_1) = 0,25$; $P(H_2) = 0,35$; $P(H_3) = 0,40$, и события H_1, H_2, H_3 образуют полную группу, если учесть, что других машин на фабрике нет и каждую деталь изготавливает только одна машина. Пусть $P(A)$ - искомая вероятность:

$$P(A) = P(H_1)P(A/H_1) + P(H_2)P(A/H_2) + P(H_3)P(A/H_3) = \\ = 0,25 \cdot 0,05 + 0,35 \cdot 0,04 + 0,4 \cdot 0,03 = 0,0385.$$

Пример 6.

Вам последовательно сдается три карты из колоды в 36 карт. Какова вероятность того, что третья карта по масти не совпадает с первыми двумя? Решение: Обозначим A - третья карта по масти отличается от первых двух. В данной задаче удобно выделить следующие две гипотезы: H_1 - первые две карты одинаковы по масти, H_2 - первые две карты различаются по масти ($H_2 = \overline{H_1}$). События H_1, H_2 образуют полную группу, т. к. $\sum_{i=1}^2 P(H_i) = 1$. Заметим, что

$$H_1 = \text{ЧЧ} + \text{ББ} + \text{ТТ} + \text{ПП},$$

где ЧЧ означает, что две первые карты червы, ББ - они бубны, ТТ - они трефы, ПП - они пики.

Вероятность вытащить две червы равна $P(\text{ЧЧ}) = \frac{C_2^9}{C_2^{36}} = \frac{9 \cdot 8}{36 \cdot 35} = \frac{2}{35}$. Аналогично, $P(\text{ББ}) = P(\text{ТТ}) = P(\text{ПП}) = \frac{2}{35}$. Следовательно,

$$P(H_1) = 4 \cdot \frac{2}{35} = \frac{8}{35}.$$

Если две первые карты одинаковы по масти, то при вытаскивании третьей карты у нас всего 34 карты, а карт, не совпадающих по масти с первыми двумя, 27. Следовательно, $P(A/H_1) = \frac{27}{34}$. Так как $H_2 = \overline{H_1}$, то $P(H_2) = 1 - P(H_1) = \frac{27}{35}$. Если две первые карты различны по масти, то при вытаскивании третьей карты у нас всего 34 карты, а карт, не совпадающих по масти с первыми двумя, 18. Следовательно, $P(A/H_2) = \frac{18}{34}$. По формуле полной вероятности

$$P(A) = P(A/H_1)P(H_1) + P(A/H_2)P(H_2) = \frac{27}{34} \cdot \frac{8}{35} + \frac{18}{34} \cdot \frac{27}{35} = \frac{26}{34} \cdot \frac{27}{35} = 0,59.$$

Пример 7.

По цели произведено три последовательных выстрела. Вероятность попадания при первом выстреле $p_1 = 0,3$; при втором $p_2 = 0,6$; при третьем $p_3 = 0,8$. При одном попадании вероятность поражения цели $r_1 = 0,4$; при двух попаданиях $r_2 = 0,7$; при трех попаданиях $r_3 = 1$. Определить вероятность поражения цели при трех выстрелах. Решение: Пусть событие A - цель поражена. Рассмотрим группу несовместных событий: H_1 - было одно попадание; H_2 - было два попадания; H_3 - было три попадания; H_4 - не было ни одного попадания. Определим вероятность каждого события. По теоремам умножения и сложения вероятностей будем иметь

$$P(H_1) = p_1(1-p_2)(1-p_3) + (1-p_1)p_2(1-p_3) + (1-p_1)(1-p_2)p_3 = 0,332.$$

$$P(H_2) = p_1p_2(1-p_3) + p_1(1-p_2)p_3 + (1-p_1)p_2p_3 = 0,468.$$

$$P(H_3) = p_1p_2p_3 = 0,144.$$

$$P(H_4) = (1-p_1)(1-p_2)(1-p_3) = 0,056.$$

События H_1, H_2, H_3, H_4 образуют полную группу, т. к. $\sum_{i=1}^4 P(H_i) = 1$. Выпишем условные вероятности поражения цели при осуществлении каждого из событий H_1, H_2, H_3 , и H_4 :

$$P(A/H_1) = 0,4, P(A/H_2) = 0,7, P(A/H_3) = 1, P(A/H_4) = 0.$$

Тогда по формуле полной вероятности

$$P(A) = P(H_1)P(A/H_1) + P(H_2)P(A/H_2) + P(H_3)P(A/H_3) + P(H_4)P(A/H_4) = \\ = 0,332 \cdot 0,4 + 0,468 \cdot 0,7 + 0,144 \cdot 1 + 0,056 \cdot 0 = 0,6044.$$

Пример 8.

Известно, что 96 % выпускаемой продукции удовлетворяет стандарту. Упрощенная схема контроля признает пригодной хорошую продукцию с вероятностью 0,98, а бракованную - с вероятностью 0,05. Определить вероятность того, что изделие, прошедшее упрощенный контроль, удовлетворяет стандарту. Решение: Событие A - изделие прошло упрощенный контроль, H_1 - изделие удовлетворяет стандарту, H_2 - изделие бракованное. $P(H_1) = 0,96$; $P(H_2) = 0,04$; и события H_1, H_2 образуют полную группу. $P(A/H_1) = 0,98$; $P(A/H_2) = 0,05$. $P(A) = P(H_1)P(A/H_1) + P(H_2)P(A/H_2) = 0,96 \cdot 0,98 + 0,04 \cdot 0,05 = 0,9428$. $P(H_1/A)$ - изделие удовлетворяет стандарту, если оно прошло упрощенный контроль. По формуле Байеса

$$P(H_1/A) = \frac{P(H_1) \cdot P(A/H_1)}{P(A)} = \frac{0,96 \cdot 0,98}{0,9428} = 0,9979.$$

Пример 9.

В первой урне находится 6 белых и 4 черных шара, во второй - 5 белых и 5 черных шаров. Из первой урны наугад берут шар и перекладывают во вторую. Затем из второй урны достают один шар. Он оказался белым. Какова вероятность того, что из первой урны во вторую переложили белый шар? Решение: Событие A - извлечен белый шар из второй урны. Введем гипотезы H_1 - достали белый шар из первой урны, H_2 - достали черный шар из первой урны. Надо найти $P(H_1/A)$.

$$P(H_1) = \frac{6}{10}, P(A/H_1) = \frac{5+1}{10+1} = \frac{6}{11}, P(H_2) = \frac{4}{10}, P(A/H_2) = \frac{5}{10+1} = \frac{5}{11}, \text{ и}$$

события H_1, H_2 образуют полную группу.

По формуле Байеса

$$P(H_1/A) = \frac{P(A/H_1)P(H_1)}{P(A/H_1)P(H_1) + P(A/H_2)P(H_2)} = \frac{\frac{6}{11} \cdot \frac{6}{10}}{\frac{6}{11} \cdot \frac{6}{10} + \frac{5}{11} \cdot \frac{4}{10}} = \frac{36}{56}.$$

Пример 10.

Один из трех стрелков вызывается на линию огня и производит два выстрела. Вероятность попадания в мишень при одном выстреле для первого стрелка равна 0,3; для второго - 0,5; для третьего - 0,8. Мишень не поражена. Найти вероятность того, что выстрелы произведены первым стрелком. Решение: Возможны три гипотезы: H_1 - «на линию огня вызван первый стрелок», H_2 - «на линию огня вызван второй стрелок», H_3 - «на линию огня вызван третий стрелок». Так как вызов на линию огня любого стрелка равновозможен, то

$$P(H_1) = P(H_2) = P(H_3) = 1/3.$$

В результате опыта наблюдалось событие A - после произведенных выстрелов мишень не поражена.

Условные вероятности этого события при сделанных гипотезах равны

$$P(A/H_1) = 0,7 \cdot 0,7 = 0,49; P(A/H_2) = 0,5 \cdot 0,5 = 0,25; P(A/H_3) = 0,2 \cdot 0,2 = 0,04.$$

По формуле Байеса находим вероятность гипотезы H_1 после опыта:

$$P(H_1/A) = \frac{0,49 \cdot 1/3}{1/3 \cdot 0,49 + 1/3 \cdot 0,25 + 1/3 \cdot 0,04} = \frac{0,49}{0,78} = 0,628.$$

Пример 11.

По линии связи с вероятностями 0,84 и 0,16 соответственно передаются два сигнала A и B . Из-за помех $1/6$ часть переданных сигналов A искажается и принимается как сигнал B , а $1/8$ часть переданных сигналов B принимается как сигнал A . 1) Какова вероятность того, что при приеме появится: а) сигнал A ; б) сигнал B ? 2) Известно, что принят сигнал A . Какова вероятность того, что он же и был передан? Решение: Введем обозначения для событий. H_1 - передан сигнал A , H_2 - передан сигнал B , A - при приеме появится сигнал A , B - при приеме появится сигнал B . Имеем

$$P(H_1) = 0,84; P(H_2) = 0,16; P(A/H_1) = 5/6;$$

$$P(A/H_2) = 1/8; P(B/H_1) = 1/6; P(B/H_2) = 7/8.$$

1. По формуле полной вероятности находим а) $P(A) = P(H_1) \cdot P(A/H_1) + P(H_2) \cdot P(A/H_2) = 0,84 \cdot \frac{5}{6} + 0,16 \cdot \frac{1}{8} = 0,72$;

б) $P(B) = P(H_1) \cdot P(B/H_1) + P(H_2) \cdot P(B/H_2) = 0,84 \cdot \frac{1}{6} + 0,16 \cdot \frac{7}{8} = 0,28$. 2. По формуле Байеса получаем

$$P(H_1/A) = \frac{P(H_1) \cdot P(A/H_1)}{P(A)} = \frac{0,84 \cdot \frac{5}{6}}{0,72} \approx 0,97.$$

Пример 12.

Два из трех независимо работающих элементов вычислительного устройства отказали. Найти вероятность того, что отказали первый и второй элементы, если вероятности отказа первого, второго и третьего элементов соответственно равны $p_1 = 0,2$; $p_2 = 0,4$; $p_3 = 0,3$. Решение: Введем обозначение A для события - отказали два элемента. Рассмотрим следующие гипотезы: H_1 - отказали первый и второй элементы, а третий элемент исправен; H_2 - отказали первый и третий элементы, а второй элемент исправен; H_3 - отказали второй и третий элементы, а первый элемент исправен; H_4 - ни один элемент не отказал; H_5 - отказал только один элемент; H_6 - отказали все три элемента. При этом (поскольку элементы работают независимо, то применима теорема умножения вероятностей):

$$P(H_1) = p_1 p_2 (1 - p_3) = 0,2 \cdot 0,4 \cdot 0,7 = 0,056;$$

$$P(H_2) = p_1 p_3 (1 - p_2) = 0,2 \cdot 0,3 \cdot 0,6 = 0,036;$$

$$P(H_3) = p_2 p_3 (1 - p_1) = 0,4 \cdot 0,3 \cdot 0,8 = 0,096.$$

Вероятности гипотез H_4, H_5, H_6 можно не считать, т. к. из условия задачи следует, что

$$P(A/H_4) = 0, P(A/H_5) = 0, P(A/H_6) = 0,$$

следовательно, равны нулю и произведения $P(H_4) \cdot P(A/H_4), P(H_5) \cdot P(A/H_5), P(H_6) \cdot P(A/H_6)$, входящие в формулу полной вероятности.

При гипотезах H_1, H_2, H_3 событие A достоверно, соответствующие условные вероятности равны единице. По формуле полной вероятности находим вероятность события A :

$$P(A) = \sum_{i=1}^6 P(H_i) \cdot P(A/H_i) = P(H_1) P(A/H_1) + P(H_2) P(A/H_2) + P(H_3) \times \\ \times P(A/H_3) = 0,056 \cdot 1 + 0,036 \cdot 1 + 0,096 \cdot 1 = 0,188.$$

Теперь искомую вероятность вычисляем по формуле Байеса:

$$P(H_1/A) = \frac{P(H_1) \cdot P(A/H_1)}{P(A)} = \frac{0,056 \cdot 1}{0,188} \approx 0,298.$$

6.2.2 Problems About the Independence of Events

- В задачах п2.15-п2.19 предполагается, что задано вероятностное пространство; требуется выяснить, зависимы или независимы некоторые события.

Ч-2.15

Брошены две игральные кости. Положим $A_l = \{ \text{число очков, выпавшее на первой кости, делится на } l \}$, $R_l = \{ \text{число очков, выпавшее на второй кости, делится на } l \}$, $C_l = \{ \text{сумма очков, выпавших на первой и второй костях, делится на } l \}$. Отправляясь от классического определения вероятности, установить, являются ли независимыми следующие пары событий а) A_l, B_k - при любых l, k ; б) A_2, C_2 ; в) A_4, C_4 ?

Ч-2.16

Игральная кость брошена 2 раза, X_1 и X_2 - числа очков, выпавшие при этих испытаниях. Рассмотрим события

$$A_1 = \{X_1 \text{ делится на } 2, X_2 \text{ делится на } 3\},$$

$$A_2 = \{X_1 \text{ делится на } 3, X_2 \text{ делится на } 2\},$$

$$A_3 = \{X_1 \text{ делится на } X_2\},$$

$$A_4 = \{X_2 \text{ делится на } X_1\},$$

$$A_5 = \{X_1 + X_2 \text{ делится на } 2\},$$

$$A_6 = \{X_1 + X_2 \text{ делится на } 3\}.$$

Найти все пары $\{A_i, A_j\}$, тройки $\{A_i, A_j, A_k\}$ и т. д. взаимно независимых событий.

Ч-2.17

Случайная точка (ξ_1, ξ_2) имеет равномерное распределение в квадрате $\{(x_1, x_2) : 0 \leq x_1, x_2 \leq 1\}$. При каких значениях r независимы события $A_r = \{|\xi_1 - \xi_2| \geq r\}$ и $B_r = \{\xi_1 + \xi_2 \leq 3r\}$?

Ч-2.18

События A и B независимы. Являются ли независимыми события: а) A и $\bar{B}$, б) $\bar{A}$ и $\bar{B}$?

Ч-2.19

Случайная точка $\xi = (\xi_1, \xi_2)$ имеет равномерное распределение в квадрате $0 \leq x, y \leq 1$. Пусть

$$A_1 = \left\{ \xi_1 \leq \frac{1}{2} \right\}, \quad A_2 = \left\{ \xi_2 \leq \frac{1}{2} \right\}$$

$$A_3 = \left\{ \left(\xi_1 - \frac{1}{2} \right) \left(\xi_2 - \frac{1}{2} \right) < 0 \right\}$$

Показать, что любые два события из A_1, A_2, A_3 независимы, но все три события A_1, A_2, A_3 зависимы. Являются ли зависимыми события $A_1 A_2$ и A_3 ?

Ч-2.20

(см. п.2.19). Обобщая пример, приведенный в предыдущей задаче, показать, что для любого целого $n \geq 4$ существует совокупность $\{A_1, \dots, A_n\}$ событий, обладающая следующими свойствами:

- события $A_1, \dots, A_n$ не являются независимыми,
- при удалении из $A_1, \dots, A_n$ любого события оставшаяся совокупность состоит из независимых событий.

Ч-2.21

События $A_1, A_2, \dots, A_n$ удовлетворяют условиям

$$P\{A_i\} = p_i, \quad P\left\{\bigcap_{j=1}^i A_j\right\} = p_1 \dots p_i, \quad i = 1, 2, \dots, n.$$

Является ли $\{A_1, \dots, A_n\}$ совокупностью независимых событий?

Ч-2.22

*. Пространство элементарных событий Ω состоит из n элементов. При каких k на подмножествах Ω можно определить вероятность P и события $A_1, \dots, A_n$ так, чтобы события $A_1, \dots, A_k$ были независимыми в совокупности и $0 < P(A_i) < 1$ ($i = 1, 2, \dots, k$)?

Ч-2.23

Пространство элементарных событий Ω состоит из $n \geq 4$ элементов. При каких k можно так определить на подмножествах Ω вероятность P и события $A_1, \dots, A_k$, что $0 < P(A_i) < 1$ ($i = 1, \dots, k$) и события $A_1, \dots, A_k$ попарно независимы?

Ч-2.28

По цели производится n независимых выстрелов. Вероятность попадания при i -м выстреле равна p_i , $i = 1, \dots, n$. Найти вероятность того, что при n выстрелах будет не менее двух попаданий.

6.2.3 Problems About the Total Probability Formula**Ч-2.29**

В первой урне находятся 1 белый и 9 черных шаров. а во второй - 1 черный и 5 белых шаров. Из каждой урны по схеме случайного выбора без возвращения удалили по одному шару, а оставшиеся шары высыпали в третью урну. Найти вероятность того, что шар, вынутый из третьей урны, окажется белым.

Ч-2.30

В первой урне лежит 1 белый шар и 4 красных, а во второй - 1 белый и 7 красных. В первую урну до-

бавляются два шара, случайно выбранных из второй урны.

- Найти вероятность того, что шар, выбранный из пополненной первой урны, будет белым.
- Пусть из пополненной первой урны по схеме случайного выбора с возвращением извлекают k шаров. Найти вероятность того, что все они будут белыми.

Ч-2.31

Из урны, содержащей 2 белых и 3 черных шара, наудачу извлекают 2 шара и добавляют в урну один белый шар.

а) Найти вероятность того, что после этого наудачу выбранный из урны шар окажется белым. б) Пусть из урны по схеме случайного выбора с возвращением извлекают k шаров. Найти вероятность того, что все они белые.

в) Найти ту же вероятность, что в п. б), для схемы выбора без возвращения.

Ч-2.32

В пункте проката имеется 10 телевизоров, для которых вероятность исправной работы в течение месяца равна 0,90, и 5 телевизоров с аналогичной вероятностью, равной 0,95. Найти вероятность того, что два телевизора, взятые наудачу в пункте проката, будут работать исправно в течение месяца.

Ч-2.33

В одной урне содержится 1 белый и 2 черных шара, а в другой урне - 2 белых и 3 черных шара. В третью урну кладут два шара, случайно выбранных из первой урны, и два шара, случайно выбранных из второй.

а) Какова вероятность того, что шар, извлеченный из третьей урны, будет белым?

б) Найти вероятность того, что при выборе с возвращением из третьей урны двух шаров один из них будет белым, а другой - черным.

в) Найти ту же вероятность, что в п. б), для схемы выбора без возвращения.

Ч-2.37

При переливании крови надо учитывать группы крови донора и больного. Человеку, имеющему четвертую группу крови, можно перелить кровь любой группы; человеку со второй или третьей группой крови можно перелить кровь либо той же группы, либо первой; человеку с первой группой крови можно перелить только кровь первой группы. Среди населения 33,7 % имеют первую, 37,5 % - вторую, 20,9 % - третью и 7,9 % - четвертую группы крови.

а) Найти вероятность того, что случайно взятому больному можно перелить кровь случайно взятого донора.

б) Найти вероятность того, что переливание можно осуществить, если имеются два донора; три донора.

Ч-2.38

Во время испытаний было установлено, что вероятность безотказного срабатывания реле при отсутствии помех равна 0,99, при перегреве - 0,95, при вибрации - 0,9, при вибрации и перегреве - 0,8. Найти вероятность P_1 отказа этого реле при работе в жарких странах (вероятность перегрева равна 0,2, вероятность вибрации 0,1) и вероятность P_2 отказа при работе в передвижной лаборатории (вероятность перегрева 0,1, вероятность вибрации 0,3), предполагая перегрев и вибрацию независимыми событиями.

Ч-2.39

(см. п.2.38). Найти границы, в которых могут изменяться вероятности P_1 и P_2 в предыдущей задаче, если отказаться от предположения о независимости перегрева и вибрации.

Ч-2.40

Имеется пять урн. В 1-й, 2-й и 3-й урнах находится по 2 белых и 3 черных шара; в 4-й и 5-й урнах - по 1 белому и 1 черному шару. Случайно выбирается урна и из нее извлекается шар. Какова условная вероятность того, что выбрана 4-я или 5-я урна, если извлеченный шар оказался белым?

6.2.4 Problems About Bernoulli's Scheme (!?!?!?!?)

(очень важно это прорешать будет!)

Пр и м е р 3.

Определить вероятность того, что в семье, имеющей 5 детей, будет не больше трех девочек. Вероятности рождения мальчика и девочки предполагаются одинаковыми. Решение: Вероятность

рождения девочки $p = \frac{1}{2}$, тогда $q = \frac{1}{2}$. Найдём вероятности того, что в семье нет девочек, родилась одна, две или три девочки:

$$P_5(0) = q^5 = \frac{1}{32}, P_5(1) = C_5^1 p^1 q^4 = \frac{5}{32}, \quad P_5(2) = C_5^2 p^2 q^3 = \frac{10}{32},$$

$$P_5(3) = C_5^3 p^3 q^2 = \frac{10}{32}.$$

Следовательно, искомая вероятность

$$P_5(k \leq 3) = P_5(0) + P_5(1) + P_5(2) + P_5(3) = \frac{13}{16}$$

Ч-2.46

Проведено 20 независимых испытаний, каждое из которых заключается в одновременном подбрасывании трех монет. Найти вероятность того, что хотя бы в одном испытании появятся три «герба».

Ч-2.47

При передаче сообщения вероятность искажения одного знака равна $1/10$. Каковы вероятности того, что сообщение из 10 знаков: а) не будет искажено, б) содержит ровно 3 искажения, в) содержит не более трех искажений?

Ч-2.48

Испытание заключается в бросании трех игральных костей. Найти вероятность того, что в пяти независимых испытаниях ровно два раза выпадет по три единицы.

Ч-2.49

Найти вероятность того, что в $2n$ испытаниях схемы Бернулли с вероятностью успеха p и неудачи $q = 1 - p$ появится $m + n$ успехов и все испытания с четными номерами закончатся успехом.

Ч-2.50

Из множества $S = \{1, 2, \dots, N\}$ случайно и независимо выбирают два подмножества A_1 и A_2 так, что каждый элемент из S независимо от других элементов с вероятностью $q = 1 - p$ не включается. Найти вероятность того, что $A_1 \cap A_2 = \emptyset$.

Ч-2.51

. По той же схеме выбора подмножеств на $S = \{1, 2, \dots, N\}$, что в задаче п2.50, независимо выбираются r подмножеств $A_1, A_2, \dots, A_r, r \geq 2$. Найти вероятность того, что выбранные подмножества попарно не пересекаются.

Ч-2.52

(см. п50). Из множества $S = \{1, 2, \dots, N\}$ независимо выбираются r подмножеств $A_1, A_2, \dots, A_r$. Механизм выбора состоит в следующем: любой элемент множества S независимо от других элементов с вероятностью p_i включается в множество A_i и с вероятностью $q_i = 1 - p_i$ не включается ($i = 1, \dots, r$). Найти вероятность того, что подмножества $A_1, A_2, \dots, A_r$ попарно не пересекаются.

Ч-2.55

В одном из матчей на первенство мира по шахматам партии не учитывались, и игра шла до тех пор, пока один из участников матча не набирал 6 очков (выигрыш - 1 очко, проигрыш и ничья - 0 очков). Считая участников матча одинаковыми по силе, а результаты отдельных игр независимыми, найти вероятность того, что при таких правилах в момент окончания матча проигравший набирает k очков, $k = 0, 1, \dots, 5$.

Ч-2.56

Обрабатываемые на станке детали сортируются по размерам на две группы. Каждая очередная деталь независимо от предыдущих с равными вероятностями попадает в первую или вторую группу. Пусть в начал: смены для каждой группы деталей приготовлено по ящику емкости r . Какова вероятность того, что в момент, когда очередную деталь будет некуда класть, в другом ящике будет m деталей?

Бракованные детали в партии

В партии содержатся 50 деталей, из которых 10 бракованных. Из партии наудачу берутся 5 деталей. Найти вероятность того, что:

1) все 5 деталей бракованные; 2) все 5 деталей доброкачественные; 3) в пятерке извлеченных деталей 3 детали бракованные и 2 детали доброкачественные; 4) среди 5-ти извлеченных деталей хотя бы одна бракованная.

Solution

Будем считать, что все детали пронумерованы, т. е. они образуют множество из $n = 50$ различных по номерам объектов, из которых 10 - бракованные, а остальные 40 - доброкачественные. Из этой партии наудачу берутся 5 деталей. В этом состоит эксперимент.

Событие A - все 5 деталей бракованные, событие B - все 5 деталей доброкачественные, событие C - в пятерке извлеченных деталей 3 детали бракованные и 2 - доброкачественные, событие D - в пятерке извлеченных деталей число бракованных деталей от 1 до 5.

Элементарный исход нашего эксперимента определяется номерами пяти взятых деталей, причем порядок указания этих номеров не имеет значения. Общее число таких элементарных событий совпадает с числом различных сочетаний из 50 элементов по 5 и выражается формулой

$$C_{50}^5 = \frac{50!}{5!45!} = \frac{46 \cdot 47 \cdot 48 \cdot 49 \cdot 50}{120} = 2118760.$$

Все эти элементарные исходы образуют полную группу и равновозможны, т. к. 5 деталей берутся наудачу и одна пятерка деталей имеет такие же шансы быть взятой, что и любая другая.

Событию A благоприятствуют лишь те пять деталей, которые взяты из 10 бракованных деталей. Так как порядок взятых 5-ти деталей не играет роли, то их число равно $C_{10}^5 = \frac{10!}{5!5!}$ и, следовательно,

$$P(A) = \frac{\tilde{N}_{10}^5}{\tilde{N}_{50}^5} = \frac{6 \cdot 7 \cdot 8 \cdot 9 \cdot 10}{120 \cdot 2118760} \approx 0,0001.$$

Событию B благоприятствуют лишь пятерки деталей, которые взяты из 40 доброкачественных деталей. Поэтому их число равно C_{40}^5 , и

$$P(B) = \frac{\tilde{N}_{40}^5}{\tilde{N}_{50}^5} = \frac{36 \cdot 37 \cdot 38 \cdot 39 \cdot 40}{120 \cdot 2118760} \approx 0,311.$$

Событию C благоприятствуют лишь те элементарные исходы (пятерки деталей), которые содержат 3 бракованные и 2 доброкачественные детали. Их число (по правилу умножения) равно $C_{10}^3 C_{40}^2$ и, следовательно,

$$P(C) = \frac{C_{10}^3 C_{40}^2}{C_{50}^5} = \frac{8 \cdot 9 \cdot 10 \cdot 39 \cdot 40}{6 \cdot 2 \cdot 2118760} \approx 0,044.$$

Событию $\bar{D}$ благоприятствуют лишь те элементарные исходы (пятерки деталей), которые не содержат бракованные (все доброкачественные детали), т. е. $\bar{D} = B$, поэтому

$$P(D) = 1 - P(\bar{D}) = 1 - P(B) = 1 - 0,311 \approx 0,689$$

Ч-2.57

По каналу связи передаются сообщения из пулей и единиц. Из-за помех вероятность правильной передачи знака равна 0,55. Для повышения вероятности правильной передачи каждый знак сообщения повторяют n раз. Полагают, что последовательности из n принятых знаков в сообщении соответствуют знак, составляющий в ней большинство. Найти вероятность правильной передачи одного знака при n -кратном повторении, если $n = 5$.

Ч-2.58

(см п2.57). Подобрать n так, чтобы вероятность правильной передачи знака была не меньше 0,99.

Ч-2.59

По каналу связи передается 1000 знаков. Каждый аяак может быть искажен независимо от остальных с вероятностью 0,005. Найти приближенное значение вероятности того, что будет искажено не более трех знаков.

Ч-2.61

Пусть ξ_n - число успехов в n независимых испытаниях Бернулли с вероятностью успеха, равной $1/2$.
 а) С помощью теоремы Муавра - Лапласа найти приближенные значения

$$P\left\{\left|\xi_n - \frac{n}{2}\right| \leq \frac{\sqrt{n}}{2}\right\}, P\left\{\left|\xi_n - \frac{n}{2}\right| \geq \frac{\sqrt{n}}{2}\right\}$$

при $n = 100$. 40

б) Вычислить те же вероятности, что в п. а), с погрешностью 10^{-5} , используя уточненную формулу Стирлинга.

Сравнить результаты пп. а) и б) и истолковать их. п2.62*. Решить предыдущую задачу, полагая $n = 128$.

Ч-2.63

Найти приближенное значение вероятности того, что число "девяток" среди 10000 случайных писел заключено между 940 и 1060.

Ч-2.64

Из таблицы случайных чисел отбирают числа, делящиеся на 3, до тех пор, пока не наберется 1025 таких чисел. Найти приближенное значение вероятности того, что потребуется таблица, содержащая не меньше 2500 чисел.

Ч-2.65

Театр, вмещающий 1000 человек, имеет два разных входа. Около каждого из входов имеется свой гардероб. Сколько мест должно быть в каждом из гардеробов для того, чтобы в среднем в 99 случаях из 100 все зрители могли раздеться в гардеробе того входа, через который они вошли? Рассмотреть два случая: а) зрители приходят парами; б) зрители приходят поодиночке. Предположить, что входы зрители выбирают с равными вероятностями.

Ч-2.66

В поселке 2500 жителей. Каждый из них примерно 6 раз в месяц ездит на поезде в город, выбирая для поездок по случайным мотивам независимо от остальных. Какой наименьшей вместимостью должен обладать поезд, чтобы он переполнялся в среднем не чаще одного раза в 100 дней (поезд ходит раз в сутки).

Ч-2.67

Пусть η_N - суммарное число появлений «5» к «6» в N бросаниях игральной кости. При $N = 1800$ найти вероятность того, что $\eta_N \geq 620$.

Ч-2.68

Две монеты подбрасывают 4800 раз. Найти приближенное значение вероятности того, что событие "герб - герб" появится меньше 1140 раз.

Ч-2.69

Вероятность попадания в цель при одном выстреле равна 0,01. Найти приближенное значение вероятности того, что при 100 выстрелах будет не больше трех попаданий.

Ч-2.70

Из урны, содержащей 1 белый и 4 черных шара, по схеме случайного выбора с возвращением проводят 2500 извлечений шаров. Найти приближенное значение вероятности того, что число появлений белого шара заключено между 480 и 540.

Ч-2.71

На одной странице 2400 знаков. При типографском наборе вероятность искажения одного знака равна $1/800$. Найти приближенное значение вероятности того, что на странице не менее двух опечаток.

Ч-2.72

Вероятность успеха в каждом испытании схемы Бернулли равна p . Найти вероятность того, что k -й по порядку успех происходит при l -м испытании.

В задачах с п2.73 по п2.78 рассматриваются бесконечные последовательности испытаний. Воспользоваться частным случаем вероятностного пространства, который определяется формулами (п15) - (п2.17) при $N = 2$.

Ч-2.73

Две игральные кости бросают до выпадения 6 и хотя бы на одной из них. Найти вероятность того, что впервые "6" появляется при k -м бросании, $k = 1, 2, 3, \dots$

Ч-2.74

Двое по очереди бросают монету. Выигрывает тот, кто первым получит «герб». Найти вероятности событий:

- а) игра закончится до 4-го бросания;
- б) выиграет начавший игру (первый игрок);
- в) выиграет второй игрок.

Ч-2.75

В схеме Бернулли p - вероятность исхода 1 и $q = 1 - p$ - вероятность исхода 0. Найти вероятность того, что цепочка 00, состоящая из двух нулей подряд, появится раньше цепочки 01. В частности, вычислить эту вероятность при $p = 1/2$.

Ч-2.76

Пусть выполнены условия задачи п2.75. Найти вероятность того, что цепочка 00 (два нуля подряд) появится раньше цепочки 10. В частности, вычислить эту вероятность при $p \approx 1/2$.

Ч-2.77

Пусть выполнены условия задачи п2.75. Найти вероятность P_{0011} того, что цепочка 00 появится раньше цепочки 111. В частности, вычислить эту вероятность при $p = 1/2$.

6.2.5 Problems About Poisson Events (??)

(мб тут и не про них, посмотрим потом, пока так, выгрузил то, что встречалось.)

Белан-л5.1.

Стрельба по цели ведется до тех пор, пока суммарное число попаданий не достигнет N . Найдите математическое ожидание, дисперсию, а также функцию распределения вероятностей потребовавшегося числа выстрелов k , если выстрелы производятся независимо друг от друга и вероятность попадания в отдельной попытке равна p .

Задача 1

Число выстрелов, требуемое чтобы получить N попаданий можно представить в виде суммы $k = k_1 + k_2 + \dots + k_N$, где k_n - это число выстрелов которое потребовалось для достижения n -го попадания, если отсчитывать от предыдущего попадания (или от начала стрельбы, если речь идет об $n = 1$). В силу отсутствия памяти о процесса Бернулли слагаемые статистически независимы. Каждое из них при этом имеет геометрическую функцию распределения вероятностей, заданную формулой (4). Воспользовавшись результатами, изложенными на третьей лекции и уравнением (5) и (7), запишем для случая произвольной вероятности p успеха в отдельном испытании

$$\mu_k = N\mu_{k_1} = \frac{N}{p},$$

$$\sigma_k^2 = N\sigma_{k_1}^2 = N\frac{1-p}{p^2}.$$

Чтобы получить функцию распределения вероятностей $p(k)$ нужно умножить вероятность получить ровно $N - 1$ попаданий в серии из первых $k - 1$ выстрелов на вероятность попадания в k -м исходе, то есть

$$p(k) = pB(N - 1, k - 1, p) = \frac{(k - 1)!}{(N - 1)!(k - N)!} p^N (1 - p)^{k - N}$$

Этот результат известен как распределение Паскаля.

Белан-л5.2.

Найти среднее число бросков симметричной монетки необходимое чтобы каждая ее сторона выпала хотя бы два раза.

Задача 2 Подбросим монетку три раза. С вероятностью $1/4$ все три исхода будут одинаковы, и тогда среднее оставшееся число попыток, требуемое для накопления двух недостающих исходов, равно $2 + 2 = 4$. С вероятностью $3/4$ не все исходы первых трех попыток одинаковы, и тогда, чтобы получить недостающий исход второй раз, в среднем придется подбросить монетку еще 2 раза. В итоге получаем, что среднее число бросков симметричной монетки необходимое чтобы каждая ее сторона выпала хотя бы два раза равно

$$\frac{1}{4}(3 + 4) + \frac{3}{4}(3 + 2) = \frac{22}{4} = 5.5$$

Белан-л5.3.

Симметричный игральный кубик подбрасывают раз за разом. Найти среднее число попыток $\langle k \rangle$, необходимое, чтобы каждая из шести граней выпала хотя бы один раз. Найти также дисперсию требуемого числа попыток.

Белан-л5.4.

Сколько в среднем раз необходимо подбросить симметричную монетку, чтобы впервые встретить в последовательности исходов реализацию вида "орел, решка"?

Задача 4 Для соответствующей случайной величины справедливо

$$k = k_1^o + k_1^p$$

а значит

$$\langle k \rangle = \langle k_1^o \rangle + \langle k_1^p \rangle = \frac{1}{p} + \frac{1}{1-p} = \frac{1}{p(1-p)}.$$

В частности, для симметричной монетки находим, что для появления пары "орел, решка" необходимо в среднем подкинуть монетку $\langle k \rangle = 4$ раз. Отметим, что среднее число попыток, необходимое для появления пары "орел,орел"равно 6.

Белан-л5.5.

Сколько в среднем раз необходимо подбросить симметричную монетку, чтобы впервые встретить в последовательности исходов реализацию вида "десять орлов подряд"?

Задача 5 Для ответа на поставленный вопрос, рассмотрим схему Бернулли, где единица соответствует выпадению орла, а ноль - выпадению решки. Будем решать задачу в общем случае, когда вероятность выпадения единицы в отдельном испытании равна p , и речь идет об ожидании цепочки исходов из N единиц подряд. Количество попыток, требуемое для реализации первой такой цепочки обозначим через k . Каждая последовательность независимых повторений эксперимента Бернулли характеризуется некоторым количеством нулей, выпавших до первого появления единицы, и количеством единиц, выпавших подряд сразу вслед за этим. Если число идущих подряд единиц в первой встретившейся цепочке единиц меньше чем N , то можно считать, что процесс ожидания цепочки из N единиц подряд начинается заново в том смысле, что оставшееся число попыток k' необходимое для выпадения N единиц подряд имеет ту же статистику, что и k . Сказанное может быть записано в форме уравнения

$$k = k_1^o + (N - 1) \cdot I(k_1^p > N - 1) + (k_1^p + k') I(k_1^p \leq N - 1)$$

где нами были введены следующие обозначения: k_1^o - номер попытки в цепочке независимых испытаний Бернулли, когда впервые выпадает; единица, k_1^p - номер попытки в цепочке независимых испытаний Бернулли, когда впервые выпадает ноль, k' - независимая реализация случайной величины k , $I(\dots)$ - индикаторная функция. Рекурсивное соотношение (43) представляет собой пример того, что называется уравнением обновления.

Произведем статистическое усреднение

$$\langle k \rangle = \langle k_1^o \rangle + (N - 1) \cdot \langle I(k_1^p > N - 1) \rangle + \langle k_1^p I(k_1^p \leq N - 1) \rangle + \langle k' I(k_1^p \leq N - 1) \rangle.$$

Словом, мы приравняли друг другу математические ожидания левой и парвой частей Ур. (43). В силу отсутствия памяти у процесса Бернулли, k_1^p и k' это статистически независимые случайные величины, а значит $\langle k' I(k_1^p \leq N - 1) \rangle = \langle k' \rangle \langle I(k_1^p \leq N - 1) \rangle = \langle k \rangle \langle I(k_1^p \leq N - 1) \rangle$, и мы можем переписать

$$\langle k \rangle = \langle k_1^o \rangle + (N - 1) \cdot \langle I(k_1^p > N - 1) \rangle + \langle k_1^p I(k_1^p \leq N - 1) \rangle + \langle k \rangle \langle I(k_1^p \leq N - 1) \rangle$$

Решая это уравнение относительно $\langle k \rangle$, находим

$$\langle k \rangle = N - 1 + \frac{\langle k_1^o \rangle + \langle k_1^p I(k_1^p \leq N - 1) \rangle}{\langle I(k_1^p > N - 1) \rangle}.$$

Остается лишь вычислить средние значения, входящие в правую часть полученного выражения.

Вспомним, что статистика случайных величин k_1^o и k_1^p описывается геометрическими функциями распределения вероятностей, $P_o(k_1^o) = p(1-p)^{k_1^o-1}$ и $P_p(k_1^p) = (1-p)p^{k_1^p}$ (см п. (4)). Значит

$$\begin{aligned}\langle k_1^o \rangle &= \frac{1}{p}, \\ \langle I(k_1^p > N - 1) \rangle &= \text{Prob}[k_1^p > N - 1] = \sum_{k_1^p=N}^{\infty} (1-p)p^{k_1^p} - 1 = p^{N-1}, \\ \langle k_1^p I(k_1^p \leq N - 1) \rangle &= \sum_{k_1^p=1}^{N-1} k_1^p (1-p)p^{k_1^p} = \frac{1 + (N-1)p^{N-Np^{N-1}}}{1-p}.\end{aligned}$$

Подстановка этих формул в Ур. (46) дает финальный ответ к задаче

$$\langle k \rangle = \frac{1}{1-p} \left(\frac{1}{p^N} - 1 \right).$$

В частности, для симметричной монетки ($p = 1/2$) имеем $\langle k \rangle = 2^{N+1} - 2$. То есть, чтобы встретить десять орлов подряд придется подбросить монетку (в среднем) $2^{11} - 2 = 2046$ раз.

Белан-л5.6.

Крупные метеориты падают на Землю в пуассоновские моменты времени с рейтом λ . Обозначим через t_n время ожидания падения n -го по счету метеорита, отсчитываемое от произвольно выбранного момента начала наблюдений. Какова функция плотности вероятности случайной величины t_n ?

Белан-л5.7.

В области полностью ионизированного водорода начинается происходить рекомбинация электронов с ионами, в результате чего газ постепенно деионизируется. Вероятность того, что электрон рекомбинируется за время dt равна $\rho(t)\beta dt$, где $\rho(t)$ это концентрация ионов в текущий момент времени t , а β - некоторый постоянный коэффициент, характеризующий акт рекомбинации. Известно, что концентрация ионов в начальный момент времени равна ρ_0 , а система в целом электронейтральна. Найдите плотность вероятностей для случайной величины T - времени, прошедшего до момента рекомбинации отдельно взятого электрона. Каково медианное значение времени рекомбинации T ?

6.2.6 Problems About Polynomial Schema

Ч-2.80

При прохождении одного порога байдарка не получает повреждений с вероятностью p_1 , полностью ломается с вероятностью p_2 , получает серьезное повреждение с вероятностью p_3 ($p_1 + p_2 + p_3 = 1$). Два серьезных повреждения приводят к полной поломке. Найти вероятность того, что при прохождении n порогов байдарка не будет полностью сломана.

Ч-2.81

Сообщения, передаваемые по каналу связи, состояются из трех знаков А, В, С. Из-за помех каждый знак принимается правильно с вероятностью 0,6 и принимается ошибочно за любой из двух других знаков с вероятностью 0,2. Для увеличения вероятности правильного приема каждый знак передается 5 раз. За переданный знак принимается знак, который чаще всего встречается в принятой пятерке знаков. Если наиболее частых знаков два, то из них выбирается равновероятно один. Найти вероятность правильного приема знака при указанном способе передачи.

Ч-2.82

Пусть $\xi_{i,\uparrow}$ - число появлений исхода i в n первых испытаниях полиномиальной схемы (см. введение К гл. 2). Найти $P\{\xi_{n,1} = m_1\}$.

Ч-2.83

В схеме, описанной в задаче п2.82, найти условную вероятность

$$P\{\xi_{n,2} = m_2, \dots, \xi_{n,N} = m_n | \xi_{n,1} = m_1\}.$$

Ч-2.84

В N ячейках, разбитых на две группы по N_1 и N_2 ячеек соответственно ($N_1 + N_2 = N$), независимо одну от другой размещают n частиц; пусть p_{ij} ($i = 1, 2; j = 1, 2, \dots, N_1$) - вероятность попадания частицы в j -ю ячейку i -й группы, а $\prod_{ij}^{(n)}$ - число частиц, попавших в j -ю ячейку i -й группы после размещения n частиц. Найти:

- $P\{\eta_{ij}^{(n)} = k_{ij}, i = 1, 2, j = 1, \dots, N\},$
- $P\{\eta_{21}^{(n)} + \dots + \eta_{2N_2}^{(n)} = n_2\},$
- $P\{\eta_{1j} = k_{1j}, j = 1^2, \dots, N_1 | \eta_{21} + \dots + \eta_{2N_2} = n_2\}.$

Ч-2.87

Исходы $\theta_1, \theta_2, \dots$ последовательности испытаний с $N = 3$ возможными исходами 1, 2, 3 и вероятностями исходов p_1, p_2, p_3 объединяются в тройки $(\theta_{3k+1}, \theta_{3k+2}, \theta_{3k+3})$. Из первой тройки $(\theta_{3v+1}, \theta_{3v+2}, \theta_{3v+3})$, в которой все исходы различны, выбирается θ_{3v+1} . Найти $P\{\theta_{3v+1} = i\}$.

Ч-2.88

Испытания в полиномиальной схеме с исходами 1, 2, 3, имеющими вероятности p_1, p_2, p_3 соответственно, заканчиваются, когда впервые не появится исход 3. Найти вероятность того, что испытания закончатся исходом 1.

Ч-2.89

Игрок A подбрасывает 3 игральные кости, а игрок B - 2 кости. Эти испытания они проводят вместе и последовательно до первого выпадения "6" хотя бы на одной кости. Найти вероятности событий:

- $A = \{\text{впервые 6 появилось у игрока } A, \text{ а не } B\};$
- $B = \{\text{впервые 46 и появилось у игрока } B, \text{ а не } A\};$
- $C = \{\text{впервые "6" появилось одновременно у } A \text{ и } B\}.$

6.3 Problems of Probability Theory in Life and Everyday Life

(мб с собеседований тоже сюда буду собирать задачи.)

6.3.1 Problems About Typical Simple Life Tasks (!!!!!)

(тут вот то, что встречается в жизни и буду решать!! заработает теория, если пройду основы нормально!!!)

(другой сборник открыл, лень писать, какой)

Время ремонта и обслуживания автомобиля после одной поездки случайно и имеет экспоненциальный закон распределения. Было замечено, что в текущем сезоне на ремонт и обслуживание автомобиля после одной поездки тратилось в среднем 5 минут. Найти вероятность того, что при очередной поездке это время не превысит 30 минут.

Пусть X - время ремонта и обслуживания автомобиля после одной поездки. По условию задачи $X \sim E(\lambda)$ и $M[X] = 5$. По формуле связи $M[X]$ и λ получаем, что $M[X] = \frac{1}{\lambda} = 5$. Отсюда находим $\lambda = 1/5$. Требуется найти $P\{X \leq 30\}$. Плотность вероятности СВ X имеет вид

$$f(x) = \begin{cases} \frac{1}{5}e^{-x/5}, & x > 0 \\ 0, & x \leq 0 \end{cases}$$

Следовательно, получаем

$$P\{X \leq 30\} = \int_0^{30} \frac{1}{5}e^{-x/5} dx = -e^{-x/5} \Big|_0^{30} = 1 - e^{-6}$$

Ответ. $P\{X \leq 30\} = 1 - e^{-6}$.

Рост взрослого мужчины удовлетворительно описывается нормальным законом распределения. По статистике средний рост составляет 175 см, а среднеквадратическое отклонение равно 7 см. Найти вероятность того, что рост наугад взятого мужчины будет отличаться от среднего роста не больше чем на 7 см.

Решение. Обозначим рост наугад взятого взрослого мужчины через X . По условию задачи $X \sim N(175; 49)$. Требуется найти $P\{|X - M[X]| \leq 7\}$. Так как интервал $(M[X] - 7, M[X] + 7)$ симметричен относительно $M[X]$, то

$$P\{|X - M[X]| \leq 7\} = 2\Phi_0\left(\frac{7}{\sqrt{D[X]}}\right) = 2\Phi_0(1) \approx 2 \cdot 0,3413 = 0,6826.$$

О т в е т. Искомая вероятность приблизительно равна 0,6826.

МФТИ-Т.18 Мячи из ящика (??!!)

В ящике находится 10 теннисных мячей, из которых 6 новые. Для первой игры наугад берут два мяча, которые после игры возвращают в ящик. Для второй игры также наугад берут 2 мяча. Найти вероятность того, что оба мяча, взятые для второй игры, новые.

Посмотрим на первое вытягивание двух мячей. Пусть A_1 - событие изъятия 2х старых мячей; A_2 - событие изъятия одного старого мяча; A_3 событие изъятия только новых мячей. С помощью комбинаторики легко посчитать эту вероятность:

$$P(A_1) = \frac{C_4^2}{C_{10}^2} = \frac{6 \cdot 5}{10 \cdot 9} = \frac{2}{15} \quad (6.4)$$

$$P(A_2) = \frac{C_6^1 C_4^1}{C_{10}^2} = \frac{8}{15} \quad (6.5)$$

$$P(A_3) = \frac{C_6^2}{C_{10}^2} = \frac{1}{3} \quad (6.6)$$

Теперь, зная каждый случай первого доставания, второе доставание - назовем его A - при котором достались два новых мяча можно записать с помощью формулы полной вероятности.

$$P(A) = P(A_1)P(A|A_1) + P(A_2)P(A|A_2) + P(A_3)P(A|A_3)$$

$$P(A) = \frac{2}{15} \frac{C_6^2}{C_{10}^2} + \frac{8}{15} \frac{C_5^2}{C_{10}^2} + \frac{1}{3} \frac{C_4^2}{C_{10}^2} = \frac{2}{45} + \frac{16}{135} + \frac{2}{45} = \frac{12 + 16}{135} = \frac{28}{135}$$

Ответ $\frac{28}{135}$.

МФТИ-Т.19 Вероятность получения букв по каналу

(мб отдельный раздел про такие задачи сделаю.)

По каналу можно передавать одну из трех последовательностей букв: $AAAA, BBBB, CCCC$, причем априорные вероятности равны 0,3, 0,4 и 0,3 соответственно. Известно, что действие шумов на приемное устройство уменьшает вероятность правильного приема каждой из переданных букв до 0,6, а вероятность приема переданной буквы за две другие увеличивается до 0,2 и 0,2. Предполагается, что буквы искажаются независимо друг от друга. Найти вероятность того, что была передана последовательность $AAAA$, если на приемном устройстве получено $ACAB$.

Нам не известно, как считать условную вероятность прошлого при условии будущего, однако есть формула Байеса, которая позволяет перейти к вероятностям будущего при условии прошлого, которые посчитать мы сможем. По формуле Байеса:

$$\begin{aligned} P(AAAA|ACAB) &= \\ &= \frac{P(AAAA)P(ACAB|AAAA)}{P(AAAA)P(ACAB|AAAA) + P(BBBB)P(ACAB|BBBB) + P(CCCC)P(ACAB|CCCC)} = \\ &= \frac{0,3 \cdot (0,6 \cdot 0,2 \cdot 0,6 \cdot 0,2)}{0,3(0,6 \cdot 0,2 \cdot 0,6 \cdot 0,2) + 0,4((0,2)^3 \cdot 0,6) + 0,3((0,2)^3 \cdot 0,6)} = \frac{9}{16} \quad (6.7) \end{aligned}$$

Ответ: $\frac{9}{16}$.

МФТИ-Т.20 автоматы и жетоны

Имеется три телефонных автомата, которые принимают специальные жетоны. Один из них никогда не работает, второй работает всегда, а третий работает с вероятностью $1/2$. Некто имеет три жетона пытается выяснить, какой из автоматов исправный (работает всегда). Он делает попытку на одном из автоматов, которая оказывается неудачной. Затем переходит к другому автомату, на котором две подряд попытки оказываются удачными. Какова вероятность, что этот автомат исправный?

У нас ситуация, когда в первом подходе мог быть или полуробочий или нерабочий автомат. Вероятность неудачи при полуробочем автомате. Так как вероятность работы полуробочего в два раза меньше, чем у рабочего (чья вероятность работы равна 1), то если мы будем подходить к автоматам, то вероятность, что неудача будет на полуробочем должна быть равна половине вероятности на рабочем, а в сумме вероятности равны 1. Из этих соображений следует, что вероятность, что первый автомат был нерабочим равна $\frac{2}{3}$, а что полуробочим - $\frac{1}{3}$.

Дальше, пусть был первый автомат не рабочий. Тогда после ухода из него остались только полуробочий и рабочий автоматы, вероятность успеха на рабочем (по аналогичным рассуждениям) равна $\frac{2}{3}$, а что полуробочим - $\frac{1}{3}$.

А если первый автомат был полуробочий, то остался только рабочий и нерабочий, вероятности успехов на первом равна 1, на втором - 0.

В итоге по формуле полной вероятности получаем (тут сперва вариант: нерабочий, рабочий, рабочий, потом: полуробочий, рабочий, рабочий):

$$P = \frac{2}{3} \frac{2}{3} \frac{2}{3} + \frac{1}{3} \cdot 1 \cdot 1 = \frac{17}{27}$$

Ответ $\frac{17}{27}$

МФТИ-Т.21 Торпеды по кораблю (!?)

(потом еще раз разберу)

Подводная лодка атакует корабль, выпуская по нему последовательно и независимо одну от другой n торпед. Каждая торпеда попадает в корабль с вероятностью p . При попадании торпеды с вероятностью $\frac{1}{m}$ затопляется один из m отсеков корабля. Определить вероятность гибели корабля, если для этого необходимо затопление не менее двух отсеков.

Попробуем просто посчитать число нужных нам событий, и сложить вероятности каждого.

Пусть попало в цель k торпед, нам нужно, чтобы попало $k \in (2, n)$, тогда $(n - k)$ штук промахнулось. Попавшие торпеды могут быть в разной последовательности, последовательностей всего C_n^k . Среди k попавших может быть разное число r пробитых отсеков $r \in (2, k)$, всего возможных последовательностей C_k^r .

Таким образом, нам остается только суммировать наши исходы:

$$P = \sum_{k=2}^n p^k (1-p)^{n-k} \sum_{r=2}^k C_k^r \left(\frac{1}{m}\right)^r \left(1 - \frac{1}{m}\right)^{k-r}$$

Вероятность исхода в первом испытании с вероятностью 60% (?????!!)

Пусть вероятность нужного исхода равна X , сколько нужно сделать испытаний, чтобы получить этот исход с вероятностью 60% (?????)

(прямо очень актуальная жизненная задача, освою теорвер - прорешаю, пока не так владею им)

6.3.2 Service Sector Tasks with Many Customers (!!!)

(это первый раздел, ибо самый прикольный. потом обязательно классификацию додумаю, чтобы прямо с ходу нужные задачи выдавать!!! пока в одном месте всё.)

 n конвертов случайным n адресатам (????)

(?? тут пока слабая теория)

В n конвертов разложено по одному письму n адресатам. На каждом конверте наудачу написан один из n адресов. Найти вероятность того, что хотя бы одно письмо пойдет по назначению.

Пусть A_i - элементарное событие, состоящее в том, что i -е письмо дошло своему адресату. Тогда нужно найти.

$$P\left(\bigcup_{i=1}^n A_i\right) = \sum_{m=1}^n (-1)^{m-1} \sum_{1 \leq i_1 < \dots < i_m \leq n} P(A_{i_1} A_{i_2} \dots A_{i_m})$$

Посмотрим на каждое слагаемое внимательнее. Каждое слагаемое - вероятность того, что у нас будет m попаданий конвертов и писем. Если зафиксировать m писем, то оставшиеся письма можно раскладывать $(n - m)!$ способами. Всего можно n писем по n конвертам разложить $n!$ способами:

$$P(A_{i_1} A_{i_2} \dots A_{i_m}) = \frac{(n - m)!}{n!}$$

при этом при любых перестановках адресатов вероятность такая же, а всего перестановок при фиксированном m будет C_n^m . В итоге формула примет вид:

$$P = \sum_{m=1}^n (-1)^{m-1} C_n^m \frac{(n - m)!}{n!} = \sum_{m=1}^n (-1)^{m-1} \frac{1}{m!}$$

$$\text{Ответ: } P = \sum_{m=1}^n (-1)^{m-1} \frac{1}{m!}$$

Хватит ли сдачи на кассе (!!!)

(!!! вот тут нужно как следует потренироваться и сделать решение быстрым!!)

У билетной кассы стоит очередь в 100 человек. Половина людей в очереди имеет 100-рублевые купюры, а вторая половина - 50-рублевые купюры. Изначально в кассе нет денег и стоимость билета - 50 рублей. Какова вероятность, что никому не придется ждать сдачи?

Для каждого человека, имеющего 100р чтобы нашлась сдача необходимо, чтобы был перед ним человек, пришедший с 50р.

Таким образом, задача эквивалентна задаче о скобочной последовательности, где есть открывающая и закрывающая скобка. и каждая открытая должна иметь после нее закрывающую.

Известно, что число скобочных последовательностей из n открывающих и n закрывающих скобок равно числу Каталана $K_n = \frac{1}{n+1} C_n^{2n}$.

Всего же возможных последовательностей - число сочетаний C_{2n}^n , Поэтому ответ: $P = \frac{1}{n+1}$, для нашего случая, $n = 50$, имеем $P = \frac{1}{51}$

$$\text{Ответ: } P = \frac{1}{51}$$

Вероятность успешной сдачи экзамена при половине подготовки

Студент выучил 50 из 100 экзаменационных вопросов. В экзаменационном билете будет 9 различных вопросов. Вычислить вероятность получения студентом билета, содержащего не более 1 вопроса из числа невыученных.

Поликлиника закреплена за двумя районами города. При этом население первого района в 3 раза больше населения второго района. Известно, что прививку от гриппа в среднем сделали каждые 39 из 100 жителей первого района и каждые 42 из ста жителей второго района. Вычислить вероятность того, что случайным образом выбранный человек проживает в первом районе при условии, что он не сделал прививку от гриппа.

6.3.3 Distance Tasks

Вероятность, что случайно расположенные на дороге три человека находятся не близко друг с другом

Вдоль дороги, длиной в 1 км расположены случайным образом три человека. Найти вероятность того, что никакие два человека не находятся друг от друга на расстоянии, меньшем $1/4$ км.

Обозначим координаты каждого человека x, y, z . Пусть они расположены в порядке $x \rightarrow y \rightarrow z$, тогда благоприятный исход A состоит в следующем:

$$A = \left\{ \begin{array}{l} y - x \geq \frac{1}{4} \\ z - y \geq \frac{1}{4} \\ \frac{1}{4} \leq y \leq \frac{3}{4} \end{array} \right.$$

(?? тут ниже еще кст нормировка какая-то идет, типа чтобы обезразмерить координаты ведь!!!)

$$P(A) = \int_{\frac{1}{4}}^{\frac{3}{4}} dy \int_0^{y-\frac{1}{4}} dx \int_{y+\frac{1}{4}}^1 dz = \int_{\frac{1}{4}}^{\frac{3}{4}} dy \left(y - \frac{1}{4} \right) \left(1 - y - \frac{1}{4} \right) = \frac{1}{48}$$

Порядок трех точек может быть получен $3! = 6$ способами, поэтому ответ есть полученная вероятность, домноженная на 6.

Ответ: $P = \frac{1}{8}$.

6.3.4 Meeting Tasks

Случайный выход пешехода и автобуса (???)

(?? тут путаница много!!!)

(?? что там с геометрией СТО???? мб там какие-то особенности есть, о которых Суслов говорил?? потом подумаю)

Расстояние от пункта А до пункта В автобус проходит за 2 минуты, а пешеход - за 15 минут. Интервал движения автобусов 25 минут. Пешеход в случайный момент времени из пункта А и отправляется в В пешком. Найти вероятность того, что в пути пешехода догонит очередной автобус.

Будем решать графически. Нарисуем 2D пространство-время, как на картинке.




Figure 6.2: для задачи T12. Здесь короткая линия под наклоном - линия движения автобуса. длинные наклонные линии - линии движения человека. заштрихованная область - область всех возможных линий, где мировая линия человека не пересечет мировую линию автобуса. (я рисовал рисунок, решая обратную задачу, поэтому там написано $\frac{12}{25}$)

Несложные соображения показывают, что незаштрихованная область, которая соответствует встрече, может быть такой такой.

Таким образом, вероятность не встречи равна отношению площади этой области к интервалу, в котором мы смотрим движение автобуса. То есть вероятность равна: $\frac{15-2}{25} = \frac{13}{25}$

Ответ: $\frac{13}{25}$

Вероятность прихода на каждую лекцию хотя бы одного студента

В течение учебного семестра проводится 14 лекций по теории вероятностей для потока, на котором учатся 60 студентов. Каждый студент независимо от остальных приходит на лекцию с вероятностью 0.25. Вычислить вероятность того, что на каждую лекцию придёт хотя бы один студент.

Считаем типично как обратная вероятность, потому что так проще. На 1 лекцию 1 студент не придет с вероятностью 0,75. Никто не придет на одну лекцию с вероятностью $0,75^{60}$. На каждую лекцию никто не придет - то есть будет 14 лекций, каждой вероятностью $0,75^{60}$ - с вероятностью $0,75^{60 \cdot 14}$. Таким образом, хотя бы один придет на каждую с вероятностью $1 - 0,75^{60 \cdot 14}$.

6.4 Problems About Random Variables (!!!!)

6.4.1 Problems About the Inequalities of Chebyshev and Markov

(тут очень мощные задачи, соберу потом! ??? куда это лучше вставить, пока не знаю????)

МФТИ-Т.1 Число посетителей салона

Известно, что число посетителей некоторого салона в день является случайной величиной со средним значением 50.

(а) Оценить вероятность того, что число посетителей в конкретный день превысит 75

Решается неравенством Маркова:

$$\mathbf{P}\{|\xi| \geq x\} \leq \frac{\mathbf{M}|\xi|}{x}$$

то

$$P(\xi \geq 75) \leq \frac{50}{75} = \frac{2}{3}$$

(b) При условии, что дисперсия числа посетителей в день равна 25, оценить вероятность того, что в конкретный день их число будет между 40 и 60.

Решается неравенством Чебышева:

$$P(|\xi - 50| \geq 10) \leq \frac{25}{100} = \frac{1}{4}$$

поэтому

$$P(|\xi - 50| < 10) = 1 - P(|\xi - 50| \geq 10) \geq \frac{3}{4}$$

(?? а если бы было между 50 и 60-ю?????)

Самарова-нЧ,збч.1

По известному «правилу трех сигм» вероятность отклонения случайной величины от своего математического ожидания более, чем на три корня из дисперсии, мала. Найти

$$P(|\xi - E\xi| < 3\sqrt{D\xi})$$

если ξ имеет:

1) нормальное распределение с параметрами a, σ ;

2) распределение Пуассона с $E\xi = 0,09$.

Сравнить результат с оценкой, полученной по неравенству Чебышева.

Solution 1.

$$E\xi = a, \quad D\xi = \sigma^2$$

$$p_\xi(x) = \frac{1}{\sigma\sqrt{2\pi}} \cdot e^{-\frac{(x-a)^2}{2\sigma^2}}$$

$$P(|\xi - E\xi| < 3\sqrt{D\xi}) = P(|\xi - a| < 3\sigma) = P(-3\sigma < \xi - a < 3\sigma) =$$

$$= P(a - 3\sigma < \xi < a + 3\sigma) = \frac{1}{\sigma\sqrt{2\pi}} \int_{a-3\sigma}^{a+3\sigma} e^{-\frac{(x-a)^2}{2\sigma^2}} dx$$

$$y := \frac{x-a}{\sigma}, \quad dy = \frac{dx}{\sigma},$$

$$P(|\xi - E\xi| < 3\sqrt{D\xi}) = \frac{1}{\sqrt{2\pi}} \int_{-3}^3 e^{-\frac{y^2}{2}} dy = \frac{2}{\sqrt{2\pi}} \int_0^3 e^{-\frac{y^2}{2}} dy$$

Используем таблицу значений функции $\Phi_0(x) = \frac{1}{\sqrt{2\pi}} \int_0^x e^{-\frac{t^2}{2}} dt$,

$$P(|\xi - E\xi| < 3\sqrt{D\xi}) = 2\Phi_0(x) = 2 \cdot 0,4987 = 0,9974$$

А по неравенству Чебышева

$$P(|\xi - E\xi| \geq 3\sqrt{D\xi}) \leq \frac{D\xi}{9D\xi} = \frac{1}{9} \approx 0,1111 \quad \Rightarrow \quad P(|\xi - E\xi| < 3\sqrt{D\xi}) \geq 0,8889$$

Результат, полученный с помощью вероятностных таблиц, входит в область, полученную с помощью неравенства Чебышева.

Solution 2

Найдем параметр λ данного распределения. Числовые характеристики случайной величины, распределенной по закону Пуассона с параметром λ , равны

$$E\xi = \lambda, \quad D\xi = \lambda$$

Поэтому

$$\lambda = 0,09; \quad D\xi = E\xi = 0,09$$

Тогда

$$\begin{aligned} P(|\xi - E\xi| < 3\sqrt{D\xi}) &= P(|\xi - 0,09| < 0,9) = P(-0,9 + 0,09 < \xi < 0,9 + 0,09) = \\ &= P(-0,81 < \xi < 0,99) = P(\xi = 0) = e^{-0,09} = 0,9139 \end{aligned}$$

Здесь мы воспользовались формулой для вероятности, с которой случайная величина, распределенная по закону Пуассона, принимает заданное значение

$$P(\xi = k) = \frac{\lambda^k}{k!} e^{-\lambda} \quad (k = 0, 1, 2, 3, \dots)$$

Найдем оценку для требуемой вероятности по неравенству Чебышева:

$$P(|\xi - E\xi| \geq 3\sqrt{D\xi}) \leq \frac{D\xi}{9D\xi} = \frac{1}{9} \approx 0,1111 \quad \Rightarrow \quad P(|\xi - E\xi| < 3\sqrt{D\xi}) \geq 0,8889$$

Видим, что точное значение вероятности входит в интервал оценки, полученный с помощью неравенства Чебышева.

Стоит заметить, что в некоторых случаях улучшить неравенство Чебышева нельзя. Например, в задаче из задания 7 г) приведен пример случайной величины, для которой оценка, полученная по неравенству Чебышева, является точной.

Самарова-нЧ,збч.2

Для случайной величины ξ с распределением

$$P(\xi = 1) = P(\xi = -1) = \frac{1}{18}, \quad P(\xi = 0) = \frac{8}{9}$$

Найти

$$P(|\xi - E\xi| < 3\sqrt{D\xi})$$

и сравнить найденную вероятность с оценкой, полученной по неравенству Чебышева.

Solution

$$E\xi = 1 \cdot \frac{1}{18} + (-1) \cdot \frac{1}{18} + 0 \cdot \frac{8}{9} = 0$$

$$D\xi = E\xi^2 - (E\xi)^2 = E\xi^2 = 1 \cdot \frac{1}{18} + 1 \cdot \frac{1}{18} + 0 \cdot \frac{8}{9} = \frac{1}{9}$$

$$P(|\xi - E\xi| < 3\sqrt{D\xi}) = P(|\xi| < 1) = P(\xi = 0) = \frac{8}{9}$$

Заметим, что неравенство Чебышева в этом случае превращается в равенство

$$P(|\xi - E\xi| \geq 3\sqrt{D\xi}) \leq \frac{D\xi}{9D\xi} = \frac{1}{9} \quad \Rightarrow \quad P(|\xi - E\xi| < 3\sqrt{D\xi}) = 1 - \frac{1}{9} = \frac{8}{9}$$

МФТИ-Т.2 Вероятность нахождения числа выпадений герба в промежутке

С помощью неравенства Чебышева оценить вероятность того, что при 1000 бросаниях монеты число выпадений герба окажется в промежутке $[450, 550]$
Дисперсия в таком эксперименте

$$E(\xi^2) = E\left(\sum_i^n 1_{A_i}\right) + E\left(\sum_{i \neq j}^n\right) = 500 + (1000^2 - 100)/4 =$$

$$D(\xi) = E(\xi^2) - (E(\xi))^2 = 250$$

Неравенство Чебышева даст:

$$P(|\xi - 500| \geq 50) \leq \frac{250}{500}$$

$$P(|\xi - 500| < 50) = 1 - P(|\xi - 500| \geq 50) \geq 0.9$$

МФТИ-Т.3 Объем партии, при котором более всего качественных изделий (!?!?!?)

Вероятность того, что изделие качественное, равна 0,9. Каков должен быть объем партии изделий, чтобы с вероятностью $\geq 0,95$ можно было утверждать, что отклонение (по абсолютной величине) доли качественных изделий от 0,9 не превысит 0,01?

$$P(|\xi - 0,9| < 0,01) = 1 - P(|\xi - 0,9| \geq 0,01) \geq 0,95$$

$$P(|\xi - 0,9| \geq 0,01) \leq 0,05 = \frac{\sigma(\xi)^2}{10^{-4}}$$

но также $\sigma(\xi) = \sqrt{D(\xi)} = \frac{1}{n} \sqrt{n \cdot 0,9 \cdot 0,1} = \frac{0,3}{\sqrt{n}}$, поэтому:

$$0,05 = \frac{0,3^2}{n \cdot 10^{-4}} \rightarrow n = \frac{0,3^2}{0,05 \cdot 10^{-4}} = 18000$$

МФТИ-Т.4 Одностороннее неравенство Чебышева

Пусть случайная величина ξ имеет нулевое среднее и дисперсию σ^2 . Показать, что для $\varepsilon > 0$ выполняется неравенство

$$P(\xi \geq \varepsilon) \leq \frac{\sigma^2}{\sigma^2 + \varepsilon^2}$$

Очевидно, что для $b > 0$

$$P(\xi \geq \varepsilon) \leq P((\xi + b)^2 \geq (\varepsilon + b)^2) \leq \frac{\mathbb{E}((\xi + b)^2)}{(\varepsilon + b)^2} = \frac{\mathbb{E}(\xi^2 + 2\xi b + b^2)}{(\varepsilon + b)^2} = \frac{\mathbb{E}(\xi^2) + b^2}{(\varepsilon + b)^2}$$

Правая часть имеет минимум при $b = \frac{\sigma^2}{\varepsilon}$. Подставляя это значение, приходим к нужному неравенству.

МФТИ-Т5 Неравенство Йенсена

Пусть $\varphi : (a, b) \rightarrow \mathbb{R}$ - дважды непрерывно дифференцируемая функция и $\varphi''(x) \geq 0$ для всех $x \in (a, b)$ (т. е. φ Выпуклая функция, и допускается $a = -\infty$ и $b = \infty$). Допустим также, что ξ - случайная величина, которая принимает значения из (a, b) и $\mathbf{M}\xi = m$, $\mathbf{M}\varphi(\xi)$ конечны. Показать, что тогда

$$\mathbf{M}\varphi(\xi) \geq \varphi(\mathbf{M}\xi) = \varphi(m)$$

$$\varphi(x) \geq \varphi'(x_0)(x - x_0) + \varphi(x_0)$$

Пусть

$$\varphi(\xi) \geq \lambda(\mathbb{E}(\xi))(\xi - \mathbb{E}(\xi)) + \varphi(\mathbb{E}(\xi))$$

Берём матожидание от обеих частей.

$$\mathbb{E}(\varphi(\xi)) \geq \lambda(\mathbb{E}(\xi))(\mathbb{E}(\xi) - \mathbb{E}(\xi)) + \mathbb{E}(\varphi(\mathbb{E}(\xi))) = \varphi(\mathbb{E}(\xi))$$

МФТИ-Т6

Пусть ξ - случайная величина, имеющая нормальное распределение с параметрами (a, σ^2) . Показать, что для $\varepsilon > 0$ выполняется неравенство

$$\begin{aligned} \mathbf{P}(|\xi - a| \geq \varepsilon \sigma) &\leq \frac{1}{\varepsilon} \sqrt{\frac{2}{\pi}} e^{-\varepsilon^2/2} \\ P(|\xi| \geq \varepsilon \sigma) &= \left(\int_{\varepsilon \sigma}^{+\infty} + \int_{-\infty}^{-\varepsilon \sigma} \right) \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{x^2}{2\sigma^2}\right) d(x) = \int_{\varepsilon \sigma}^{+\infty} \frac{\sqrt{2}}{\sqrt{\pi}\sigma} \exp\left(-\frac{x^2}{2\sigma^2}\right) d(x) = \\ &= \sqrt{\frac{2}{\pi}} \int_{\varepsilon}^{+\infty} \exp\left(-\frac{x^2}{2}\right) d(x) \leq \sqrt{\frac{2}{\pi}} \frac{1}{\varepsilon} \exp\left(-\frac{\varepsilon^2}{2}\right) \end{aligned}$$

МФТИ-Т.7 еще неравенство (!?!?)

Пусть $\xi_1, \xi_2, \dots$ - последовательность одинаково распределенных случайных величин такая, что $\mathbf{M}\xi_k = a$, $\mathbf{D}\xi_k = \sigma^2$ и $\text{cov}(\xi_i, \xi_j) = (-1)^{i-j} \sigma^2$ $i \neq j$. Доказать, что для всякого $\varepsilon > 0$ выполняется предельное соотношение

$$\lim_{n \rightarrow \infty} \mathbf{P} \left(\left| \frac{1}{n} \sum_{k=1}^n \xi_k - a \right| \geq \varepsilon \right) = 0$$

МФТИ-Т.8 о чем-то схожим с Збч (???)

Пусть $\{\xi_n\}$ - последовательность неотрицательных случайных величин с конечными математическими ожиданиями и $\xi_n \xrightarrow{\text{п.н.}} \xi$, $\mathbf{M}\xi < \infty$, $\mathbf{M}\xi_n \rightarrow \mathbf{M}\xi$ при $n \rightarrow \infty$.

Покажите, что тогда $\mathbf{M}|\xi_n - \xi| \rightarrow 0$ при $n \rightarrow \infty$ т. е. $\xi_n \xrightarrow{L} \xi$
 $\forall n \in \mathbb{N}$:

$$|\mathbb{E}(\xi_n)| < +\infty, |\mathbb{E}(\xi)| < +\infty \rightarrow \mathbb{E}|\xi_n - \xi| < +\infty$$

$\forall \varepsilon > 0$:

$$\begin{aligned} \mathbb{E}|\xi_n - \xi| &= \mathbb{E}|(\xi_n - \xi) (I(\sup_{k \geq n} (|\xi_k - \xi| > \varepsilon)) + I(\sup_{k \geq n} (|\xi_k - \xi| \leq \varepsilon)))| = \\ &= \mathbb{E}|(\xi_n - \xi) I(\sup_{k \geq n} (|\xi_k - \xi| > \varepsilon))| + \mathbb{E}|(\xi_n - \xi) I(\sup_{k \geq n} (|\xi_k - \xi| \leq \varepsilon))| \end{aligned}$$

Первый индикатор в пределе принимает значение 1 только на множестве меры 0. Второй индикатор в пределе принимает значение 1 почти всюду. Но тогда под вторым математическим ожиданием стоит выражение $|\xi_k - \xi| \leq \varepsilon \rightarrow \mathbb{E}|\xi_k - \xi| \leq \mathbb{E}(\varepsilon) = \varepsilon$

6.4.2 Problems of Probability Distribution of Random Variables (!?!?)

(обязательно прорешаю скоро, просто пока не до этого. уже немного решал, в 1ю часть добавлю потом методы.)

Белан-ВШЭ-л1.1. (!?!?!?)

В диапазоне от 1 до N включительно равновероятным образом выбирается натуральное число. Оцените вероятность того, что это число окажется простым, если $\ln N \gg 1$.

$$\pi(N) \approx \int_2^N \frac{dk}{\ln k} \approx \frac{N}{\ln N}.$$

Solution

(?? тут какой-то гениальный метод, который я пока не понимаю.)

Белан-ВШЭ-л1.2. Первая цифра во времени жизни атома (!?)

Оцените вероятности различных значений, которые может принимать первая значащая цифра в десятичной записи случайного времени жизни радиоактивного атома неизвестного вещества.

Solution

Случайное времени жизни T радиоактивного атома вещества с рейтом распада λ имеет экспоненциальную функцию распределения $p(t) = \lambda \exp(-\lambda t)$. Обозначим первую значащую цифру случайной величины T через a_1 . От нас требуется вычислить вероятности $\Pr[a_1 = d|\lambda]$, где $d = 1, 2, \dots, 9$. Легко сообразить, что

$$\begin{aligned} \Pr[a_1 = d|\lambda] &= \sum_{n=-\infty}^{+\infty} \int_{10^n d}^{10^{n(d+1)}} p(t) dt = \lambda \sum_{n=-\infty}^{+\infty} \int_{10^n d}^{10^{n(d+1)}} \exp(-\lambda t) dt = \\ &= \sum_{n=-\infty}^{+\infty} [1 - \exp(-10^n \lambda)] \exp(-10^n d \lambda) \end{aligned}$$

Отметим, что полученное выражение строго инвариантно относительно замены $\lambda \rightarrow 10\lambda$. Этот факт с необходимостью следует из постановки задачи; ведь если используется именно десятичная запись числа, то умножение случайной величины на любую целочисленную степень десятки не изменит первую значащую цифру. Из Ур. (36) видно, что $\Pr[a_1 = d|\lambda] > \Pr[a_1 = d+1|\lambda]$, ведь $\exp(-10^n d \lambda) > \exp(-10^n (d+1) \lambda)$ для любого n . То есть, значение первой значащей цифры не равнораспределено, как можно было бы подумать. К сожалению, суммирование в формуле (36) невозможно произвести аналитическими методами. Однако, имеется возможность получить приближенный ответ, воспользовавшись подсказками численного анализа. А именно, построение графика выражения (36), рассматриваемого как функция от $\log_{10} \lambda$ показывает, что значение $\Pr[a_1 = d|\lambda]$ достаточно слабо зависит от этого параметра. Это означает, что искомое распределение первой значащей цифры в десятичной записи времени жизни атома практически не зависит от того, в секундах или в годах это время измеряется.

H.-A. Engel, C. Leuenberger / Statistics & Probability Letters 63 (2003) 361–365




Figure 6.3: На этом графике приняты следующие обозначения: $x = \log_{10} \lambda$, $h_d(x) = \Pr[a_1 = d|10^x]$

Так как функция $\Pr[a_1 = d|\lambda]$ периодична по переменной $x = \log_{10} \lambda$ с единичным периодом и слабо меняется, то для целей численной оценки она может быть заменена на амплитуду своей нулевой Фурье-гармоники. Иными словами, можно усреднить по периоду, записав

$$\begin{aligned} \Pr[a_1 = d|\lambda] &\approx \int_0^1 \Pr[a_1 = d|10^x] dx = \int_0^1 \sum_{n=-\infty}^{+\infty} [1 - \exp(-10^{n+x} \lambda)] \exp(-10^{n+x} d \lambda) dx + \text{37} \\ &= \int_{-\infty}^{+\infty} e^{-10^{x'}} d (1 - e^{-10^{x'}}) dx' = \frac{1}{\ln 10} \int_0^{+\infty} \frac{d\lambda}{\lambda} e^{-d\lambda} (1 - e^{-\lambda}) = \\ &= \frac{1}{\ln 10} \lim_{\varepsilon \rightarrow +0} \left[\int_0^{+\infty} d\lambda \lambda^{\varepsilon-1} e^{-d\lambda} - \int_0^{+\infty} d\lambda \lambda^{\varepsilon-1} e^{-(d+1)\lambda} \right] = \\ &= \frac{1}{\ln 10} \lim_{\varepsilon \rightarrow +0} \left[d^{-\varepsilon} \int_0^{+\infty} d\lambda' \lambda'^{\varepsilon-1} e^{-d\lambda'} - (d+1)^{-\varepsilon} \int_0^{+\infty} d\lambda' \lambda'^{\varepsilon-1} e^{-(d+1)\lambda'} \right] = \\ &= \frac{1}{\ln 10} \lim_{\varepsilon \rightarrow +0} [(d^{-\varepsilon} - (d+1)^{-\varepsilon}) \Gamma(\varepsilon)] = \frac{\ln \frac{d+1}{d}}{\ln 10} = \log_{10} \frac{d+1}{d} \end{aligned}$$

где

$$\Gamma(z) \equiv \int_0^{+\infty} t^{z-1} e^{-t} dt,$$

$$d^{-\varepsilon} - (d+1)^{-\varepsilon} = e^{-\varepsilon \ln d} - e^{-\varepsilon \ln(d+1)} \approx \varepsilon \ln \frac{d+1}{d},$$

$\Gamma(\varepsilon) \approx \int_0^{+\delta} t^{\varepsilon-1} dt = \frac{\delta^\varepsilon}{\varepsilon} \approx \frac{1}{\varepsilon}$ (Так как при $\varepsilon \rightarrow +0$ интеграл расходится в точке $t = 0$, то при малом ε можно сказать, что доминирующий вклад набирается в малой области $t \in (0, +\delta)$, где $\delta \ll 1$ и $\delta^\varepsilon \approx 1$)
Полученное распределение называется законом Бенфорда. Очевидно выполнение условия нормировки:
 $\sum_{d=1}^9 \Pr[a_1 = d] = \sum_{d=1}^9 \log_{10} \frac{d+1}{d} = 1.$

оставшиеся вопросы

как так мы численно считаем???? пропишу этот метод, что если период-я функция, то так вот делаем!!!!
это самое гениальное в задаче.

Белан-ВШЭ-л1.3. (!!!!!!)

Рассмотрим случайную величину, равномерно распределенную между 0 и 1. Какова вероятность того, что ровно две из пяти цифр после запятой в десятичной записи этой величины заняты цифрами меньше 6?
(решения нет, я пока хз.)

Solution

Белан-ВШЭ-л1.4. (!!!!!!)

Найдите $\langle \max(X_1, X_2) \rangle$, где X_1 и X_2 - две независимые реализации гауссовой случайной величины, с первым моментом μ и дисперсией σ .

Solution

(решения нет, я пока хз.)

Белан-ВШЭ-л1.5. (!!!!!!)

Найдите $\langle \min(|X|, 1) \rangle$, где X - случайная величина с функцией распределения $p(x) = \frac{1}{\pi} \frac{1}{1+x^2}$.

Solution

$$\begin{aligned} \langle \min(|X|, 1) \rangle &= \langle |X| \cdot I(|X| \leq 1) \rangle + \langle 1 \cdot I(|X| > 1) \rangle = \\ &= \int_{-1}^{+1} dx |x| p(x) + \int_{-\infty}^{-1} dx p(x) + \int_{+1}^{+\infty} dx p(x) = \frac{2}{\pi} \int_0^{+1} \frac{x dx}{1+x^2} + \frac{2}{\pi} \int_{+1}^{+\infty} \frac{dx}{1+x^2} = \\ &= \frac{1}{\pi} \ln(1+x^2) \Big|_0^1 + \frac{2}{\pi} \arctan x \Big|_1^{+\infty} = \frac{\ln 2}{\pi} + \frac{1}{2}. \end{aligned}$$

Белан-ВШЭ-л1.7. (!!!!!!)

Случайная величина X имеет монотонно возрастающую кумулятивную функцию распределения $F(x)$. Найдите функцию распределения случайной величины $Y = F(X)$.

Solution

Обозначим функцию плотности вероятности случайной величины X как $p(x)$. Так как кумулятивная функция распределения $F(x) = \int_0^x dx' p(x')$ является монотонно возрастающей, то функция распределения $\rho(y)$ случайной величины $Y = F(X)$ равна

$$\rho(y) = p(x(y)) \frac{dx(y)}{dy} = p(x) \left(\frac{dF(x)}{dx} \right)^{-1} |_{x=x(y)} = p(x(y)) \frac{1}{p(x(y))} = 1$$

для $0 \leq y \leq 1$

Белан-ВШЭ-л1.8. (?!?!?!?)

Выразите математическое ожидание и дисперсию случайной переменной $Z = aX + bY$, где X и Y - независимые случайные величины с известными первыми двумя статистическими моментами, a и b - некоторые вещественные константы.

Solution**Белан-ВШЭ-л1.9. (?!?!?!?)**

Неотрицательные случайные величины X_1 и X_2 независимы и одинаково распределены с плотностью вероятности $p(x) = a \exp(-ax)$. Найдите функцию плотности вероятности случайной величины $Y = X_1/X_2$.

Solution

Найдем кумулятивную функцию распределения случайной переменной $Y = X_1/X_2$

$$\begin{aligned} F(y) &= \Pr \left[\frac{X_1}{X_2} \leq Y \right] = \Pr [X_1 \leq Y X_2] = \int_0^{+\infty} dx_2 p(x_2) \int_0^{y x_2} dx_1 p(x_1) = \\ &= a^2 \int_0^{+\infty} dx_2 e^{-a x_2} \int_0^{y x_2} dx_1 e^{-a x_1} = a \int_0^{+\infty} dx_2 e^{-a x_2} (1 - e^{-a y x_2}) = \\ &= a \int_0^{+\infty} dx_2 (e^{-a x_2} - e^{-a(y+1)x_2}) = 1 - \frac{1}{y+1} \end{aligned}$$

Значит искомая функция плотности вероятности равна

$$\rho(y) = \frac{dF(y)}{dy} = \frac{1}{(y+1)^2}$$

ФР-2023. Распределение порядковых статистик.

Пусть $\xi_1, \dots, \xi_n$ - независимые случайные величины из одного и того же распределения. В предположении, что распределение каждой из ξ_k - абсолютно непрерывно с функцией распределения $F(x)$ и плотностью $p(x)$, найдите распределение k -той порядковой статистики $\xi_{(k)}$.

Solution**ФР-2023. Интеграл**

Используя результат предыдущей задачи (???), найти в зависимости от n интеграл

$$I(n) = \int_{[0,1]^n} f(x_1, \dots, x_n) d^n x,$$

где $f(x_1, \dots, x_n)$ равно второму по величине сверху из чисел $x_1, \dots, x_n$.

Solution**Показательность для случайной величины без старения (??)**

Докажите, что если X - непрерывная неотрицательная случайная величина, для которой отсутствует старение (т.е. для всех $t, s > 0$ $P(X > t + s | X > t) = P(X > s)$), то она показательно распределена.

(Олег решил, но я еще слишком тупой, чтобы понять решение, так что пропускаю, подрасту - вернусь к ней.)

Распределение $\frac{x}{X-Y}$ от показательных X и Y

Пусть X и Y - независимые случайные величины, показательно распределенные с параметром λ . Найдите распределение величины $\frac{X}{X-Y}$.

(переучу потом, тут разные новые для меня методы есть.)


Solution через интегрирование, грамотно определяя область

Как всегда, в силу независимости $P(XY) = P(X)P(Y)$ делим нужную вероятность на произведение двух:

$$P\left(X; Y = \left(1 - \frac{1}{k}\right)X\right) = P(X)P\left(Y = \left(1 - \frac{1}{k}\right)X\right) = \int_X \int_Y \lambda e^{-\lambda x} \cdot \lambda e^{-\lambda y} dx dy$$

Определим область интегрирования. Оно будет свое для каждого k , который мы повышаем от $-\infty$ до $+\infty$. Наше исходное множество - область $X > 0, Y > 0$, именно по некоторым подмножествам его будут наши интегралы. Прямая $Y = \left(1 - \frac{1}{k}\right)X$ разделяет его на подмножества, если её коэффициент более нуля: $\left(1 - \frac{1}{k}\right) \geq 0$. Это возможно только при $k < 0, k > 1$, именно эти случаи и будем рассматривать.

Пусть $k < 0$. Прямая $Y = \left(1 - \frac{1}{k}\right)X$ при $k < 0$ проходит выше прямой $Y = X$. Рассмотрим подмножество $X > Y$, исходное неравенство $\frac{X}{X-Y} \leq k$ дает область $Y \geq \left(1 - \frac{1}{k}\right)X$, которая не пересекается с подмножеством $X > Y$. Рассмотрим подмножество $X < Y$, исходное неравенство $\frac{X}{X-Y} \leq k$ дает область $Y \leq \left(1 - \frac{1}{k}\right)X$, таким образом, получаем область между прямыми $X = Y$ и $Y = \left(1 - \frac{1}{k}\right)X$:




Считаем:


```
Integrate[L^2 Exp[-L(y+x)], {x,0,Infinity},{y,x,(1-1/k)x}, Assumptions->{Im[L]
==0,L>0,k<0}]

Plot[1/(2 - 4 k), {k, -10, 0}, PlotRange -> {0, 1}]
```

Получаем:




Пусть $k > 1$. Прямая $Y = (1 - \frac{1}{k})X$ проходит ниже прямой $Y = X$. Рассмотрим подмножество $X > Y$, исходное неравенство $\frac{X}{X-Y} \leq k$ дает на нем область $Y \leq (1 - \frac{1}{k})X$, то есть это множество, заключенное между прямой $Y = (1 - \frac{1}{k})X$ и осью X . Рассмотрим подмножество $X < Y$, исходное неравенство $\frac{X}{X-Y} \leq k$ дает на нем область $Y \geq (1 - \frac{1}{k})X$, то есть это вся область $Y > X$. В итоге для $k > 1$ интегрируем по $Y > X$ и $0 < Y \leq (1 - \frac{1}{k})X$:



Считаем:

```
Integrate[L^2 Exp[-L(y+x)], {y,0,Infinity},{x,0,y}, Assumptions->{Im[L] == 0, L > 0}]
+
Integrate[L^2 Exp[-L(y+x)], {x,0,Infinity},{y,0,(1-1/k)x}, Assumptions->{Im[L]
==0,L>0,k>1}]
Plot[3/2 + k/(1 - 2 k), {k, 1, 10}, PlotRange -> {0, 1}]
```

Получаем:



Получаем, что область интегрирования красиво меняется при увеличении k от $-\infty$ до 0 из прямой $Y = X$ до треугольника между осью Y и прямой $Y = X$, а дальше при увеличении k от 1 до $+\infty$, область продолжает увеличиваться, только уже от оси X , и до заполнения всей четверть плоскости $X > 0, Y > 0$.

Ответ: функция распределения имеет вид: $F(k) = \frac{1}{2-4k}$ при $k < 0$, $F(k) = \frac{1}{2}$ при $k \in [0, 1]$, $F(k) = \frac{3}{2} + \frac{k}{1-2k}$ при $k > 1$.

В общем, здесь по типичным методам происходит переход к интегралу от двух функций, а далее самое сложное: четко преобразовать основное неравенство для определения множества интегрирования. Важно понимать, что для каждого k мы должны определить свою область интегрирования, чем мы и занимаемся, аккуратно решая неравенство. Важно тут быть очень внимательным и последовательным! Поэтому эта задача сложная. После определения области интегрирования, важно правильно подставить её границы в интеграл, для этого нужна подготовка по математическому анализу. Также важно решать через вольфрам, без этого может быть огромное количество ошибок.

(мб потом улучшу вольфрамский код, пока прога слабее нужного. как соединить графики в один?? как более компактно записать код?? потом улучшу скорее всего, пока суть итак прекрасно понятна.)
(??? напишу в теорию, что константой можно соединить эти отрезки)

Solution через переход к нужной функции распределения через последовательные шаги

(?!!!! вобщю эти шаблоны и идеи переходы по шагам в теорию. пока на этот метод забил, потому что итак метод есть рабочий.)

Сперва находим распределение Y/X , далее $1 - Y/X$, в итоге $1/(1-Y/X)$

Вот некачественная копия этих рассуждений:

$$\begin{aligned}
 &= \int_0^{+\infty} (1 - e^{-\lambda q t}) \cdot e^{-\lambda q} d(\lambda t) = \left(-e^{-\lambda q} + e^{-\lambda q(1+t)} \cdot \frac{1}{1+t} \right) \Big|_0^{+\infty} = \\
 &= 1 - \frac{1}{1+t} = \frac{t}{1+t}.
 \end{aligned}$$

$$\text{Uow, } F_{\frac{y}{x}}(t) = \begin{cases} 0, & t < 0, \\ \frac{t}{1+t}, & t \geq 0. \end{cases}$$

$$\text{Uak, } F_{\frac{y}{x}}(t) = \begin{cases} 0, & t < 0, \\ \frac{t}{1+t}, & t \geq 0. \end{cases} \text{ Cuyraidear becusuma } 1 - \frac{y}{x} \text{ pacmpegenera:}$$

$$\begin{aligned}
 F_{1-\frac{y}{x}}(t) &= P\left(1 - \frac{y}{x} < t\right) = P\left(\frac{y}{x} > 1 - t\right) = 1 - P\left(\frac{y}{x} \leq 1 - t\right) = \\
 &= 1 - F_{\frac{y}{x}}(1 - t) = \begin{cases} 1 - \frac{1-t}{2-t}, t \leq 1, \\ 1, t > 1 \end{cases} = \begin{cases} \frac{1}{2-t}, t \leq 1, \\ 1, t > 1. \end{cases}
 \end{aligned}$$

$$\frac{\frac{x}{x-y}}{\frac{x}{x-y}} = \frac{1}{1-\frac{y}{x}}. \quad \text{Ocranscs} \quad \text{nomers} \quad \text{coort.} \quad \text{vemgy} \quad \text{cuyracitsodx}$$

$$\begin{aligned}
 F_{\frac{1}{2}}(t) &= P\left(\frac{1}{q} < t\right) = \begin{cases} P\left(\eta \in \left(\frac{1}{t}; 0\right)\right), & t < 0, \\ P\left(\eta \in (-\infty; 0)\right), & t = 0, \\ P\left(\eta \in (-\infty; 0) \cup \left(\frac{1}{t}; +\infty\right)\right), & t > 0 \end{cases} = \\
 &= \begin{cases} F_q(0) - F_q\left(\frac{1}{t}\right), & t < 0, \\ F_q(0), & t = 0, \\ 1 + F_q(0) - F_q\left(\frac{1}{t}\right), & t > 0. \end{cases}
 \end{aligned}$$

$$\text{B ramese ayrae: } \eta = 1 - \frac{y}{x}. \text{ Uak,}$$

$$\begin{aligned}
 F_{\frac{x}{x-y}}(t) &= \begin{cases} \frac{1}{2} - \frac{1}{2-\frac{1}{t}}, t < 0, \\ \frac{1}{2}, t = 0, \\ 1 + \frac{1}{2} - 1, t \in (0; 1), \\ 1 + \frac{1}{2} - \frac{1}{2-\frac{1}{t}}, t \geq 1 \end{cases} = \begin{cases} \frac{1}{2(2t-1)}, t < 0, \\ \frac{1}{2}, t = 0, \\ \frac{1}{2}, t \in (0, 1), \\ \frac{4t-3}{4t-2}, t \geq 1 \end{cases} \\
 &= \begin{cases} -\frac{1}{4t-2}, t < 0, \\ \frac{1}{2} \in [0; 1], \\ 1 - \frac{1}{4t-2}, t > 1. \end{cases}
 \end{aligned}$$

Распределение $x/(x+y)$ от показательных и y (!!!!)

(!!! вобью в 1ю часть, мощные важные методы!!! очень важно отработать и вбить!!!)

Another way

Another way is to do it is to perform the following transformation. Let

$$\begin{aligned}
 U &= \frac{X}{X+Y} \\
 V &= X+Y
 \end{aligned}$$

Then

$$\begin{aligned}
 x(u, v) &= uv \\
 y(u, v) &= v - uv.
 \end{aligned}$$

The determinant of the Jacobian of the transformation is

$$\begin{vmatrix} v & u \\ -v & 1-u \end{vmatrix} = v$$

Then, the joint pdf of U and V is

$$f_{U,V}(u, v) = f_{X,Y}(uv, v - uv)|v| = ve^{-v}, \text{ for } v \geq 0, 0 \leq u \leq 1.$$

Finally, "integrate out" the v to obtain the marginal pdf of $U = X/(X+Y)$ as

$$f_U(u) = \int_0^{\infty} f_{U,V}(u, v) dv = \int_0^{\infty} ve^{-v} dv = 1, \quad \text{for } 0 \leq u \leq 1.$$

Hence, as Michael Hardy showed, $X/(X+Y)$ is uniformly distributed on $[0, 1]$.

1 method

Obviously $X/(X+Y)$ is between 0 and 1. Let (lower-case) w be a number between 0 and 1.

$$\frac{X}{X+Y} \leq w \text{ if and only if } (1-w)X \leq wY$$

We will find this probability. The conditional probability that $Y \geq \frac{(1-w)}{w}X$, given the value of X , is

$$e^{-(1-w)X/w}$$

The probability we seek is then the expected value of that:

$$\begin{aligned} Ee^{-(1-w)X/w} &= \int_0^\infty e^{-(1-w)x/w} (e^{-x} dx) \\ &= \int_0^\infty e^{-x/w} dx \\ &= w \end{aligned}$$

In other words, this random variable is uniformly distributed between 0 and 1.

Second method:

The equality $Y = \frac{1-w}{w}X$ is a straight line through the (X, Y) -plane, passing through $(0, 0)$ and having positive slope. We can let y go from 0 to ∞ and then for any fixed value of y , let x go from 0 to $\frac{wy}{1-w}$

$$\begin{aligned} &\int_0^\infty \left(\int_0^{wy/(1-w)} e^{-y} e^{-x} dx \right) dy \\ &= \int_0^\infty e^{-y} \left(1 - e^{-wy/(1-w)} \right) dy \\ &= \int_0^\infty e^{-y} - e^{-y/(1-w)} dy \\ &= 1 - (1-w) \\ &= w \end{aligned}$$

Распределение $2x/(1-x^2)$, при x по Коши

Пусть X - случайная величина, распределенная по Коши $(0, 1)$. Найдите распределение величины $2X/(1-X^2)$.

(Олег решил, но я еще слишком тупой, чтобы понять решение, так что пропускаю, подрасту - вернусь к ней.)

ВШЭ вероятность первой значащей цифры (!!!)

Задача 2 (7 баллов) Случайная величина X имеет функцию распределения

$$p(x) = \begin{cases} \frac{\alpha-1}{x^\alpha}, & \text{если } x \geq 1, \\ 0, & \text{если } x < 1. \end{cases}$$

где $\alpha > 1$. Найдите вероятности различных значений, которые может принимать первая цифра в десятичной записи величины X . На основе точного ответа отдельно исследуйте случаи $\alpha - 1 \ll 1$ и $\alpha \gg 1$.

$$P(d, a) = \sum_{n=0}^{+\infty} \int_{10^n d}^{10^{n+1} d} p(x) dx$$

(тут потом напишу, почему так, там выше задача похожая, тут очень важная её модификация)

$$P(d, \alpha) = \sum_{n=0}^{+\infty} (-x^{-\alpha+1})_{10^n d}^{107(\alpha+1)} = \sum_{n=0}^{\infty} 10^{n(1-\alpha)} (d^{-\alpha+1} - (d+1)^{-\alpha+1})$$

имеем геом прогрессию с $q = 10^{(1-\alpha)}$

$$p(d, \alpha) = \frac{d^{-\alpha+1} - (d+1)^{-\alpha+1}}{1 - 10^{1-\alpha}}$$

Далее один предел

$$\begin{aligned} 2-1) \text{ ирен } \alpha \gg 1 \\ d^{-\alpha} = d^{-\alpha} \gg (d+1)^{\alpha} \\ 1 > 10^{-\alpha} \\ 1 > 10^{-\alpha} \\ p(d, \alpha) \approx \frac{d^{-\alpha}}{1} = (d^{-\alpha}) \end{aligned}$$

второй предел

2.2)

$$\begin{aligned} P(d, \beta) &= \frac{d^{-\beta} - (d+1)^{-\beta}}{2 - 10^{-\beta}} = \frac{d^{-\beta} \left(1 - \left(\frac{d+1}{d}\right)^{-\beta}\right)}{1 - 10^{-\beta}} = \\ &= \frac{d^{-\beta}(-1) \left(\frac{d+1}{d}\right)^{-\beta} \log\left(\frac{d+1}{d}\right)}{(-1)10^{-\beta} \log 10} = \\ &= \left(\frac{d+1}{10}\right)^{-\beta} \frac{\log\left(\frac{d+1}{d}\right)}{\log 10} \approx \frac{\log \frac{d+1}{d}}{\log 10} = \\ &= \log_{10} \frac{d+1}{d}. \end{aligned}$$

Оставшиеся вопросы (!!!!!)

(почему вольфрам криво считает интеграл???? вообще просто хз, когда-то пороюсь, пока не знаю.)

```
In[92]:= Integrate[(a - 1)/x^a, {x, 10^n d, 10^n (d + 1)},
Assumptions -> {a > 1, d > 0, 10^n >= 0}]

Out[92]
= 10^n (100^n d (1 + d))^-a (- (10^n d)^a (1 + d) +
d (10^n (1 + d))^a)
```

и вот дальше вообще херня пойдет, если эту формулу использовать.

(!!!! впишу, что именно так мы и ищем предельные случаи!!! именно такая гениальная идея находить первую значащую цифру!!! часа 6 ушло на эту задачу, важно прописать этим методы потом!!! просто по сути нужно аккуратно проделать очевидные математические преобразования!!!!!! поэтому эта задача и жесткая!!!)

(!! пропишу методы аппроксимации, что так вот по лопиталю нужно аккуратно брать производные!!! что они через логарифм берутся знать важно. потом когда-то дооформилю.)

ВШЭ вероятность первой значащей цифры в периоде радиоактивного распада

Solution

Основные ходы, которые в 1ю часть добавлю потом!!!

(пока просто слабоватая структура.)

Магистратура Мфти 2019.7.

Случайные величины X и Y независимы. Случайная величина X имеет плотность распределения

$$\rho_X(t) = \begin{cases} e^t, & t < 0 \\ 0, & t \geq 0 \end{cases}$$

Случайная величина Y имеет плотность распределения

$$\rho_Y(t) = \begin{cases} 0, & t < 0, \\ e^{-t}, & t \geq 0. \end{cases}$$

- а) Найти функцию распределения случайной величины $Z = X + Y$.
 б) Вычислить математическое ожидание случайной величины Z .
 в) Вычислить дисперсию случайной величины Z .

Ответ: а) $P(Z < t) = \begin{cases} 1 - \frac{\exp(-t)}{2}, & t > 0, \\ \frac{\exp(t)}{2}, & t \leq 0, \end{cases}$ $\rho_Z(t) = \frac{\exp(-|t|)}{2}$, б) $MZ = 0$, в) $DZ = 2$.

Для любого $t \in \mathbb{R}$ имеем

$$\begin{aligned} P(Z < t) &= \iint_{\substack{\xi + \eta < t \\ \xi < 0, \eta > 0}} e^{\xi - \eta} d\xi d\eta = \int_{-\infty}^{\min\{t, 0\}} e^{\xi} d\xi \int_0^{t-\xi} e^{-\eta} d\eta = \int_{-\infty}^{\min\{t, 0\}} e^{\xi} (1 - e^{\xi-t}) d\xi = \\ &= \exp(\min\{t, 0\}) - \frac{\exp(2 \min\{t, 0\} - t)}{2} = \begin{cases} 1 - \frac{\exp(-t)}{2}, & t > 0, \\ \frac{\exp(t)}{2}, & t \leq 0, \end{cases} \end{aligned}$$

Следовательно, плотность распределения случайной величины Z равна $\rho_Z(t) = \frac{\exp(-|t|)}{2}$.

$$MZ = \int_{-\infty}^{+\infty} t \rho_Z(t) dt = \int_{-\infty}^{+\infty} \frac{t \exp(-|t|)}{2} dt = 0,$$

$$DZ = MZ^2 - (MZ)^2 = \int_{-\infty}^{+\infty} t^2 \rho_Z(t) dt = \int_{-\infty}^{+\infty} \frac{t^2 \exp(-|t|)}{2} dt = \int_0^{+\infty} t^2 \exp(-t) dt = 2.$$

Ч-3.пр.1 Распределение координат точки в квадрате

Пусть точка $A = (u, v)$ равномерно распределена в квадрате $\Omega = \{(u, v) : 0 \leq u \leq 1, 0 \leq v \leq 1\}$. Положим $\xi_1 = \xi_1(u, v) = u$,

$$\xi_2 = \xi_2(u, v) = \begin{cases} 1 & \text{при } u \geq v, \\ -1 & \text{при } u < v. \end{cases}$$

Найти:

- а) функцию распределения и плотность распределения вероятностей случайной величины ξ_1 ;
 б) функцию распределения и вероятности значений дискретной величины ξ_2 .

Событие $\{\xi_1 \leq x\} = \{(u, v) : u \leq x, (u, v) \in \Omega\}$ при $x > 1$ совпадает с Ω , а при $x < 0$ — с $\emptyset$. При $0 < x < 1$ множество точек, определяющих событие $\{\xi_1 \leq x\}$, образует прямоугольник со сторонами 1 и x .

Таким образом,

$$F_{\xi_1}(x) = P\{\xi_1 \leq x\} = \begin{cases} 0 & \text{при } x \leq 0, \\ x & \text{при } 0 < x < 1, \\ 1 & \text{при } x \geq 1. \end{cases}$$

Производная $F'_{\xi_1}(x)$ в ее точках непрерывности совпадает с плотностью распределения:

$$p_{\xi_1}(x) = \begin{cases} 0 & \text{при } x \notin [0, 1], \\ 1 & \text{при } x \in [0, 1]. \end{cases}$$

Функция распределения ξ_2 находится аналогично. Если $x \geq 1$, то $\{\xi_2 \leq x\} = \Omega$; $\{\xi_1 \leq x\} = \emptyset$ при $x < -1$, а при $-1 \leq x < 1$

$$\{\xi_2 \leq x\} = \{(u, v) : u < v, (u, v) \in \Omega\}.$$

Значит,

$$F_{\xi_2}(x) = P(\xi_2 \leq x) = \begin{cases} 0 & \text{при } x < -1, \\ 1/2 & \text{при } -1 \leq x < 1, \\ 1 & \text{при } x \geq 1. \end{cases}$$

Случайная величина ξ_2 дискретна; ее распределение сосредоточено в двух точках: -1 и 1 . Действительно,

$$P\{\xi_2 = 1\} = P\{(u, v) : u > v, (u, v) \in \Omega\} = 1/2,$$

$$P\{\xi_2 = -1\} = P\{(u, v) : u < v, (u, v) \in \Omega\} = 1/2.$$

Ч-3.пр.2 Функция от точки в квадрате

Найти распределение величины $\eta = \xi_1^2$, где случайная величина ξ_1 та же, что в примере п3.1.

Функция распределения η при $0 \leq x \leq 1$ определяется следующей формулой: $F_\eta(x) = P\{\eta \leq x\} = P\{\xi_1^2 \leq x\} = P\{-\sqrt{x} \leq \xi_1 \leq \sqrt{x}\} =$

$$\Rightarrow F_{\xi_1}(\sqrt{x}) - F_{\xi_1}(-\sqrt{x}) = \sqrt{x}.$$

Очевидно, что $F_\eta(x) = 0$ при $x \leq 0$ и $F_\eta(x) = 1$ при $x > 1$. При $0 < x < 1$

$$p_\eta(x) = F'_{\xi_1}\left(\sqrt{x}\right) = \frac{1}{2\sqrt{x}};$$

$p_\eta(x) = 0$ при $x < 0$ и при $x > 1$. Таким образом, величина η абсолютно непрерывна и ее плотность распределения определяется формулой

$$p_\eta(x) = 1/(2\sqrt{x}) \quad \text{при} \quad x \in (0, 1), \\ p_\eta(x) = 0 \quad \text{при} \quad x \notin (0, 1).$$

Ч-3.пр.7

Для случайной величины η , определенной в примере 3.2, найти $M\eta$ и $D\eta$.

Если плотность распределения известна, то для вычисления $M\eta$ удобно воспользоваться формулой

$$D\eta = \int_{-\infty}^{\infty} (x - M\eta)^2 p_\eta(x) dx$$

которая является частным случаем (3.9) с $g(x) \Rightarrow (x - M\eta)^2$. Подставляя в эти формулы плотность распределения $p_\eta(x)$, вычисленную в примере 3.2, получим

$$M_\eta = \int_0^1 x \frac{1}{2\sqrt{x}} dx = \frac{1}{3}, \quad D_\eta = \int_0^1 \left(x - \frac{1}{3}\right)^2 \frac{1}{2\sqrt{x}} dx = \frac{4}{45}.$$

Можно было также воспользоваться формулами (3.9) при $g(x) = x^2$, $g(x) = (x - M\eta)^2$ с заменой $p_4(x)$ на плотность $p_{\xi_1^2}(x)$, определенную в примере 3.1:

$$M_\eta = M_{\xi_1^2}^2 - \int_{-\infty}^{\infty} x^2 p_{\xi_1}(x) dx = \int_0^1 x^2 dx = \frac{1}{3} \\ = \int_0^1 \left(x - \frac{1}{3}\right)^2 dx = \frac{4}{45}.$$

Ч-3.пр.3

Случайная величина ξ имеет плотность распределения $p_\xi(x) = e^{-x}$ ($x > 0$), $p_\xi(x) = 0$ ($x \leq 0$). Найти распределение случайной величины $\eta = \eta(\xi) = [\xi]^2$, где $[\xi]$ обозначает целую часть ξ .

Решение. Так как $\{\eta = k^2\} = \{k \leq \xi < k+1\}$, то вероятность $P\{\eta = k^2\}$ определяется равенствами $P\{\eta = k^2\} = P\{k \leq \xi < k+1\} =$

$$= \int_k^{k+1} p_\xi(x) dx = \int_k^{k+1} e^{-x} dx = e^{-k} (1 - e^{-1}), k = 0, 1, 2, \dots$$

Ч-3.пр.4

Случайная величина ξ_1 принимает значения 0 и 1, а случайная величина ξ_2 - значения -1, 0 и 1. Вероятности $P\{\xi_1 = i, \xi_2 = j\}$ задаются таблицей:

$i \backslash j$	-1	0	1
0	1/16	1/4	1/16
1	1/16	1/4	5/16

Найти распределение случайной величины $\eta = \xi_1 \xi_2$

Произведение $\xi_1 \xi_2$ равно нулю, если равен нулю хотя бы один из сомножителей:

$$\{\eta = 0\} = \{\xi_1 = 0, \xi_2 = -1\} \cup \{\xi_1 = \xi_2 = 0\} \cup \{\xi_1 = 0, \xi_2 = 1\} \cup \{\xi_1 = 1, \xi_2 = 0\},$$

поэтому суммируем нужные ячейки:

$$P\{\eta = 0\} = \frac{1}{16} + \frac{1}{4} + \frac{1}{16} + \frac{1}{4} = \frac{5}{8}.$$

Оставшиеся два значения (1, -1) и (1, 1) пары величин (ξ_1, ξ_2) приводят к двум значениям η : -1 и 1. Следовательно,

$$P\{\eta = -1\} = 1/16, \quad P\{\eta = 1\} = 5/16.$$

Таким образом, величина η дискретна; ее распределение сосредоточено на значениях -1, 0, 1, и вероятности этих значений равны соответственно 1/16, 5/8, 5/16.

Ч-3.пр.5

Случайные величины ξ_1 и ξ_2 независимы; их плотности распределения определяются формулами

$$p_{\xi_1}(x) = \begin{cases} 1 & \text{при } x \in [0, 1], \\ 0 & \text{при } x \notin [0, 1], \end{cases} \quad p_{\xi_2}(x) = \begin{cases} \frac{1}{2} & \text{при } x \in [0, 2], \\ 0 & \text{при } x \notin [0, 2]. \end{cases}$$

Найти функцию распределения и плотность распределения величины $\eta = \xi_1 \xi_2$.


Так как величины ξ_1 и ξ_2 независимы, то двумерная плотность распределения вектора (ξ_1, ξ_2) есть $p_1 \xi_2(u, v) = p_{\xi_1}(u) p_{\xi_2}(v)$. Найдем сначала функцию распределения η :

$$F_n(x) = P\{\eta \leq x\} = P\{\xi_1 \xi_2 \leq x\} = \iint_{D_x} p_{\xi_1}(u) p_{\xi_2}(v) du dv,$$

где $D_x = \{(u, v) : uv \leq x\}$. Отсюда

$$F_n(x) = \frac{1}{2} \iint_{D_x^*} du dv$$

где $D_x^* = D_x \cap \{(u, v) : 0 \leq u \leq 1, 0 \leq v \leq 2\}$.



Эта область в свою очередь $D_x^* \emptyset$ при $x < 0$ и $D_x^* = \{(u, v) : 0 \leq u \leq 1, 0 \leq v \leq 2\}$ при $x \geq 2$; следовательно,

$$P\{\eta \leq x\} = 0 \quad (x \leq 0), \quad P\{\eta \leq x\} = 1 \quad (x \geq 2).$$

При $0 < x < 2$

$$\iint_{D_x^*} dudv = 2 \cdot \frac{x}{2} + \int_{x/2}^1 \frac{x}{u} du = x - x \ln \frac{x}{2}.$$

Таким образом,

$$F_n(x) = \begin{cases} 0 & \text{при } x \leq 0. \\ \frac{x}{2} \left(1 - \ln \frac{x}{2}\right) & \text{при } 0 < x < 2 \\ 1 & \text{при } x \geq 2. \end{cases}$$

56 Отсюда находим плотность распределения

$$p_n(x) = \begin{cases} 0 & \text{при } x \notin (0, 2], \\ -\frac{1}{2} \ln \frac{x}{2} & \text{при } x \in (0, 2]. \end{cases}$$

Ч-3.пр.6.

Найти плотность распределения величины $\eta = \xi_1 + \xi_2$, где ξ_1, ξ_2 независимы и каждая из них равномерно распределена на отрезке $[0, 1]$.

По формуле (3.4)

$$p_\eta(x) = \int_{-\infty}^{\infty} p_{\xi_1}(y) p_{\xi_2}(x-y) dy = \int_0^1 p_{\xi_2}(x-y) dy.$$

Подынтегральная функция $p_{\xi_2}(x-y)$ положительна и равна 1, если $0 \leq x-y \leq 1$. Отрезок интегрирования $[0, 1]$ отрезок $[x-1, x]$ значений y , при которых $p_{\xi_2}(x-y) > 0$, не пересекаются, если $x < 0$ или $x > 2$. В этих слу-

$$p_n(x) = \int_0^x dy = x.$$

Если $1 \leq x \leq 2$, то $[0, 1] \cap [x-1, x] = [x-1, 1]$ и

$$p_n(x) = \int_{x-1}^1 dy = 2 - x.$$

Таким образом,

$$p_\eta(x) = \begin{cases} 0, & \text{если } x < 0 \text{ или } x > 2, \\ x, & \text{если } 0 \leq x \leq 1 \\ 2 - x, & \text{если } 1 \leq x \leq 2. \end{cases}$$

Ч-3.пр.8 (!?!?!?!?) про Бернулли

Пусть η_n - число переходов от успеха к неудаче или обратно в n испытаниях схемы Бернулли, в которой вероятность успеха в отдельном испытании равна p . Найти M_{η_n} и D_{η_n} .

Решение. Представим η_n в виде суммы индикаторов

$$\eta_n = \xi_1 + \xi_2 + \dots + \xi_{n-1},$$

где $\xi_i = 1$, если исходами i -го и $(i+1)$ -го испытаний были соответственно "успех" и "неудача" или "неудача" в «успех». В противном случае $\xi_i = 0$. Тогда

$$M\xi_1 = P\{\xi_i = 1\} = p(1-p) + (1-p)p = 2p(1-p)$$

и, следовательно,

$$M\eta_n = M\xi_1 + \dots + M\xi_{n-1} = 2(n-1)p(1-p).$$

Найдем теперь η_n^2 . Так как

$$\eta_n^2 = \sum_{i=1}^{n-1} \xi_i^2 + \sum_{\substack{i,j=1 \\ i \neq j}}^{n-1} \xi_i \xi_j = \sum_{i=1}^{n-1} \xi_i + 2 \sum_{i=1}^{n-2} \xi_i \xi_{i+1} + \sum_{\substack{i,j=1 \\ |i-j| \geq 2}}^{n-1} \xi_i \xi_j$$

и

$$M\xi_i\xi_{i+1} = p(1-p)p + (1-p)p(1-p) = p(1-p)_3$$

$$M_{\xi_i\xi_j} = M\xi_i M\xi_j = 4p^2(1-p)^2 \quad (|i-j| \geq 2)_s$$

То

$$M\eta_n^2 = 2p(1-p)(n-1) + 2(n-2)p(1-p) + 4p^2(1-p)^2(n^2 - 5n + 6).$$

Отсюда, воспользуя обозначение $q = 1 - p$, получаем

$$D_{\eta_{1n}} = M\eta_n^2 - (M\eta_n)^2 = 4pq(1-3pq)n - 2pq(3-10pq).$$

В задачах п3.1-п3.15 рассматриваются одномерные рас пределения; в п3.16-п3.38 - законы совместного распреде тения нескольких случайных величин; в **3.49–3.53** - случайные величины, связанные с последовательностям пслытаний.

Ч-3.1

Распределение дискретной случайной величинн & опрделается формулами

$$P\{\xi = i\} = \frac{1}{5}, \quad t = -2_s - 1_i 0_s 2.$$

Найти распределения величин $\eta_1 = -\xi, \eta_2 = |\xi|$.**Ч-3.2**

Распределение случайной пеличины ξ определя ется формулами $P\{\xi = k\} = C/k(k+1), k = 1, 2, \dots$
Найти: а) постоянную C ; б) $P\{\xi \leq 3\}$; в) $P\{n_1 \leq \xi \leq n_2\}$.

Ч-3.3

. Распределение дискретной случайной величины ξ определяется формулами: $P\{\xi = k\} = C/k(k+1)(k+2), k = 1, 2, \dots$
Найти: а) постоянную C ; б) $P\{\xi \geq 3\}$;
в) $P\{n_1 \leq \xi \leq n_2\}$.

Ч-3.4Плотность распределения случайной величины ξ задана формулами:

$$p_2(x) = C/x^4 \quad \text{при } x \geq 1;$$

$$p_2(x) = 0 \quad \text{при } x < 1.$$

Найти: а) постоянную C ; б) плотность распределения $\eta = 1/\xi$; в) $P\{0, 1 < \eta < 0, 3\}$.**Ч-3.5**

Случайная величина ξ имеет показательное распределение с параметром $\alpha : P\{\xi \leq x\} = 1 - e^{-\alpha x} (x \geq 0)$.
Найти плотности распределения случайных величин:

- а) $\eta_1 = \sqrt{\xi}$
- б) $\eta_2 = \xi^2$;
- в) $\eta_3 = \frac{1}{\alpha} \ln \xi$.

Ч-3.6

Случайная величина & распределена так же, как в задаче п3.5. Найти плотности распределения величин:

- а) $\eta_1 = \{\xi\}(\{\xi\} - \text{дробная доля } \xi)$; б) $\eta_2 = 1 - e^{-\alpha \xi}$.

Ч-3.7

Случайная величина ξ равномерно распределена на отрезке $[0, 1]$. Найти плотности распределения величин: а) $\eta_1 = 2\xi + 1$; б) $\eta_2 = -\ln(1 - \xi)$.

Ч-3.8

Случайная точка B имеет равномерное распределение на окружности $x^2 + (y - a)^2 = r^2$ с центром в точке $A = (0, a)$, а случайная точка $C = (\xi, 0)$ является пересечением оси абсцисс с прямой, проходящей через A и B . Найти функцию распределения и плотность распределения случайной величины ξ . (Распределение ξ называется распределением Коши.)

Ч-3.9

Случайная величина ξ имеет распределение Коши с плотностью $p_\xi(x) = \frac{1}{\pi} \cdot \frac{1}{1+x^2}$. Найти плотность распределения величин $\eta = \xi^2 / (1 + \xi^2)$, $\xi = 1 / (1 + \xi^2)$.

Ч-3.10

Случайная величина ξ имеет такое же распределение, как в задаче п3.9. Найти плотность распределения случайных величин $\eta = 2\xi / (1 - \xi^2)$, $\zeta = 1/\xi$.

Ч-3.11

Пусть $h(x)$ - плотность распределения случайной величины ξ ; функция $g(x)$ дифференцируема и на интервале $(0, 1)$ монотонно возрастает от $-\infty$ до $+\infty$. Найти: а) такую функцию (x) , что случайная величина $\eta = \varphi(\xi)$ имеет своей плотностью распределения $f(x) = h(g(x))g'(x)$; б) распределение случайной величины $\zeta = g(\eta)$.

Ч-3.12

Случайная величина ξ имеет непрерывную функцию распределения $F(x) = P\{\xi \leq x\}$. Показать, что случайная величина $\eta = F(\xi)$ имеет равномерное распределение на отрезке $[0, 1]$.

Ч-3.13

Пусть η имеет равномерное распределение на $[0, 1]$, а

$$F_{-1}(y) = \sup\{x : F(x) \leq y\}, \quad 0 \leq y \leq 1,$$

- Функция, обратная к функции распределения $F(x)$ (не обязательно непрерывной!). Доказать, что случайная величина $\xi = F_{-1}(\eta)$ имеет функцию распределения $F(x)$.

Ч-3.14

Построить пример такого абсолютно непрерывного распределения случайной величины ξ с плотностью $p_\xi(x)$ и такой непрерывной функции $g(x)$, что распределение случайной величины $\eta = g(\xi)$ не вырождено и дискретно.

Ч-3.15

Функция распределения $F(x)$ непрерывна в каждой точке. Доказать, что она равномерно непрерывна на всей прямой $-\infty < x < \infty$.

Ч-3.17

Случайные величины ξ и η независимы. Найти $P\{\xi = \eta\}$, если: а) ξ и η имеют одно и то же дискретное распределение $P\{\xi = x_k\} = P\{\eta = x_k\} = p_k$, $k = 0, 1, \dots$; б) функция распределения непрерывна.

Ч-3.18

Совместное распределение ξ, η является равномерным в единичном круге $x^2 + y^2 \leq 1$. Найти $P\{|\xi| \leq 3/4, |\eta| \leq 3/4\}$.

Ч-3.19

Плотность совместного распределения величин ξ, η определяется равенствами: $p_{\xi, \eta}(u, v) = 1$ при $(u, v) \in G$, $p_{\xi, \eta}(u, v) = 0$ при $(u, v) \notin G$, где $G = \{(u, v) : 0 \leq u \leq 2, 0 \leq v < 1 - \frac{1}{2}u\}$. Найти плотность распределения $p_\xi(x)$ случайной величины ξ .

Ч-3.20

Плотность совместного распределения $p_{\xi_1, \xi_2}(u, v)$ величин ξ_1, ξ_2 определяется равенствами $p_{\xi_1, \xi_2}(u, v) = -C(u+v)$ при $0 \leq u \leq 1, 0 \leq v \leq 1$ и $p_{\xi_1, \xi_2}(u, v) = 0$, в остальных случаях.

Найти: а) постоянную C ;

б) одномерные плотности распределения ξ_1 и ξ_2 ;

в) плотность распределения $\eta = \max(\xi_1, \xi_2)$.

Ч-3.21

Плотность совместного распределения величин ξ_1, ξ_2 определяется равенствами: $p_{\xi_1, \xi_2}(u, v) = \frac{2}{\pi(u^2+v^2)^3}$ при $u^2 + v^2 \geq 1$ и $p_{\xi_1, \xi_2}(u, v) = 0$ в остальных случаях. Найти плотность распределения $\eta = \sqrt{\xi_1^2 + \xi_2^2}$.

Ч-3.23

Неотрицательные случайные величины ξ_1, ξ_2 независимы и имеют одну и ту же плотность распределения $p(x), x \geq 0$. Найти плотность $q(u, v)$ совместного распределения случайных величин $\eta_1 = \xi_1 - \xi_2, \eta_2 = \sqrt{\xi_1^2 + \xi_2^2}$.

Ч-3.24

Случайные величины ξ_1, ξ_2 независимы и имеют одну и ту же плотность распределения $p(x)$. Найти совместную плотность распределения $q(r, \varphi)$ полярных координат (r, φ) точки (ξ_1, ξ_2) .

Ч-3.25

Случайные величины ξ_1 и ξ_2 независимы и имеют одно и то же показательное распределение: $P\{\xi_i \leq x\} = 1 - e^{-x}, x \geq 0; i = 1, 2$. Найти $P\{|\xi_1 - \xi_2| \leq 1\}$.

Ч-3.26

Случайные величины ξ и η независимы и имеют равномерное распределение на отрезке $[0, a]$. Найти плотности распределения случайных величин: а) $\xi + \eta$; б) $\xi - \eta$; в) $\xi\eta$; г) ξ/η . 69

Ч-3.27

Случайные величины ξ и η независимы и имеют показательное распределение с плотностью $e^{-x} (x \geq 0)$ каждая. Найти плотность распределения: а) $\xi + \eta$; б) $\xi - \eta$; в) $\xi\xi - \eta$; г) ξ/η .

Ч-3.28

Найти плотность распределения суммы $\xi + \eta$, если ξ и η независимы, ξ имеет равномерное распределение в отрезке $[0, 1]$, а η - равномерное распределение в отрезке $[0, 2]$.

Ч-3.29

Найти плотность распределения суммы независимых случайных величин. ξ и η , если ξ равномерно распределена в $[0, 1]$, а η имеет показательное распределение с плотностью $e^{-x} (x \geq 0)$.

Ч-3.30

Случайные величины ξ_1, ξ_2, ξ_3 независимы и имеют равномерное распределение в $[0, 1]$.

Найти плотности распределения сумм: а) $\xi_1 + \xi_2$

б) $\xi_1 + \xi_2 + \xi_3$. Найти $P\{0,5 \leq \xi_1 + \xi_2 + \xi_3 \leq 2,5\}$.

Ч-3.31

Точка (ξ_1, ξ_2) имеет равномерное распределение в квадрате $\{(x, y) : 0 \leq x \leq a, 0 \leq y \leq a\}$. Показать, что распределения случайных величин $|\xi_1 - \xi_2|$ и $\min\{\xi_1, \xi_2\}$ совпадают, т. е. что для любого t

$$P\{|\xi_1 - \xi_2| \leq t\} = P\{\min\{\xi_1, \xi_2\} \leq t\}.$$

Ч-3.32

Найти распределение суммы двух независимых слагаемых ξ_1 и ξ_2 , если слагаемые распределены:

- а) показательно с одним и тем же параметром a ;
 б) по закону Пуассона с параметрами λ_1 и λ_2 .

Ч-3.33

Случайные величины $\xi_1, \dots, \xi_n$ независимы и распределены показательно с одинаковым параметром α .

Найти плотность распределения величин:

- а) $\eta_1 = \max \{\xi_1, \dots, \xi_n\}$ б) $\eta_2 = \min \{\xi_1, \dots, \xi_n\}$.

Ч-3.34

Случайные величины ξ и ξ_2 независимы и имеют гамма-распределения: ξ_1 — с параметрами (λ, α) , ξ_2 — с параметрами (λ, β) . Пользуясь формулой

$$\int_0^1 t^{\alpha-1} (1-t)^{\beta-1} dt = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}.$$

найти плотность распределения случайной величины $\xi_1 + \xi_2$.

Ч-3.37

Найти плотность распределения случайной величины $\eta = \frac{\xi_1}{\xi_1 + \xi_2}$, если ξ_1, ξ_2 независимы и равномерно распределены в отрезке $[0, 1]$.

Ч-3.38

Случайные величины $\xi_1, \xi_2, \dots, \xi_n, n \geq 2$, независимы и имеют показательное распределение с плотностью $e^{-x} (x \geq 0)$. Обозначим $\eta = \frac{\xi_1}{\xi_1 + \dots + \xi_n}$. Найти плотность распределения η .

Ч-3.39

Величины ξ_1, ξ_2 независимы; $P\{\xi_1 = 0\} = P\{\xi_1 \Rightarrow 1\} = 1/2$, ξ_2 равномерно распределена на отрезке $[0, 1]$. Найти закон распределения величины $\xi_1 + \xi_2$.

Ч-3.40

Пусть случайные величины ξ и η независимы, ξ имеет функцию распределения $F(x)$, а η равномерно распределена в интервале $[a, b]$. Показать, что $\xi + \eta$ имеет плотность $\frac{F(x-a) - F(x-b)}{b-a}$.

Ч-3.41

Сумму двух независимых равномерно распределенных на $\{0, 1, \dots, 9\}$ случайных однозначных чисел ξ и η можно записать в виде $\xi + \eta = 10\xi_2 + \zeta_1$ ($0 \leq \zeta_i \leq 9$). Найти законы распределения. Зависимы ли ξ_1 и ξ_2 ?

Ч-3.42

Произведение двух независимых равномерно распределенных на $\{0, 1, \dots, 9\}$ однозначных чисел ξ и η можно записать в виде $\xi\eta = 10\xi_2 + \zeta_1$, где ζ_1, ζ_2 — целые числа, принимающие значения от 0 до 9. Зависимы ли ξ_1 и ξ_2 ?

Ч-3.43

Случайная точка $A = (\xi_1, \xi_2, \xi_3) \in R^3$ имеет равномерное распределение на сфере $x^2 + y^2 + z^2 = 1$. Найти распределение проекции (ξ_1, ξ_2) точки A на плоскость (x, y) и проекции ξ точки A на ось x .

Ч-3.44

Ввести на сфере в качестве координат широту и долготу, считая их изменяющимися в отрезках $[-\pi/2, \pi/2]$ и $[-\pi, \pi]$ соответственно. Найти плотность $p_\theta(x)$ распределения широты ξ случайной точки, имеющей равномерное распределение на сфере.

Ч-3.47

В переговорном пункте телефоны-автоматы расположены в t залах: в i -м зале n_i автоматов ($t = 1, 2, 3; n = n_1 + n_2 + n_3$). После перерыва посетители одновременно заняли все автоматы. Введем события $A_i = \{\text{посетитель, закончивший разговор первым, находился в } i\text{-м зале}\}, i = 1, 2, 3$. Найти вероятности событий $A_i, i = 1, 2, 3$, если времена разговора посетителей являются независимыми одинаково распределенными случайными величинами с непрерывной функцией распределения.

Ч-3.48

Случайная величина ξ с равномерным распределением на $[0, 1]$ записывается в виде бесконечной десятичной дроби: $\xi = \sum_{n=1}^{\infty} \xi_n 10^{-n}$, где ξ_n - целые, $0 \leq \xi_n \leq 9$. Доказать, что случайные величины $\xi_1, \xi_2, \dots$ независимы.

Ч-3.49

В схеме Бернулли с вероятностью успеха p обозначим через $v_k (k = 1, 2, \dots)$ номер испытания, при котором происходит k -й успех, и положим $\tau_1 = v_1, \tau_k = v_k - v_{k-1} (k = 2, 3, \dots)$. Найти совместное распределение величин τ_1, τ_2 . Являются ли эти величины независимыми?

Ч-3.50

В полиномиальной схеме с исходами $(1, 2, \dots, N)$ вероятность i -го исхода в каждом испытании равна $p_i, i = 1, 2, \dots, N$. Положим

$$\varepsilon_{s,i} = \begin{cases} 1, & \text{если в } s\text{-м испытании появился } i\text{-й исход} \\ 0 & \text{в противном случае.} \end{cases}$$

Являются ли независимыми случайные величины:

- а) $\varepsilon_{a,i}, \varepsilon_{s,i} \quad (s, i, j - \text{фиксированы});$
 б) $\varepsilon_{a,i}, \varepsilon_{t,j} \quad (s \neq t)$ в) $\varepsilon_{1,i}, \varepsilon_{2,i}, \dots, \varepsilon_{n,i};$
 г) $\sum_{k=1}^N \varepsilon_{1k} c_k, \sum_{k=1}^N \varepsilon_{2k} c_k, \dots, \sum_{k=1}^N \varepsilon_{nk} c_k$, если $c_1, c_2, \dots, c_N$ некоторые постоянные?
 (тут еще задачи были, но я не вставил их)

Ч-3.59

Доказать, что двумерная функция распределения $F_{6,n}(x, y)$ непрерывна в точке (x_0, y_0) , если соответствующие одномерные функции распределения $F_8(x), F_n(y)$ непрерывны в точках x_0 и y_0 соответственно.

Пусть $\xi_1, \xi_2, \dots, \xi_n$ - независимые одинаково распределенные случайные величины. При каждом $\omega \in \Omega$ расположим числа $\xi_n(\omega), k = 1, \dots, n$, в порядке возрастания и перенумеруем их заново: $\xi_{(1)} \leq \xi_{(2)} \leq \dots \leq \xi_{(n)}$. Полученная последовательность случайных величин называется вариационным рядом, а сами случайные величины $\xi_{(k)}$ - членами вариационного ряда. Таким образом, в частности, $\xi_{(1)} = \min_{1 \leq k \leq n} \xi_k, \xi_{(n)} = \max_{1 \leq k \leq n} \xi_k$. Задачи п3.60-п3.66 связаны с вариационным рядом.

Ч-3.60

Случайные величины $\xi_1, \xi_2, \dots, \xi_n (n \geq 2)$ независимы и одинаково распределены с функцией распределения $F(x)$. Найти: а) функцию распределения $\xi_{(1)}$; б) функцию распределения $\xi_{(n)}$; в) двумерную функцию распределения $\xi_{(1)}, \xi_{(n)}$.

Ч-3.61

По независимым одинаково распределенным случайным величинам $\xi_1, \dots, \xi_n$, имеющим функцию распределения $F(x)$ и плотность $p(x)$, построен вариационный ряд $\xi_{(1)} \leq \xi_{(2)} \leq \dots \leq \xi_{(n)}$. Найти: а) плотность распределения $\xi_{(m)}$; б) совместную плотность распределения $\xi_{(h)}$ и $\xi_{(m)} (k < m)$.

Ч-3.62

Случайные величины $\xi_1, \dots, \xi_k, \xi_{k+1}$ независимы и имеют одну и ту же непрерывную функцию распределения. Пусть $\xi_{(1)} \leq \xi_{(2)} \leq \dots \leq \xi_{(k)}$ - вариационный ряд величин $\xi_1, \dots, \xi_k$

$$\begin{aligned} P \{ \xi_{k+1} \in [\xi_{(l)}, \xi_{(l+1)}] \}, \quad l = 1, \dots, k-1, \\ P \{ \xi_{k+1} < \xi_{(1)} \}, \quad P \{ \xi_{k+1} > \xi_{(k)} \}. \end{aligned}$$

Ч-3.63

Случайные величины $\xi_1, \dots, \xi_n$ независимы и имеют одинаковое распределение с плотностью $p(x)$. Найти n -мерную плотность $p_n(x_1, \dots, x_n)$ распределения членов вариационного ряда $\xi_1, \xi_{(2)}, \dots, \xi_{(n)}$.

Ч-3.64

*. Случайные величины $\xi_1, \xi_2, \dots, \xi_n$ независимы и имеют показательное распределение с параметром α :

$$P\{\xi_1 \leq x\} = 1 - e^{-\alpha x}, \quad x \geq 0, \quad i = 1, \dots, n,$$

а $\xi_{(1)} \leq \xi_{(2)} \leq \dots \leq \xi_{(n)}$ - значения $\xi_1, \dots, \xi_n$, расположенные в порядке неубывания (вариационный ряд). Показать, что случайные величины

$$\Delta_1 = \xi_{(1)}, \quad \Delta_i = \xi_{(i)} - \xi_{(i-1)}, \quad i = 2, \dots, n,$$

независимы и что

$$P\{\Delta_1 \leq x\} = 1 - e^{-\alpha(n-i+1)x}, \quad i = 1, \dots, n.$$

Ч-3.65

Случайные величины $\xi_1, \xi_2, \dots, \xi_n$ независимы и имеют одно и то же показательное распределение с параметром λ . Доказать, что случайные величины

$$\max\{\xi_1, \xi_2, \dots, \xi_n\} \text{ и } \xi_1 + \frac{\xi_2}{2} + \frac{\xi_3}{3} + \dots + \frac{\xi_n}{n}$$

одинаково распределены.

Ч-3.66

Пусть $\xi_n = (\xi_{n,1}, \dots, \xi_{n,n})$, $n = 1, 2, \dots$, - последовательность независимых векторов, у которых координаты $\xi_{n,1}, \dots, \xi_{n,h}$ - независимые случайные величины, имеющие одну и ту же непрерывную функцию распределения $F(x)$. Положим

$$A_n = \left\{ \xi_{n,i} < \min_{1 \leq m \leq n} \xi_{m,t}, t = 1, \dots, k \right\}.$$

Доказать, что при $k \geq 2$

$$P\left\{ \bigcup_{n=2}^{\infty} A_n \right\} < \frac{1}{2^{k-3} \cdot g^0}$$

Ч-3.70

Случайная величина ξ_1 распределена по закону Пуассона с параметром λ_1 :

$$P\{\xi_1 = k\} = \frac{\lambda_1^k}{k!} e^{-\lambda_1}, \quad k = 0, 1, 2, \dots$$

а случайная величина ξ_2 распределена по закону Пуассона с параметром $\lambda_2 > \lambda_1$. Доказать, что при любом целом k

$$P\{\xi_1 \leq k\} > P\{\xi_2 \leq k\}.$$

Ч-3.71

Функции распределения $F_1(i)$ и $F_2(t)$ удовлетворяют условию

$$F_1(t) \leq F_2(t) \text{ для любого } t.$$

Показать, что можно так задать на одном вероятностном пространстве случайные величины ξ_1 и ξ_2 с функциями распределения $F_1(t)$ и $F_2(t)$ соответственно, что

$$P\{\xi_1 \geq \xi_2\} = 1.$$

равномерное распределение на отрезке $[0, 1]$, а случайная величина ζ удовлетворяет условиям

$$P\{\zeta = \xi\} = P\{\zeta = \eta\} = 1/2.$$

Найти максимально и минимально возможные значения $P\left\{ \left| \zeta - \frac{1}{2} \right| \leq t \right\}$ и указать, при каких совместных распределениях ξ, η, ζ эти значения достигаются.

6.4.3 Problems About Mathematical Expectations, Covariance, Correlation

(почему-то ватеры не делят на типы эти задачи, я мб потом разделю, посмотрим.)

В задачах п3.75-п3.112 используются прямые способы вычисления математических ожиданий; задачи п3.113-п3.131 иллюстрируют возможности метода индикаторов (см. введение к гл. 3). Задачи п3.132-п3.138 содержат полезные Формулы для математических ожиданий различных функций от случайных величин. Разными методами решаются задачи п3.139-п3.188.

Ч-3.75

Распределение дискретной случайной величины ξ определяется формулами $P\{\xi = i\} = 1/5, i = -2, -1, 0, 1, 2$. Найти математические ожидания величин $\eta_1 = -\xi, \eta_2 = |\xi|$. Распределение дискретной случайной величины ξ определяется формулами $P\{\xi = k\} = \frac{4}{k(k+1)(k+2)}, k = 1, 2, \dots$. Найти математическое ожидание случайной величины ξ .

Ч-3.77

Плотность распределения случайной величины & задана Формулами $p_z(x) = 0$ при $x < 1, p_s(x) = 3/x^4$ при $x \geq 1$. Найти математические ожидания и дисперсии случайных величин ξ и $\eta = 1/\xi$.

Ч-3.78

Найти $M\xi, D\xi, M\xi^{[k]} (k = 1, 2, \dots)$, если:

- $P\{\xi = m\} = \frac{\lambda^m}{m!} e^{-\lambda}, m = 0, 1, 2, \dots;$
- $P\{\xi = m\} = C_n^m p^m q^{n-m}, q = 1 - p, m = 0, 1, 2, \dots, n.$

Ч-3.79

Плотность совместного распределения $p_{\xi_1, \xi_2}(u, v)$ величин ξ_1, ξ_2 определяется равенствами $p_{q_1, z_2}(u, v) = u + v$ при $0 \leq u \leq 1, 0 \leq v \leq 1; p_{\xi_1, \xi_2}(u, v) = 0$ в остальных случаях. Найти $M\xi_1, M\xi_2, D\xi_1, D\xi_2, \text{cov}(\xi_1, \xi_2)$.

Ч-3.80

Совместное распределение ξ_1, ξ_2 определяется условиями $P\{\xi_1 \xi_2 = 0\} = 1, P\{\xi_i = 1\} = P\{\xi_i = -1\} = 1/4, i = 1, 2$. Найти $M\xi_1, M\xi_2, D\xi_1, D\xi_2, \text{cov}(\xi_1, \xi_2)$.

Ч-3.81

Случайная величина ξ равномерно распределена на отрезке $[0, 2\pi]; \eta_1 = \cos \xi, \eta_2 = \sin \xi$. Найти $M, M_2, \text{cov}(\eta_1, \eta_2)$. Являются ли η_1 и η_2 независимыми?

Ч-3.84

Плотность совместного распределения случайных ξ_1, ξ_2 равна $\frac{2}{\pi(u^2 + v^2)^3}$ при $u^2 + v^2 \geq 1$ и $p_{\xi_1, \xi_2}(u, v) = 0$ в остальных случаях. Найти $M\sqrt{\xi_1^2 + \xi_2^2}$.

Ч-3.85

Случайная величина ξ имеет показательное распределение: $P\{\xi > x\} = e^{-x}$ при $x \geq 0$. Найти $M\xi(1 - e^{-\alpha \xi})$.

Ч-3.86

Случайная точка A имеет равномерное распределение в круге радиуса R . Найти математическое ожидание и дисперсию расстояния ξ точки A от центра круга.

Ч-3.87

Случайная величина ζ имеет равномерное распределение на отрезке $[0, 1]$. Найти коэффициент корреляции случайных величин η_1, η_2 , если:

а) $\eta_1 = a\xi, \eta_2 = b\zeta \quad (a, b > 0)$,

б) $\eta_1 = a\zeta, \eta_2 = b\zeta \quad (a < 0 < b)$,

в) $\eta_1 = \zeta, \quad \eta_2 = \zeta^2$, г) $\eta_1 = \zeta - \frac{1}{2}, \quad \eta_2 = \left(\zeta - \frac{1}{2}\right)^2$,

д) $\eta_1 = \sin\left(\frac{\pi}{2}\zeta\right), \quad \eta_2 = \cos\left(\frac{\pi}{2}\zeta\right)$. Имеющей равномерное распределение в области $D \subset R^2$. Найти коэффициент корреляции $\rho(\xi, \eta)$, если:

а) D -часть единичного круга, лежащая в первом квадранте: $x^2 + y^2 \leq 1, x \geq 0, y \geq 0$.

б) D - треугольник: $x + y \leq 1, x \geq 0, y \geq 0$.

Ч-3.89

Пусть $\xi_1, \xi_2, \dots, \xi_n$ - независимые случайные величины, имеющие равномерное распределение на отрезке $[0, 1]$, а $\xi_{(1)} \leq \xi_{(2)} \leq \dots \leq \xi_{(n)}$ - построенный по $\xi_1, \dots, \xi_n$ вариационный ряд, т. е. значения $\xi_1, \dots, \xi_n$ - расположенные в порядке неубывания. Найти плотность $p_i(x)$ распределения $\xi_i, M\xi_{(i)}^h (k = 1, 2, \dots), D_{(i)}$.

Ч-3.91

Доказать, что если случайные величины то $M(\xi + \eta)^3 = M\xi^3 + M\eta^3$.

Ч-3.92

Случайные величины ξ и η некоррелированы. Доказать, что $M\xi\eta = M\xi M\eta$.

Ч-3.93

Случайные величины ξ, η и ξ попарно некоррелированы. Верно ли равенство $M\xi\eta\xi = M\xi M\eta M\xi$?

Ч-3.94

Случайные величины ξ, η, ξ имеют нулевые математические ожидания, дисперсии σ^2 и попарно некоррелированы. Чему равны минимальное и максимальное значения $M\xi\eta(???)$?

Ч-3.95

Найти ковариационную матрицу случайного вектора $\xi = (\xi_1, \xi_2, \xi_3)$, если:

а) ξ_1, ξ_2, ξ_3 независимы и имеют стандартное нормальное распределение;

б) вектор ξ имеет равномерное распределение в кубе $\{(x_1, x_2, x_3) : \max_{1 \leq i \leq 3} |x_i| \leq \sqrt{3}\}$;

в) вектор ξ с вероятностью $1/6$ принимает каждое из 6 значений $(0, 0, \pm\sqrt{3}), (0, \pm\sqrt{3}, 0), (\pm\sqrt{3}, 0, 0)$.

Ч-3.96

Какие из приведенных ниже матриц могут, а какие не могут быть ковариационными для случайного вектора $\xi = (\xi_1, \xi_2, \xi_3)$:

а) $\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$,

б) $\begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix} * \text{В)} \begin{pmatrix} 1 & 1 & 1 \\ -1 & 1 & 1 \\ -1 & -1 & 1 \end{pmatrix}$,

в) $\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}$,

г) $\begin{pmatrix} 1 & 2 & 3 \\ 2 & 3 & 4 \\ 3 & 4 & 5 \end{pmatrix}$ е) $\begin{pmatrix} 1 & 1 & -1 \\ 1 & 1 & 1 \\ -1 & 1 & 1 \end{pmatrix}$

Ч-3.98.

Случайный вектор (ξ_1, ξ_2, ξ_3) имеет ковариационную матрицу $\begin{pmatrix} 1 & \rho_{12} & \rho_{13} \\ \rho_{12} & 1 & \rho_{23} \\ \rho_{13} & \rho_{23} & 1 \end{pmatrix}$. Доказать, что

$$|\rho_{13} - \rho_{12}\rho_{23}| \leq \sqrt{(1 - \rho_{12}^2)(1 - \rho_{23}^2)}$$

Ч-3.99

а) Показать, что если распределение случайного вектора (ζ_1, ζ_2) совпадает с распределением вектора $(-\xi_1, -\xi_2)$ и $M\xi_1^2 + M\xi_2^2 < \infty$, то

$$\text{cov}(\zeta_1, \zeta_2) = 2M(\zeta_1 \max(0, \zeta_2)).$$

б) Пусть (ξ_1, η_1) и (ξ_2, η_2) - независимые одинаково распределенные векторы, $M(\xi_1^2 + \eta_1^2) < \infty$. Доказать, что $\text{cov}(\xi_1, \eta_1) = M((\eta_1 - \eta_2) \max(0, \xi_1 - \xi_2))$.

Ч-3.100

Случайные векторы $\xi = (\xi_1, \dots, \xi_n) \in R^n$ и $\xi^* = (\xi_1^*, \dots, \xi_n^*) \in R^n$ независимы, $\xi = (m_1, \dots, m_n) = m$, $M\xi^* = (m_1^*, \dots, m_n^*) = m$, матрицы ковариаций ξ и ξ^* равны $\sigma = \|\sigma_{ij}\|$ и $\sigma^* = \|\sigma_{ij}^*\|$ соответственно.

Найти: а) математическое ожидание и дисперсию случайной величины $\zeta = \xi_1 + \dots + \xi_n$;

б) математическое ожидание и дисперсию случайной величины $\eta = (\xi, a) = a_1\xi_1 + \dots + a_n\xi_n$, где $a = (a_1, \dots, a_n)$ - данный неслучайный вектор;

в) математические ожидания и ковариационные матрицы векторов $\xi + \xi^*, \xi - \xi^*, a\xi + b\xi^*$ (a, b - постоянные).

Ч-3.104

Случайные величины $\xi_1, \xi_2, \dots$ независимы и имеют равномерное распределение на отрезке $[0, 1]$. Найти математическое ожидание случайной величины

$$\eta_n = \sum_{i=1}^n |\xi_{i+1} - \xi_i|.$$

Ч-3.105

Найти дисперсию случайной величины η_n , введенной в задаче

Ч-3.104

Сравнить ее с $nD|\xi_2 - \xi_1| = D(|\xi_2 - \xi_1| + |\xi_4 - \xi_3| + \dots + |\xi_{2n} - \xi_{2n-1}|)$.

Ч-3.106

Случайные величины $\xi_1, \dots, \xi_n, \eta_1, \dots, \eta_n$ независимы. Положим

$$\zeta_n = \xi_1 + \dots + \xi_n, \zeta_n^* = \xi_1\eta_1 + \dots + \xi_n\eta_n.$$

Найти $M\zeta_n, M\zeta_n^*, D\zeta_n, D\zeta_n^*, \text{cov}(\zeta_n, \zeta_n^*)$, если

$$\begin{aligned} M\xi_k &= a, & D\xi_k &= \sigma^2, & P\{\eta_k = 1\} &= p, \\ P\{\eta_k = 0\} &= q = 1 - p & & & (k = 1, 2, \dots, n). \end{aligned}$$

Ч-3.107

В экспедиции, рассчитанной на n дней, ежедневно от запаса продуктов нужно отделять соответствующую часть: в 1-й день - $1/n$ -ю часть, во 2-й день $1/(n-1)$ -ю часть от остатка и т. д. В действительности нужная часть продуктов отделяется с ошибкой. Пусть $\eta_k (k = 1, 2, \dots, n-1)$ - часть от остатка продуктов, которая отделяется в k -й день. Предполагается, что величины η_k независимы, $M_k = a_k = \frac{1}{n-k+1}$. Найти математическое ожидание случайной величины ξ , равной части продуктов, оставшейся к последнему дню:

$$\zeta = (1 - \eta_1)(1 - \eta_2) \dots (1 - \eta_{n-1}).$$

Ч-3.112

Случайные величины $\xi_1, \xi_2, \dots$ независимы $M\xi_i = 0, D\xi_i = \sigma^2 < \infty, i = 1, 2, \dots$. Найти математическое ожидание и дисперсию случайных величин

$$S_n = \xi_1 + \xi_2 + \dots + \xi_n \text{ и } T_n = \xi_{v_1} + \xi_{v_2} + \dots + \xi_{v_n}$$

где $v_1, v_2, \dots$ независимые и все зависящие от $\xi_1, \xi_2, \dots$ случайные величины, имеющие равномерное распределение на множестве $\{1, 2, \dots, k\}$.

Ч-3.113

Обозначим через θ_r число циклов длины r в подстановке, случайно выбранной из множества всех $n!$ подстановок степени n . Найти: а) $M\theta_1, D\theta_1$; б) $M\theta_2$; в) $M\theta_r (r \geq 1)$.

Ч-3.114

В урне содержится M_1 шаров с номером 1, M_2 шаров с номером 2, ..., M_N шаров с номером N . По схеме случайного выбора без возвращения выбирается n шаров. Найти математическое ожидание числа появившихся номеров.

Ч-3.115

Из урны, содержащей M белых и $N - M$ черных шаров, по схеме выбора без возвращения извлекается гипергеометрическое распределение: $P\{\xi = m\} = \frac{C_M^m C_{N-M}^{n-m}}{C_N^n}$. Найти $M\xi$ и D .

Ч-3.116

В N ячейках случайно размещаются n частиц. Каждая частица независимо от остальных с вероятностью $1/N$ может попасть в любую фиксированную ячейку. Обозначим через $\mu_0(n, N)$ число пустых ячеек. Найти $M\mu_0(n, N), D\mu_0(n, N)$ и асимптотические формулы для них при $n, N \rightarrow \infty, n/N \rightarrow \alpha \in (0, \infty)$.

Ч-3.126

В схеме Бернулли задачи

Ч-3.125

найти формулы для математического ожидания и дисперсии числа μ_{111} появлений цепочки 111. (Цепочка 111 появляется при i -м испытании, если исходами $(i-2)$ -го, $(i-1)$ -го и i -го испытаний были единицы.) Сравнить их с формулами для математического ожидания и дисперсии числа μ_1 появлений исхода 1 в $n-2$ независимых испытаниях схемы Бернулли, когда вероятность исхода 1 равна p^3 .

Ч-3.127

Пусть проведено n испытаний по той же схеме Бернулли, что в задачах Ч-3.125 и Ч-3.126. Назовем серией единиц цепочку исходов последовательных испытаний вида 011...110 (при этом считается, что исходами дополнительных испытаний с номерами 0 и $n+1$ были нули). Пусть η_n - число серий единиц. Найти: а) $P\{\eta_n = 0\}$;

Ч-3.128

На бесконечный лист клетчатой бумаги (сторона клеточки равна 1) случайно бросается круг единичного радиуса. Считая, что центр круга равномерно распределен на том единичном квадрате, на который он попал, найти математическое ожидание числа точек с целочисленными координатами (x, y) , покрытых этим кругом.

Ч-3.129

*. Каждую целочисленную точку числовой оси независимо от остальных назовем белой с вероятностью q и черной с вероятностью $p = 1 - q$. Пусть S - множество всех таких целочисленных точек x , что расстояние от x до ближайшей черной точки (включая x , если точка x черная) не меньше расстояния от x до начала координат. Найти математическое ожидание числа $|S|$ элементов множества S .

Ч-3.137

Неотрицательная случайная величина ξ имеет функцию распределения $F(x) = P\{\xi \leq x\}$. Доказать, что для любого действительного $\alpha \neq 0$

$$M\xi^\alpha = |\alpha| \int_0^\infty x^{\alpha-1} (1 - F(x)) dx.$$

Ч-3.138

Случайная величина ξ имеет функцию распределения $F(x) = P\{\xi \leq x\}$. Показать, что если

$$F_{-1}(y) = \sup\{x : F(x) \leq y\}.$$

то для любой интегрируемой функции $h(x)$

$$Mh(\xi) = \int_0^1 h(F_{-1}(y)) dy.$$

Замечание, в 1ю унесу

При решении задач Ч-3.139-Ч-3.143 можно использовать следующее простое замечание.

Если случайные величины $\xi_1, \dots, \xi_n$ независимы и одинаково распределены, а число $k < n$, то набор $\xi_{i_1}, \xi_{i_2}, \dots, \xi_{i_k}$, $1 \leq i_1 < i_2 < \dots < i_k \leq n$, одинаково распределены и для любой измеримой функции $f(x_1, \dots, x_k)$ случайные величины $\eta_{i_1, \dots, i_k} = f(\xi_{i_1}, \dots, \xi_{i_k})$ где $1 \leq i_1, \dots, i_k \leq n, i_r \neq i_s, (r \neq s)$, тоже одинаково распределены.

Ч-3.139

Пусть вероятностное распределение P на плоскости R^2 таково, что если точки $X_1, X_2, X_3 \in R^2$ независимы и имеют распределение P , то

$$P\{X_1, X_2, X_3 \text{ лежат на одной прямой}\} = 0.$$

Найти математическое ожидание угла $X_1 X_2 X_3$.

Ч-3.140

Случайные точки A_1, A_2, A_3, A_4 независимы и имеют равномерное распределение в выпуклой плоской фигуре $C \subset R^2$, площадь которой равна 1. Доказать, что вероятность того, что выпуклая оболочка точек A_1, A_2, A_3, A_4 есть треугольник, в 4 раза больше математического ожидания площади треугольника $A_1 A_2 A_3$.

Ч-3.137

Неотрицательная случайная величина ξ имеет функцию распределения $F(x) = P\{\xi \leq x\}$. Доказать, что для любого действительного $\alpha \neq 0$

$$M\xi^\alpha = |\alpha| \int_0^\infty x^{\alpha-1} (1 - F(x)) dx.$$

Ч-3.138

Случайная величина ξ имеет функцию распределения $F(x) = P\{\xi \leq x\}$. Показать, что если

$$F_{-1}(y) = \sup\{x : F(x) \leq y\}.$$

то для любой интегрируемой функции $h(x)$

$$Mh(\xi) = \int_0^1 h(F_{-1}(y)) dy.$$

При решении задач Ч-3.139-Ч-3.143 можно использовать следующее простое замечание.

Если случайные величины $\xi_1, \dots, \xi_n$ независимы и одинаково распределены, а число $k < n$, то набор $\xi_{i_1}, \xi_{i_2}, \dots, \xi_{i_k}$, $1 \leq i_1 < i_2 < \dots < i_k \leq n$, одинаково распределены и для любой измеримой функции $f(x_1, \dots, x_k)$ случайные величины $\eta_{i_1, \dots, i_k} = f(\xi_{i_1}, \dots, \xi_{i_k})$ где $1 \leq i_1, \dots, i_k \leq n, i_r \neq i_s, (r \neq s)$, тоже одинаково распределены.

Ч-3.139

Пусть вероятностное распределение P на плоскости R^2 таково, что если точки $X_1, X_2, X_3 \in R^2$ независимы и имеют распределение P , то

$$P\{X_1, X_2, X_3 \text{ лежат на одной прямой}\} = 0.$$

Найти математическое ожидание $X_1 X_2 X_3$.

Ч-3.140

Случайные точки A_1, A_2, A_3, A_4 независимы и имеют равномерное распределение в выпуклой плоской фигуре $C \subset R^2$, площадь которой равна 1. Доказать, что вероятность того, что выпуклая оболочка точек A_1, A_2, A_3, A_4 есть треугольник, в 4 раза больше математического ожидания площади треугольника $A_1 A_2 A_3$.

Ч-3.151

Доказать, что если $M|\xi|^h < \infty$, то

$$\lim_{x \rightarrow \infty} x^k P\{|\xi| \geq x\} = 0.$$

Ч-3.152

Показать, что если ξ - действительная случайная величина с конечным математическим ожиданием и $f(x)$ - функция, выпуклая вниз, то

$$Mf(\xi) \geq f(M\xi),$$

в если $f(x)$ выпукла вверх, то

$$Mf(\xi) \leq f(M\xi).$$

Ч-3.153

Показать, что для любых случайных величин $\xi_1, \xi_2, \dots, \xi_h$ с конечными r -ми ($r \geq 1$) моментами справедливо соотношение

$$M|\xi_1 + \dots + \xi_h|^r \leq k^{r-1} (M|\xi_1|^r + \dots + M|\xi_h|^r).$$

Ч-3.154

Доказать, что при $1 \leq r \leq 2$ справедливо неравенство

$$|x+y|^r + |x-y|^r \leq 2(|x|^r + |y|^r), \quad -\infty < x, \quad y < \infty,$$

и с его помощью показать, что если случайные величины ξ и η независимы и распределение η симметрично (т. е. распределения η и $-\eta$ совпадают), то при любом $r, 1 \leq r \leq 2$,

$$M|\xi + \eta|^r \leq M|\xi|^r + M|\eta|^r.$$

Ч-3.155

Показать, что если случайные величины $\xi_1, \dots, \xi_n$ независимы, имеют симметричные распределения и $M|\xi_i|^r < \infty, i = 1, \dots, n$, для некоторого $r \in [1, 2]$, то

$$M|\xi_1 + \dots + \xi_n|^r \leq M|\xi_1|^r + \dots + M|\xi_n|^r.$$

,

Ч-3.158

Случайный вектор $\xi = (\xi_1, \dots, \xi_k)$ принимает значения в R^k , и существует такой набор чисел

$$(\alpha_0, \alpha_1, \dots, \alpha_k) \neq (0, 0, \dots, 0),$$

что $P\{\alpha_1 \xi_1 + \dots + \alpha_k \xi_k + \alpha_0 = 0\} = 1$. Доказать, что если все элементы матрицы ковариаций $B = \|\sigma_{ijl}\|$ компонент вектора ξ конечны, то:

а) $\det B = 0$;

б) $(\alpha_1, \dots, \alpha_k) B = \left(B(\alpha_2, \dots, \alpha_k)^\top \right)^\top = 0_k$ где $^\top$ - знак транспонирования.

Ч-3.159

Случайный вектор $\xi = (\xi_1, \dots, \xi_h)$ принимает значения в R^h имеет математическое ожидание $m \in R^h$ и матрицу ковариаций $B = \|b_{ij}\|$. Доказать, что: а). матрица B неотрицательно определена, т. е.

$$\sum_{i,j=1}^h b_{ij} a_i a_j = (a, B) \geq 0 \text{ для любого } a \in R^h;$$

б)* если ранг матрицы B равен r , то существует r -мерная гиперплоскость $L_r \subset R^h$, для которой $P\{\xi \in L_r\} = 1$, и $P\{\xi \in L_{r-1}\} < 1$ для любой $(r-1)$ -мерной гиперплоскости $L_{r-1} \subset R^h$.

Ч-3.160

Случайные величины $\xi_1, \xi_2, \dots, \xi_n$ независимы и равномерно распределены в отрезке $[0, 1]$. Найти $P\{\xi_1 + \xi_2 + \dots + \xi_n \leq x\}$ при $0 \leq x \leq 1$.

Ч-3.161

Случайные величины $\xi_1, \xi_2, \dots$ независимы и равномерно распределены в отрезке $[0, 1]$. Определим случайную величину v равной тому значению k , при котором впервые сумма $\xi_1 + \xi_2 + \dots + \xi_k$ превзойдет 1. Найти Mv .

Ч-3.162

Случайные величины $\xi_1, \dots, \xi_n$ независимы; $D\xi_i = \sigma_i^2, i = 1, \dots, n$. При каких $c_1, \dots, c_n$, удовлетворяющих условиям $c_k \geq 0, c_1 + \dots + c_n = 1$, случайная величина $\eta_n = c_1 \xi_1 + \dots + c_n \xi_n$ имеет минимальную дисперсию? Найти минимальную дисперсию.

Ч-3.163

По известному «правилу трех сигм» вероятность отклонения случайной величины от своего математического ожидания более чем на три корня из дисперсии мала. Найти $P\{|\xi - \xi| < 3\sqrt{D}\}$, если ξ имеет:

- нормальное распределение;
- показательное распределение;
- равномерное распределение на отрезке $[-1, 1]$;
- $P\{\xi = -1\} = P\{\xi = 1\} = 1/18, P\{\xi = 0\} = 8/9$;
- распределение Пуассона с $M\xi = 0,09$.

Ч-3.164

Вычислить математическое ожидание и дисперсию ξ_{ij} - независимые случайные величины с $M\xi_i^5 = 0$ и $D\xi_{ij} = \sigma^2$.

Ч-3.165

Пусть матрица $\Xi = \|\xi_i \xi_j\|_{i,j=1}^n$, где $\xi_1, \dots, \xi_n$ - независимые случайные величины, имеющие симметричные распределения (т. е. распределение ξ совпадает с распределением $-\xi, i = 1, \dots, n$). Показать, что если $M|\xi_j|^k < \infty, i = 1, \dots, n$, для целого $k \geq 1$, то $M\Xi^k$ - диагональная матрица.

Ч-3.166

Неотрицательная случайная величина ξ имеет монотонно убывающую выпуклую вниз плотность распределения $f(x), f''(x) > 0$. Что можно сказать о вехах величин $M \sin \xi, M \cos \xi$?

Ч-3.167

Обозначим $\eta_n = \max\{\mu_n, n - \mu_n\}$, где μ_n - число успехов в схеме Бернулли с n испытаниями и с вероятностью успеха p . Найти $\lim_{n \rightarrow \infty} n^{-1} M\eta_n$.

Ч-3.171

Координата x случайной точки A на действительной прямой имеет непрерывную функцию распределения. Найти на этой прямой такую точку B , для которой математическое ожидание длины отрезка AB минимально.

Ч-3.172

Случайные величины ξ, η имеют непрерывную двумерную плотность распределения. Как выбрать точку $B = (x, y) \in R^2$, чтобы величина $\varphi(x, y) = M(|\xi - x| + |\eta - y|)$ была минимальной?

Ч-3.173

Случайная величина ξ имеет конечный второй момент $M\xi^2$. Найти $\min_x M(\xi - x)^2$ и то значение x , при котором этот минимум достигается.

Ч-3.174

Математическое ожидание квадрата расстояния случайной точки $X \in R^2$ от начала координат конечно. Для какой точки A минимально $M|AX|^2$?

Ч-3.175

Уровень весеннего паводка на реке является случайной величиной ξ с непрерывной функцией распределения $F(x) = P\{\xi \leq x\}$. Плотина рассчитана так, чтобы выдерживать паводок уровня не выше z . Предполагая, что уровни паводков в разные годы независимы и одинаково распределены, найти:

а) минимальное значение z , при котором математическое ожидание времени до разрушения плотины паводком будет не меньше $T = 100$ лет;

б) минимальное значение z , при котором вероятность разрушения плотины паводком за $T = 100$ лет будет не больше $\alpha = 1/100$.

Ч-3.176

Наблюдения за уровнями весенних паводков в течение T лет дали значения $\xi_1, \dots, \xi_T$. На реке построена плотина, которая может выдержать паводок, если только его уровень не превосходит $\zeta_T = \max\{\xi_1, \dots, \xi_T\}$. Пусть $\tau_T = \min\{t : \hat{J}_{T++} > \zeta_T\}$ - время до разрушения плотины паводком. Предполагая, что случайные величины $\xi_1, \xi_2, \dots$ независимы и имеют одну и ту же для $M\tau_T, P\{\tau_T \leq u\}$ и численные значения этих величин при $T = 100, u = 10$.

Ч-3.177

Случайные величины ξ_1, ξ_2, ξ_3 независимы и имеют одно и то же распределение с конечным математическим ожиданием, $\xi_{(1)} \leq \xi_{(2)} \leq \xi_{(3)}$ - их вариационный ряд (см. задачу 3.60).

а) Доказать, что

$$M(\xi_{(3)} - \xi_{(1)}) = \frac{3}{2}M|\xi_1 - \xi_2|.$$

б) Доказать, что

$$M\{\min(\xi_{(2)} - \xi_{(1)}, \xi_{(3)} - \xi_{(2)})\} \leq \frac{3}{4}M|\xi_1 - \xi_2|.$$

Ч-3.178

Будем говорить, что случайная величина ξ сосредоточена на отрезке $[a, b]$, если $P\{a \leq \xi \leq b\} = 1$ и при любом $\varepsilon > 0$

$$P\{a \leq \xi < a + \varepsilon\} > 0 \text{ и } P\{b - \varepsilon < \xi \leq b\} > 0.$$

Доказать, что дисперсия случайной величины, сосредоточенной на отрезке длины l , не превосходит $l^2/4$.

Ч-3.183

Случайные величины ξ и η заданы на одном вероятностном пространстве. Описать множество возможных значений $P\{\xi \leq \eta\}$ в следующих случаях:

а) $M\xi = 1, M\eta = 10$;

б) ξ и η одинаково распределены; мерно на $[0, a], a \neq 1$. распределение на отрезке $[0, 1]$. Доказать, что при любом характере зависимости между ξ и η

$$M|\xi - \eta| \leq 1/2.$$

Ч-3.185

Случайные величины ξ и η независимы и имеют равномерное распределение на отрезке $[0, 1]$, а случайная величина ζ удовлетворяет условию

$$P\{\xi = \zeta\} = P\{\zeta = \eta\} = 1/2.$$

кажать совместные распределения ξ, η, ζ , при которых достигаются экстремальные значения $M\zeta$, и пайти максимальное и минимальное возможные значения $M\zeta$.

Ч-3.186

Случайная величина ξ имеет непрерывную функцию распределения $F(x)$, случайная величина χ принимает только значения 0 и 1: $P\{\chi = 1\} = a, P\{\chi = 0\} = 1 - a$. Указать совместные распределения ξ и χ , при которых достигаются экстремаль

Ч-3.187

Случайные величины ξ и η независимы и имеют равномерное распределение на отрезке $[0, 1]$, а случайная величина ζ удовлетворяет условию

$$P\{\xi = \zeta\} = p \leq 1/2, \quad P\{\xi = \eta\} = 1 - p.$$

Указать совместные распределения ξ, η, ζ , при которых достигаются экстремальные значения $M\zeta$, и найти эти экстремальные значения.

МФТИ-Т.31.

Игральная кость подбрасывается n раз. Пусть ξ - число появлений единицы, а η - число появлений шестёрки. Найти коэффициент корреляции этих случайных величин.

Обозначим A_i - выпадение на i -м броске единицы, B_i - шестёрки. $\xi = \sum_{i=1}^n \mathbb{1}_{A_i}, \eta = \sum_{i=1}^n \mathbb{1}_{B_i}$;

$$E\mathbb{1}_{A_i} = E\mathbb{1}_{B_i} = \frac{1}{6};$$

$$\text{cov}(\xi, \eta) = E(\xi - E\xi)(\eta - E\eta) = E(\xi\eta) - E\xi E\eta$$

$$E(\xi\eta) = E\left(\sum_{i=1}^n \mathbb{1}_{A_i}\right)\left(\sum_{k=1}^n \mathbb{1}_{B_k}\right) = E\left(\sum_{i=1}^n \mathbb{1}_{A_i} \mathbb{1}_{B_i}\right) + E\left(\sum_{i \neq k} \mathbb{1}_{A_i} \mathbb{1}_{B_k}\right) = \sum_{i \neq k} E\mathbb{1}_{A_i} \mathbb{1}_{B_k} = \frac{n(n-1)}{36}$$

$$\text{cov}(\xi, \eta) = \frac{n(n-1)}{36} - \frac{n^2}{36} = -\frac{n}{36}$$

$$D\xi = \sum_{i=1}^n D\mathbb{1}_{A_i} = \sum_{i=1}^n (E\mathbb{1}_{A_i}^2 - (E\mathbb{1}_{A_i})^2) = n\left(\frac{1}{6} - \frac{1}{6} \cdot \frac{1}{6}\right) = \frac{5}{36}n$$

$$D\eta = \frac{5n}{36}$$

$$\text{cov}(\xi, \eta) = \frac{-\frac{n}{36}}{\frac{5n}{36}} = -\frac{1}{5}$$

МФТИ-Т.32.

Подбрасывают две игральные кости. Пусть ξ_1 - число очков на первой игральной кости, а ξ_2 - на второй. Определим $\eta_1 = \xi_1 + \xi_2$, $\eta_2 = \xi_1 - \xi_2$. Найти $\text{cov}(\eta_1, \eta_2)$ и выяснить, являются ли η_1 и η_2 независимыми.

$$\text{cov}(\eta_1, \eta_2) = E(\eta_1, \eta_2) - E(\eta_1)E(\eta_2) = E(\eta_1^2) - E(\eta_2^2) - (E\eta_1)^2 + (E\eta_2)^2 = D\eta_1 - D\eta_2 = 0$$

Зависимость следует из контрпримера. Из подсчета количества благоприятных успехов, видим, что

$$P(\eta_1 = 2) = \frac{1}{36}; P(\eta_2 = 0) = \frac{1}{6}$$

и в то же время

$$P(\eta_1 = 1, \eta_2 = 0) = P(\eta_1 = 1, \eta_2 = 1) = \frac{1}{36}$$

МФТИ-Т.33.

Доказать, что если случайные величины ξ и η принимают только по два значения каждая, то из некоррелируемости следует их независимость.

$$\xi \in \{a_1, a_2\}; \eta \in \{b_1, b_2\},$$

$$\text{cov}(\xi, \eta) = E\xi\eta - E\xi E\eta = 0$$

$$E\xi\eta = \sum_{i,k=1}^2 a_i b_k E(\mathbb{1}_{A_i} \mathbb{1}_{B_k}) = \sum_{i,k=1}^2 a_i b_k P(A_i, B_k)$$

$$E\xi E\eta = \left(\sum_{i=1}^2 a_i P(A_i)\right) \left(\sum_{k=1}^2 b_k P(B_k)\right) = \sum_{i,k=1}^2 a_i b_k P(A_i)P(B_k)$$

Поэтому

$$\text{cov}(\xi, \eta) = 0 = \sum_{i,k=1}^2 a_i b_k (P(A_i, B_k) - P(A_i)P(B_k)) \rightarrow P(A_i B_k) = P(A_i)P(B_k)$$

МФТИ-Т.34.

Авария происходит в точке X , которая равномерно распределена на дороге длиной L . Во время аварии машина скорой помощи находится в точке Y , которая также равномерно распределена на дороге. Считая, что X и Y независимы, найдите математическое ожидание расстояния между машиной скорой помощи и точкой аварии.

$$E|X - Y| = \int_{[0,L]^2} |x - y| \rho_X(x) \rho_Y(y) dx dy = \int_{[0,L]^2} \frac{|x - y|}{L^2} dx dy = \frac{2}{L} \int_0^L dx \int_x^L (y - x) dy = \frac{2}{L} \frac{L^3}{6} = \frac{L}{3}$$

6.4.4 Problems About Conditional Distributions**Ч-3.189**

Случайные величины ξ и η независимы;

$$P\{\xi = k\} = P\{\eta = k\} = pq^{k-1},$$

Найти: $q = 1 - p$, $0 < p < 1$, $k = 1, 2, \dots$

- а) $P\{\xi = \eta\}$; б) $P\{\xi > \eta\}$; в) $P\{\xi < \eta\}$; г) $P\{\xi = k | \xi > \eta\}$; д) $P\{\xi = k | \xi < \eta\}$; е) $P\{\xi = k | \xi = \eta\}$; ж) $P\{\xi = k | \xi + \eta = l\}$; з) $M(\xi | \xi + \eta = l)$, $l \geq 2$.

Ч-3.190

Найти распределение целочисленной неотрицательной случайной величины ξ , если:

- а) $P\{0 < \xi < \infty\} = 1$, $P\{\xi = k + 1 | \xi > k\} = p$, $k =$

Ч-3.191

Случайные величины ξ , η независимы и одинаково распределены. Найти условную плотность $p_{z|+n=z}(x)$ распределения ξ при условии $\xi + \eta = z$ в следующих случаях: а) ξ и η имеют показательное распределение с плотностью $p(x) = \lambda e^{-\lambda x}$, $x \geq 0$; б) ξ и η равномерно распределены в $[0, 1]$; в) ξ и η имеют распре-

Ч-3.192

Найти условную, $x \geq 0$. случаях а), б), в), определенных в дию $D(\xi | \xi + \eta = z)$ в

Ч-3.193

Случайные величины в предыдущей задаче. паково распределены. Найти $M(\xi | \xi + \eta = z)$. Разобрать отдельно случай, когда ξ имеет дискретное распределение или положительную плотность.

Ч-3.194

Плотность совместного распределения величин ξ, η определяется равенствами: $p_{\xi, \eta}(u, v) = 1$ при $(u, v) \in G$, $p_{\xi, \eta}(u, v) = 0$ при $(u, v) \notin G$, где $G = \{(u, v) : 0 \leq u \leq 2, 0 < v < 1 - \frac{1}{2}u\}$. Найти плотность $p_{\eta|\xi-\varepsilon}(x)$ условного распределения η при условии $\xi = 2$.

Ч-3.195

Пусть $\xi_1, \xi_2, \dots, \xi_n$ - результаты n последовательных испытаний Бернулли, $P\{\xi_i = 1\} = p, P\{\xi_i = 0\} = q = 1 - p$. Доказать, что для любого $k (0 \leq k \leq n)$ условное распределение набора $\xi = (\xi_1, \dots, \xi_n)$ при условии $\xi_1 + \dots + \xi_n = k$ является равномерным на множестве всех C_n^k наборов, состоящих из k единиц и $n - k$ нулей.

Ч-3.196

Случайная величина X_1 имеет биномиальное распределение с параметрами (n, p) , случайная величина X_2 при условии X_1 имеет биномиальное распределение с параметрами (X_1, p) , X_3 - при условии X_2 имеет биномиальное распределение с параметрами (X_2, p) , ..., X_k при условии X_{k-1} имеет биномиальное распределение с параметрами (X_{k-1}, p) . Доказать, что безусловным распределением X_n является биномиальное распределение с параметрами (n, p^k) .

Ч-3.197

Случайные величины $\xi_1, \xi_2, \dots, \xi_N$ независимы и распределены по закону Пуассона с параметрами $\lambda_1, \lambda_2, \dots, \lambda_N$. Найти: а) $P\{\xi_1 + \dots + \xi_k = m | \xi_1 + \dots + \xi_N = n\}$; б) $M\{\xi_1 + \dots + \xi_n | \xi_1 + \dots + \xi_N = n\}$.

Ч-3.19

8. з урны, содержащей M белых и $N - M$ черных шаров, сначала извлекается без возвращения выборка объема n , а затем из этой выборки извлекается без возвращения выборка объема $n_0 < n$. Найти закон распределения числа ξ белых шаров во второй выборке. Зависит ли этот закон от n ?

Ч-3.199

Случайные величины ξ_1, ξ_2, ξ_3 независимы и распределены нормально с одинаковыми параметрами $a = 0, \sigma = 1$; положим $\eta = \frac{\xi_1 + \xi_2 \xi_3}{\sqrt{1 + \xi_3^2}}$. Найти распределение η .

Ч-3.20

2. Пусть $\xi_1, \xi_2, \dots$ последовательность независимых одинаково распределенных случайных величин с математическим ожиданием $a > 0$ и дисперсией $\sigma^2 < \infty$, случайная величина v не зависит от $\xi, \xi, \dots$ и принимает целые положительные значения, $Mv = b, a$

$$\tau = \xi_1 + \xi_2 + \dots + \xi_v.$$

Найти τ и (при дополнительном условии $Dv = \delta^2$) $D\tau$.

Ч-3.20

3. Случайные величины ξ, η независимы. Показать, что для любого x

$$P\{\xi \geq \max(\eta, x)\} \geq P\{\xi \geq \eta\}P\{\xi \geq x\}.$$

Ч-3.204

Случайные величины ξ, η, ζ независимы. Показать, что

$$P\{\xi \geq \max(\eta, \zeta)\} \geq P\{\xi \geq \eta\}P\{\xi \geq \zeta\}.$$

Ч-3.205

Случайные величины $\xi_1, \xi_2, \dots$ независимы и имеют равномерное распределение на отрезке $[0, 1]$. Пусть $\delta_1 = 1$, и если $n \geq 2$, то

$$\delta_n = \begin{cases} 1 & \text{при } \xi_n < \min \{\xi_1, \dots, \xi_{n-1}\}, \\ 0 & \text{в противном случае.} \end{cases}$$

Построим последовательность τ_k моментов появления минимальных значений ξ_n и последовательность их величин η_k , т. е. положим $\tau_0 = 0, \eta_0 = 1, \tau_k \Rightarrow \min \{n : n > \tau_{k-1}, \delta_n = 1\}, \eta_k = \xi_{\tau_k}, k = 1, 2, \dots$

а) Доказать, что при любых $k \geq 1$ и $x \in [0, 1]$ условное распределение η_k при условии $\eta_{k-1} = x$ является равномерным на $[0, x]$.

б) Найти математическое ожидание и дисперсию случайной величины $\xi_h = -\ln \eta_h$.

в) Найти условное распределение, математическое ожидание и дисперсию случайной величины $\Delta_h = \tau_h - \tau_{h-1}$ при условии $\eta_{k-1} = x$.

г) Найти τ_h при $k \geq 1$.

Ч-3.207

Случайная величина ξ имеет нормальное распределение с математическим ожиданием 0 и дисперсией σ^2 . Показать, что при любых $x > 0, C > 0$

$$\mathbf{P} \left\{ \xi - x > \frac{C}{x} \mid \xi > x \right\} < e^{-C/\sigma^2}.$$

Ч-3.20

8. В схеме Бернулли с вероятностью p исхода 1 и вероятностью $q = 1 - p$ исхода 0 найти математическое ожидание числа v_{00} испытаний до первого появления дешифровки из двух нулей. В частности, вычислить $M v_{00}$ при $p = 1/2$.

Ч-3.209

В схеме Бернулли предыдущей задачи найти математическое ожидание числа v_{111} испытаний до первого появления цепочки из трех единиц. В частности, вычислить это математическое ожидание при $p = 1/2$.

Ч-3.210

В схеме Бернулли задачи п3.208 найти математическое ожидание числа v_{01} испытаний до первого появления цепочки 01. В частности, вычислить это математическое ожидание при $p = 1/2$.

Ч-3.211

Точки $A_1, A_2, \dots, A_n$ независимы и имеют равномерное распределение на окружности единичной длины. Найти вероятность того, что длина наименьшей дуги, содержащей все эти точки, не больше $x \leq 1/2$.

Ч-3.212

Точки $A_1, \dots, A_n$ независимы и имеют равномерное распределение на окружности S с центром O . Найти вероятность того, что O лежит внутри выпуклого многоугольника с вершинами $A_1, \dots, A_n$.

Ч-3.213

Точки $A_1, \dots, A_n$ независимы и имеют равномерное распределение на окружности с центром O , а случайная величина v равна наименьшему n , при котором выпуклый многоугольник с вершинами $A_1, \dots, A_n$ содержит O . Найти Mv и Dv .

Ч-3.216

Пусть выполнены условия задачи п3.215 и ξ_i длина дуги $A_{(i)}A_{(i+1)}$. Найти совместное распределение и коэффициент корреляции ξ_1, ξ_2 .

Ч-3.217

Пусть выполнены условия задачи п3.216. Найти закон совместного распределения $\xi_1, \dots, \xi_n$ при $k \leq n$.

Ч-3.21

8. Пусть выполнены условия задачи п3.216. Показать, что совместное распределение $\xi_{i_1}, \xi_{i_2}, \dots, \xi_{i_k}$ с $1 \leq i_1 < i_2 < \dots < i_k \leq n$ совпадает с распределением

Ч-3.219

Пусть выполнены условия задачи п3.216 и v - число дуг $A_{(i)}A_{(i+1)}$ длины больше Δ , $0 \leq \Delta \leq 2\pi r$. Найти Mv , Dv и асимптотические формулы для них при $r = n/2\pi\lambda$, $n \rightarrow \infty$.

Ч-3.220

Пусть выполнены условия задачи п3.216. Найти закон распределения случайной величины

$$\eta_n = \max \{ \xi_1, \dots, \xi_n \}.$$

Ч-3.221

Пусть выполнены условия задачи п3.220. Найти $M\eta_n$ и асимптотическую формулу для η_n при $n \rightarrow \infty$.

Ч-3.222

Точки A, B, C независимы и имеют равномерное распределение на окружности единичного радиуса. Пусть S - площадь $\triangle ABC$, p - его периметр, r - радиус вписанного круга. Найти MS , Mp , Mr .

Ч-3.223

Стороны прямоугольника $ABCD$ параллельны осям координат. Найти математическое ожидание площади S прямоугольника $ABCD$ в следующих случаях:

а) координаты (a_1, a_2) точки A фиксированы, $0 \leq a_1, a_2 \leq 1$, а точка C имеет равномерное распределение на диагонали единичного квадрата, соединяющей его вершины $(0, 0)$ и $(1, 1)$;

б) точки A и C независимы; точка A равномерно распределена в единичном квадрате $[0, 1] \times [0, 1]$, а точка C имеет то же распределение, что в п. а).

МФТИ-Т22 Совместное распределение по заданным вероятностям

Найдем совместное распределение случайных величин ξ и η , принимающих значения $-1, 0, 1$, если известно:

$$P(\xi\eta = 0) = 1; \quad P(\xi = 1) = P(\xi = -1) = P(\eta = 1) = P(\eta = -1) = \frac{1}{4}$$

Из условия следует, что $P(\xi \neq 0) = \frac{1}{2}$, $P(\eta \neq 0) = \frac{1}{2}$, поэтому $P(\xi = 0) \leq \frac{1}{2}$, $P(\eta = 0) \leq \frac{1}{2}$, то есть $P(\xi = 0) + P(\eta = 0) \leq 1$.

В то же время $1 = P(\xi\eta = 0) = P((\xi = 0) \cup (\eta = 0)) \leq P(\xi = 0) + P(\eta = 0)$, поэтому $P(\xi = 0) = \frac{1}{2}$, $P(\eta = 0) = \frac{1}{2}$.

Для плотности распределения получаем таблицу:

$\xi \downarrow; \eta \rightarrow$	-1	0	1
-1	0	1/4	0
0	1/4	1/2	1/4
1	0	1/4	0

Table 9: T22

А плотность распределения имеет вид:

(???)

Она заполняется легко: нули ставятся, так как если бы был не ноль в этих местах, то не было бы $P(\xi\eta = 0) = 1$.

Дальше оставшиеся равенства из условия дают нам заполнение боковых средних ячеек.

Ячейка посередине легко находится.

	$\xi < -1$	$\xi \in [-1; 0)$	$\xi \in [0, 1)$	$\xi \leq 1$
$\eta < -1$	0	0	0	0
$\eta \in [-1; 0)$	0	0	1/4	1/4
$\eta \in [0, 1)$	0	1/4	3/4	1
$\eta \leq 1$	0	1/4	1	1

Table 10: T22

МФТИ-Т23 Закон распределения $\xi + \eta$ при известных плотности ξ и вероятности η (??)

Найти закон распределения случайной величины $\xi + \eta$, если ξ и η независимы; ξ имеет плотность распределения $f_\xi(x)$, а $\eta \in \{-1, 0, 1\}$, $\mathbf{P}(\eta = 0) = \mathbf{P}(\eta = 1) = \mathbf{P}(\eta = -1) = \frac{1}{3}$.

Имеем:

$$\begin{aligned} F_{\xi+\eta}(x) &= P(\xi + \eta \leq x) = \\ &= P(\eta + \xi \leq x | \eta = 0)P(\eta = 0) + P(\eta + \xi \leq x | \eta = 1)P(\eta = 1) + P(\eta + \xi \leq x | \eta = -1)P(\eta = -1) \end{aligned}$$

$$P(\xi \leq x | \eta = 0) = \frac{P(\{\xi \leq x\} \cap (\eta = 0))}{P(\eta = 0)} = \frac{P(\xi \leq x)P(\eta = 0)}{P(\eta = 0)} = P(\xi \leq x)$$

Поэтому

$$\begin{aligned} F_{\xi+\eta}(x) &= 1 \frac{1}{3} (F_\xi(x) + F_\xi(x-1) + F_\xi(x+1)) = \\ &= \frac{1}{3} \left(\int_{-\infty}^x f_\xi(y) dy + \int_{-\infty}^{x-1} f_\xi(y) dy + \int_{-\infty}^{x+1} f_\xi(y) dy \right) \end{aligned}$$

Это и ответ.

МФТИ-Т.24 (?!!!) Функция распределения и плотность суммы координат точки в квадрате

В квадрат $\{(x_1, x_2) : 0 \leq x_i \leq 1; i = 1, 2\}$ наудачу брошена точка. Пусть ξ_1, ξ_2 - ее координаты. Найти функцию распределения и плотность случайной величины $\eta = \xi_1 + \xi_2$

$$F_{\xi_i}(x) = P(\xi_i \leq x) = \begin{cases} 0, & \text{если } x < 0 \\ x, & \text{если } x \in [0, 1) \\ 1 & \text{если } x \in [1, +\infty) \end{cases}$$

$$\rho_{\xi_i} = \begin{cases} x, & \text{если } x \in [0, 1) \\ 1, & \text{иначе} \end{cases}$$

ξ_1, ξ_2 - независимые случайные величины.

$$\rho(\xi_1 + \xi_2) = \rho(\xi_1)\rho(\xi_2)$$

Посчитаем

$$\rho(\xi_1 + \xi_2 \leq x) = \iint_S \rho_{\xi_1 \xi_2}(u, s) du ds$$

где S - площадь, которую пересекает заштрихованная часть на рисунке с единичным квадратом

При $x < 0$ $\rho(\xi_1 + \xi_2 \leq x) = 0$.

При $x > 0$ $\rho(\xi_1 + \xi_2 \leq x) = 1$.

При $x \in [0, 1]$

$$\rho(\xi_1 + \xi_2 \leq x) = \int_0^x d\xi_1 \int_0^{x-\xi_1} d\xi_2 = \int_0^x (x - \xi) d\xi = x^2 - \frac{x^2}{2} = \frac{x^2}{2}$$

При $x \in [1, 2]$ нужно посчитать площадь, заштрихованную розовым, для этого от площади квадрата отнимем оставшуюся площадь, которая легко находится, зная длину сторон этой области.

$$\rho(\xi_1 + \xi_2 \leq x) = 1 - \frac{(2-x)^2}{2} = 2x - \frac{x^2}{2} - 1$$




Figure 6.4: МФТИ-Точка в квадрате

В итоге функция распределения имеет вид:

$$F_{\xi_1+\xi_2}(x) = P(\xi_1 \leq x) = \begin{cases} 0, & \text{если } x < 0 \\ \frac{x^2}{2}, & \text{если } x \in [0, 1) \\ 2x - \frac{x^2}{2} - 1, & \text{если } x \in [1, 2) \\ 1 & \text{если } x \in [2, +\infty) \end{cases}$$

Найдем плотность распределения:

$$f_{\xi_1+\xi_2}(x) = P'(\xi_1 \leq x) = \begin{cases} 0, & \text{если } x < 0 \\ x, & \text{если } x \in [0, 1) \\ 2 - x, & \text{если } x \in [1, 2) \\ 0 & \text{если } x \in [2, +\infty) \end{cases}$$

МФТИ-? (?!?) распределение суммы случайных величин по условному распределению одной и $\xi_1 + \xi_2$

Пусть ξ_1, ξ_2 - независимые случайные величины с распределением Пуассона. Найти распределение их суммы и условное распределение ξ_1 , если известна сумма $\xi_1 + \xi_2$.

$$\begin{aligned} P(\xi_1 + \xi_2 = k) &= \sum_{i=0}^k P(\xi_1 = i; \xi_2 = k - i) \stackrel{\text{независимость}}{=} \sum_{i=0}^k P(\xi_1 = i)P(\xi_2 = k - i) = \\ &= \sum_{i=0}^k \frac{\lambda_1^i \lambda_2^{k-i}}{i!(k-i)!} e^{-\lambda_1 - \lambda_2} = \frac{1}{k!} e^{-\lambda_1 - \lambda_2} \sum_{i=0}^k C_k^i \lambda_1^i \lambda_2^{k-i} = \frac{(\lambda_1 + \lambda_2)^k}{k!} e^{-(\lambda_1 + \lambda_2)} \end{aligned}$$

Теперь найдем

$$P(\xi_1 = k | \eta = n) = \frac{P(\xi_1 = k \cap \xi_2 = n - k)}{P(\eta = n)} = \frac{\xi_1(k) \xi_2(n - k)}{\eta(n)} = C_n^k \frac{(\lambda_1 / \lambda_2)^k}{(\lambda_1 / \lambda_2 + 1)^n}$$

МФТИ-? (?!?) Распределение $\eta = f_\xi(\xi)$

Известно, что случайная величина ξ имеет строго возрастающую непрерывную функцию распределения $F_\xi(x)$. Найти распределение случайной величины $\eta = F_\xi(\xi)$.

Из непрерывности следует, что $\exists F_\xi^{-1}(x)$, поэтому просто преобразуем искомую функцию распределения:

$$F_\eta(x) = P(\eta \leq x) = P(F_\xi(\xi) \leq x) = P(\xi \leq F_\xi^{-1}(x)) = F_\xi(F_\xi^{-1}(x)) = x.$$

В итоге

$$F_\xi(\xi) = \begin{cases} 0, & \text{если } x < 0 \\ x, & \text{если } x \in [0, 1] \\ 1, & x \in [1, +\infty) \end{cases}$$

МФТИ-? Параметры ξ^2 по ξ

Пусть ξ имеет стандартное нормальное распределение. Найти функцию распределения и плотность случайной величины ξ^2

Функция распределения:

$$F_{\xi^2}(x) = P(\xi^2 \leq x) = P(\xi \leq \sqrt{x}) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x}$$

А ее плотность есть производная $F_{\xi^2}(x)$:

$$f_{\xi^2}(x) = -\frac{1}{2\sqrt{2\pi}} e^{-\frac{1}{2}x}$$

6.4.5 Problems About Normal Distribution**Ч-3.225**

Случайная величина ξ имеет нормальное распределение с математическим ожиданием 0 и дисперсией 1. Какое из двух событий: $\{|\xi| \leq 0,7\}$ или $\{|\xi| \geq 0,7\}$ - имеет большую вероятность?

Ч-3.226

Случайная величина ξ имеет нормальное распределение с параметрами $(0, 1)$. Что больше:

$$P\{-0,5 \leq \xi \leq -0,1\} \text{ или } P\{1 \leq \xi \leq 2\}?$$

Ч-3.227

Случайная величина ξ имеет нормальное распределение с математическим ожиданием 0 и дисперсией σ^2 . Найти $M\xi^k, k = 1, 2, \dots$

Ч-3.228

Случайная величина ξ имеет нормальное распределение с математическим ожиданием 0 и дисперсией σ^2 . Показать, что при любом $x > 0$

$$\frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma^2}} \left\{ \frac{\sigma}{x} - \frac{\sigma^3}{x^3} \right\} < P\{\xi \geq x\} < \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma^2}} \frac{\sigma}{x}.$$

Ч-3.229

Случайная величина ξ распределена нормально с параметрами (a, σ^2) . Найти:

- плотность распределения - величины $\eta_1 = \xi^2$ при $a = 0$
- плотность распределения величины $\eta_2 = e^3$ при произвольных $a, 0$.

Ч-3.230

Случайные величины ξ_1 и ξ_2 независимы и имеют нормальное распределение с параметрами $(0, \sigma^2)$. Найти функцию распределения случайной величины $\eta = \xi_1^2 + \xi_2^2$

Ч-3.231

Случайные величины $\xi_1, \dots, \xi_n$ независимы и имеют нормальное распределение с математическим ожиданием 0 и дисперсией 1. Распределение случайной величины $\eta_n = \xi_1^2 + \dots + \xi_n^2$ называется χ^2 -распределением с n степенями свободы. Найти плотность распределения η_n .

Какое из распределений, перечисленных в конце введения к гл. 3, при соответствующем выборе параметров совпадает с распределением η_n ?

Ч-3.232

Случайная величина η_n имеет χ^2 -распределение с n степенями свободы (см. задачу п3.231). Найти $M\eta_n, D\eta_n$

Ч-3.233

Найти $M\xi$ и $D\xi$, если $\ln \xi$ имеет нормальное распределение с параметрами (a, σ^2) (в этом случае говорят, что ξ имеет логарифмически нормальное распределение).

Ч-3.234

Для случайной величины ξ , определенной в задаче **3.233**, найдите точку, в которой максимальна плотность распределения ξ (эта точка называется модой распределения). Найдите отношение математического ожидания ξ к его моде.

Ч-3.235

Некоторая категория людей имеет средний вес m кг и среднее квадратическое отклонение веса 3 кг. Для случаев $m = 60$ и $m = 10$ определить вероятность того, что вес случайно взятого человека отличается от m не более чем на 5 кг; если: а) вес имеет нормальное распределение; б) вес имеет логарифмически нормальное распределение.

Ч-3.236

Для случайных величин η_1, η_2 , определенных в п3.229, найти $M\eta_1^k, M\eta_2^k$.

Ч-3.237

Случайная величина ξ имеет нормальное распределение с нулевым математическим ожиданием и единичной дисперсией. Найти $M\xi \cos \xi$, $M \frac{\xi}{1+\xi^2}$, $M \sin \xi$.

Ч-3.242

Случайная величина η нормально распределена с параметрами $(0, 1)$. Положим

$$\eta = \begin{cases} \xi, & \text{если } |\xi| \leq 1, \\ -\xi, & \text{если } |\xi| > 1. \end{cases}$$

- а) Найти закон распределения η .
б) Имеет ли величина $\xi + \eta$ нормальное распределение?

Ч-3.243

Случайные величины ξ и η независимы и имеют нормальное распределение с математическим ожиданием 0 и дисперсией 1. Найти

$$P\{|\xi - \eta| \leq 1\}.$$

Ч-3.244

Случайные величины ξ_1, ξ_2, ξ_3 независимы и нормально распределены с параметрами $(1, 1)$, $(2, 5)$, $(0, 7)$ соответственно. Найти: а) $P\{2\xi_1 - \xi_2 < 0\}$, б) $P\{-3 < 2\xi_1 - \xi_2 < 5\}$, в) $P\{1 < 2\xi_1 - \xi_2 + \xi_3 < 4\}$.

Ч-3.245

Случайный вектор (ξ_1, ξ_2) имеет сферически симметричное нормальное распределение с $D\xi_1 = D\xi_2 = \sigma^2$. Найти распределение вектора (ζ_1, ζ_2) , если

$$\xi_1 = \zeta_1 \cos \varphi + \zeta_2 \sin \varphi, \quad \xi_2 = -\zeta_1 \sin \varphi + \zeta_2 \cos \varphi.$$

симметричное нормальное распределение с $D_1 = D_2 = 1$. Доказать, что случайные величины $\xi_1 \xi_2$ и $\frac{1}{2}(\xi_1^2 - \xi_2^2)$ одинаково распределены.

Ч-3.247

Случайные величины ξ и η независимы и имеют нормальное распределение с параметрами $(0, 1)$. Найти совместное распределение случайных величин $\zeta_1 = a\xi_1 + b\xi_2$, $\zeta_2 = a\xi_1 - b\xi_2$ при $a, b \neq 0$.

Ч-3.250

Случайные величины ξ_1 и ξ_2 независимы и имеют нормальное распределение с математическим ожиданием 0 и дисперсией σ^2 . Случайные величины η_1 и η_2 определяются соотношением

$$\eta_1 + i\eta_2 = \frac{(\xi_1 + i\xi_2)^k}{(\xi_1^2 + \xi_2^2)^{(k-1)/2}}$$

где $i = \sqrt{-1}$, а $k > 0$ - целое число. Найти совместное распределение величин η_1 и η_2 .

Ч-3.251

Случайные величины ξ_1, ξ_2 независимы и имеют нормальное распределение с параметрами $(0, 1)$. Найти математическое ожидание величины

$$\eta = e^{(\xi_1^2 + \xi_2^2)/2} (1 + \xi_1^2 + \xi_2^2)^{-3/2}$$

Ч-3.252

Случайные величины ξ, η независимы и нормально распределены с параметрами $(0, \sigma_1^2), (0, \sigma_2^2)$ соответственно. Вычислить при $\sigma_1 = 1, \sigma_2 = 2$ вероятность попадания случайной точки $(\xi, \eta) \in R^2$ в следующие области: а) прямоугольник $|x| \leq 1, |y| \leq 2$;

б) прямоугольник $0 \leq y \leq 4$;

г) трапецию $x + y \leq 0, |x| \leq 1, y \geq -2$;

д) область $\frac{x^2}{a^2} + \frac{y^2}{b^2} \leq 1$, ограниченную эллипсом, вписанным в прямоугольник $|x| \leq 1, |y| \leq 2$;

е) область $\frac{x^2}{c^2} + \frac{y^2}{4c^2} \leq 1$, ограниченную эллипсом, описанным около прямоугольника $|x| \leq 1, |y| \leq 2$.

Ч-3.253

Мост через реку представляет собой прямоугольник, координаты которого в декартовой системе координат удовлетворяют неравенствам: $|x| \leq 10, |y| \leq 100$. При артиллерийском обстреле моста точка попадания снаряда (ξ, η) в той же системе координат имеет «вумерное нормальное распределение с независимыми координатами и со средними квадратическими отклонениями $\sigma_\xi = 10, \sigma_\eta = 40$. «Точкой прицеливания» назовем $(M\xi, M\eta)$. Определить вероятность попадания в мост при одном выстреле, если точка прицеливания равна: а) $(0, 0)$; б) $(10, 0)$; в) $(5, 20)$.

Ч-3.257

Случайная точка (ξ, η) имеет сферически симметричное нормальное распределение с $D\xi = D\eta = 1$. Найти вероятность попадания (ξ, η) :

а) в квадрат $C = \{(x, y) : |x| \leq 3, |y| \leq 3\}$;

б) во вписанный в C круг;

в) в описанный около C круг.

Ч-3.258

Случайные величины ξ, η независимы и нормально распределены с $M = M\eta = 0, D\xi = D\eta = 4$. Найти вероятность того, что случайная точка (ξ, η) попадет в:

а) кольцо $\{(x, y) : 2 \leq \sqrt{x^2 + y^2} \leq 3\}$; б) область $\{(x, y) : 2 \leq \min(|x|, |y|), \max(|x|, |y|) \leq 3\}$;

в) область $\{(x, y) : 2 \leq |x| + |y| \leq 3\}$.

Ч-3.259

Случайные точки $A_1 = (\xi_1, \eta_1)$ и $A_2 = (\xi_2, \eta_2)$ на плоскости R^2 независимы и имеют сферически симметричное нормальное распределение с единичной матрицей ковариаций. Найти функцию распределения длины отрезка A_1A_2 .

Ч-3.260

Случайные точки $A_1 = (\xi_1, \eta_1), A_2 = (\xi_2, \eta_2), A_3 = (\xi_3, \eta_3)$ на плоскости R^2 независимы и имеют нормальное распределение с нулевым вектором математических ожиданий и единичной матрицей ковариаций. Найти функцию распределения длины медианы A_1M_1 треугольника $A_1A_2A_3$.

Ч-3.261

Доказать, что в условиях задачи п3.260 длина сторон A_2A_3 и длина медианы A_1M_t треугольника $A_1A_2A_3$ - независимые случайные величины.

Ч-3.262

В условиях задачи п3.260 найти вероятность того, что треугольник $A_1A_2A_3$ - тупоугольный.

Ч-3.264

Случайный вектор $\xi = (\xi_1, \xi_2) \in R^2$ имеет двумерное нормальное распределение с нулевым вектором математических ожиданий, область A - угол с вершиной в начале координат и раствором α . Доказать, что если область A' симметрична области A относительно начала координат, то $P\{\xi \in A'\} = P\{\xi \in A\}$.

Ч-3.265

Случайный вектор $\xi = (\xi_1, \xi_2) \in R^2$ имеет двумерное нормальное распределение с нулевым вектором математических ожиданий и матрицей ковариаций $\begin{pmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{pmatrix}$. Найти $P\{|\xi_1| > a|\xi_2|\}, a > 0$.

Ч-3.266

Случайный вектор $\xi = (\xi_1, \xi_2) \in R^2$ имеет невырожденное двумерное нормальное распределение с нулевым вектором математических ожиданий и матрицей ковариаций $\|\sigma_{ij}\|_{i,j=1}^2, |\sigma_{12}|^2 < \sigma_{11}\sigma_{22}$. Найти

$$\begin{aligned} p_{00} &= P\{\xi_1 \geq 0, \xi_2 \geq 0\}, & p_{01} &= P\{\xi_1 \geq 0, \xi_2 \leq 0\}, \\ p_{10} &= P\{\xi_1 \leq 0, \xi_2 \geq 0\}, & p_{11} &= P\{\xi_1 \leq 0, \xi_2 \leq 0\}. \end{aligned}$$

Ч-3.267

Случайный вектор $\xi = (\xi_1, \xi_2) \in R^2$ имеет двумерное нормальное распределение с нулевым вектором математических ожиданий и матрицей ковариаций $\begin{pmatrix} \sigma^2 & \alpha \\ \alpha & \sigma^2 \end{pmatrix}$ $\xi \alpha < \sigma^2$. Найти

$$P\{0 \leq \xi_1 \leq x\xi_2\}, \quad x > 0.$$

Ч-3.268

Случайный вектор $\xi = (\xi_1, \xi_2) \in R^2$ имеет невырожденное нормальное распределение с нулевым вектором математических ожиданий и матрицей ковариаций $\begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{12} & \sigma_{22} \end{pmatrix}$. Найти:

$$\text{а) } P\{\xi_1 > a\xi_2\}, -\infty < a < \infty; \text{ б) } P\{\xi_1 > a\xi_2 + b\}, -\infty < a, b < \infty.$$

Ч-3.271

Случайные величины ξ_1, ξ_2, ξ_3 и $\xi_{(1)}, \xi_{(2)}, \xi_{(3)}$ те же, что в задаче п3.270.

- Найти распределение вектора $(\xi_2 - \xi_1, \xi_3 - \xi_2)$.
- Найти плотность $p(x_1, x_2)$ распределения вектора $(\xi_{(2)} - \xi_{(1)}, \xi_{(3)} - \xi_{(2)})$.
- Найти распределение случайной величины $\xi = (\xi_{(2)} - \xi_{(1)}) / (\xi_{(3)} - \xi_{(2)})$.

Ч-3.272

Случайные величины ξ_1, ξ_2, ζ независимы и имеют стандартное нормальное распределение. Положим

$$\xi_3 = \begin{cases} \zeta, & \text{если } \xi_1\xi_2 = 0, \\ |\zeta| \operatorname{sgn}(\xi_1\xi_2), & \text{если } \xi_1\xi_2 \neq 0, \end{cases}$$

где $\operatorname{sgn}(x) = x/|x|$ при $x \neq 0$.

- Найти вектор математических ожиданий и матрицу ковариаций вектора (ξ_1, ξ_2, ξ_3) .
- Найти распределение ξ_3 и совместные распределения векторов (ξ_1, ξ_3) и (ξ_2, ξ_3) .
- Найти плотность $p(x_1, x_2, x_3)$ распределения вектора (ξ_1, ξ_2, ξ_3) . Н является ли оно нормальным?

Ч-3.273

Используя предыдущую задачу, построить пример случайного вектора $(\xi_1, \xi_2, \dots, \xi_n)$, распределение которого не является нормальным, но любые $n-1$ его координат взаимно независимы и имеют нормальное распределение с математическим ожиданием 0 и дисперсией 1.

Ч-3.274

Случайный вектор $\xi = (\xi_1, \dots, \xi_k)$ имеет сферически симметричное нормальное распределение с нулевым вектором математических ожиданий и единичной ковариационной матрицей. Найти распределение вектора $\eta = \xi A$, где A -действительная матрица с k строками и m столбцами.

Ч-3.277

Случайные величины $\xi_1, \xi_2, \dots, \xi_n$ независимы; ξ_i имеет нормальное распределение с параметрами $(0, \sigma_i^2)$, $i = 1, 2, \dots, n$. Найти распределение вектора

$$\eta = (\xi_1, \xi_1 + \xi_2, \xi_1 + \xi_2 + \xi_3, \dots, \xi_1 + \xi_2 + \dots + \xi_n).$$

Ч-3.278

Случайные величины $\xi_1, \xi_2, \dots$ независимы и имеют нормальное распределение с параметрами $(0, 1)$. Найти минимальное значение k , при котором

$$P\{\max\{|\xi_1|, |\xi_2|, \dots, |\xi_k|\} \geq 2\} \geq \frac{1}{2}.$$

Ч-3.279

Случайный вектор $\xi = (\xi_1, \dots, \xi_k)$ имеет k -мерное нормальное распределение с математическим ожиданием $a \in R^A$ и ковариационной матрицей A . Что больше: $\ln |M\xi_1 \dots \xi_k|$, $M \ln |\xi_1 \dots \xi_k|$ или $\ln M |\xi_1 \dots \xi_k|$?

Ч-3.280

Случайный вектор $\xi = (\xi_1, \xi_2) \in R^2$ имеет двумерное нормальное распределение с нулевым вектором математических ожиданий и ковариационной матрицей $\begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{21} & \sigma_{22} \end{pmatrix}$. Найти условное распределение ξ_1 при условии $\xi_2 = x$.

Ч-3.281

Распределение случайного вектора $\xi = (\xi_1, \xi_2) \in R^2$ таково, что при любом $x \in (-\infty, \infty)$ условное распределение ξ_1 при условии $\xi_2 = x$ и условное распределение ξ_2 при условии $\xi_1 = x$ нормальны. Является ли нормальным распределение вектора ξ ?

6.5 Problems About Limit Theorems That Produce and Characteristic Functions

6.5.1 Problems About the law of Large Numbers and the Borel-cantelli Lemma (????)

(уже года 2 как знаю, что такое есть, но только сейчас до задач дошел.)

Ч-4.1

. Пусть функция $g(x), x \geq 0$, неотрицательна и монотонно возрастает. Показать, что для любой действительной случайной величины ξ справедливо неравенство

$$P\{|\xi| \geq x\} \leq \frac{Mg(|\xi|)}{g(x)}.$$

Ч-4.2

. Пусть случайная величина η_n равна сумме очков, появившихся при n бросаниях симметричной игральной кости. Используя неравенство Чебышева, оценить сверху

$$P \left\{ \left| \frac{\eta_n}{n} - 3,5 \right| > \varepsilon \right\}, \quad \varepsilon > 0.$$

Ч-4.3

. Пусть $\xi_1, \xi_2, \dots, \xi_{n+1}$ - результаты $n+1$ испытаний схемы Бернулли ($P\{\xi_i = 1\} = p, P\{\xi_i = 0\} = 1 - p$) и случайная величина η_n равна числу таких $i, 1 \leq i \leq n$, что $\xi_i = \xi_{i+1} = 1$. Используя неравенство Чебышева, оценить сверху

$$P \left\{ \left| \frac{\eta_n}{n} - p^2 \right| > \varepsilon \right\}, \quad \varepsilon > 0.$$

Ч-4.4

Последовательности $\xi_1, \xi_2, \dots$ и $\eta_1, \eta_2, \dots$ образованы одинаково распределенными случайными величинами, независимыми внутри каждой последовательности (случайные величины ξ_i и η_i могут не быть независимы для больших чисел для последовательности $\zeta_1, \zeta_2, \dots$

$$\xi_{2h-1} = \xi_h, \quad \zeta_{2h} = \eta_h, \quad k = 1, 2, \dots$$

Ч-4.5

. Случайные величины $\xi_1, \xi_2, \dots$ независимы и имеют стандартное нормальное распределение,

$$\eta_n = \cos \frac{\varepsilon_n}{\xi_{n+1}}, \quad n = 1, 2, \dots$$

$\bigvee_n$ удовлетворяют ли последовательности $\eta_3, \eta_5, \dots$ и $\eta_1, \eta_2, \eta_3, \dots$ закону больших чисел?

Ч-4.7

Последовательность $\xi_1, \xi_2, \dots$ образована независимыми случайными величинами,

$$M\xi_k = 0, \quad D\xi_k = Ck^\alpha, \quad C > 0, \quad \alpha \geq 0, \quad k = 1, 2, \dots$$

При каких значениях α последовательность $\xi_1, \xi_2, \dots$ может удовлетворять закону больших чисел и при каких значениях α может не удовлетворять ему?

Ч-4.8

Последовательность $\xi_1, \xi_2, \dots$ состоит из независимых одинаково распределенных случайных величин, $M\xi_k = 0, D\xi_k = \sigma^2 < \infty, k = 1, 2, \dots$. Удовлетворяют ли закону больших чисел последовательности случайных величин

$$\zeta_n = \xi_n + \xi_{n+1}, \quad \eta_n = \sum_{k=0}^{\lfloor n/3 \rfloor} \xi_{n+k}, \quad n = 1, 2, \dots?$$

Ч-4.9

Случайные величины $\xi_1, \xi_2, \dots$ независимы, $M\xi_1 = a, D\xi_1 = \sigma^2 < \infty$. Пусть $\zeta_n = \sum_{1 \leq i < j \leq n} \xi_i \xi_j$.

Показать, что последовательность ζ_n удовлетворяет закону больших чисел: для каждого $\varepsilon > 0$

$$\lim_{n \rightarrow \infty} P \left\{ \left| \frac{\zeta_n}{c_n^2} - a^2 \right| > \varepsilon \right\} = 0.$$

Ч-4.10

Показать, что утверждение предыдущей задачи останется справедливым, если в ее формулировке заменить условие независимости $\xi_1, \xi_2, \dots$ условием их некоррелированности:

$$M(\xi_i - a)(\xi_j - a) = 0, \quad 1 \leq i < j < \infty.$$

Ч-4.11

Пусть функция $f(x)$, $0 \leq x \leq 1$, ограничена и интегрируема, а $\xi_1, \xi_2, \dots$ - независимые случайные величины, равномерно распределенные на отрезке $[0, 1]$.

а) Доказать, что для любого $\varepsilon > 0$

$$\lim_{n \rightarrow \infty} P \left\{ \left| \frac{f(\xi_1) + \dots + f(\xi_n)}{n} - \int_0^1 f(x) dx \right| > \varepsilon \right\} = 0.$$

б) Найти

$$\lim_{n \rightarrow \infty} P \left\{ \left| \frac{f(\xi_0) + \dots + f(\xi_n)}{n} - \int_0^1 f(x) dx \right| > \frac{\varepsilon}{\sqrt{n}} \right\}, \quad \varepsilon > 0.$$

114

Ч-4.12

Пусть ограниченная интегрируемая функция $f(x)$, $-\infty < x < \infty$, имеет период 1, а случайные величины ξ и η независимы и равномерно распределены на отрезке $[0, 1]$. Построим последовательность случайных величин

$$\xi_k = f(\xi + k\eta), \quad k = 1, 2, \dots$$

а) Доказать, что случайные величины $\xi_1, \xi_2, \dots$ попарно независимы и одинаково распределены. Найти $M\xi_1, D\xi_1$.

б) Показать, что для любого $\varepsilon > 0$

$$\lim_{n \rightarrow \infty} P \left\{ \left| \frac{\xi_1 + \dots + \xi_n}{n} - \int_0^1 f(x) dx \right| > \varepsilon \right\} = 0.$$

в) Показать, что если $\left| f(y) - \int_0^1 f(x) dx \right| > \varepsilon$ при всех $y \in [a, b] \subset (0, 1)$, то при любом $n = 1, 2, \dots$

$$P \left\{ \left| \frac{\xi_1 + \dots + \xi_n}{n} - \int_0^1 f(x) dx \right| > \varepsilon \right\} \geq 2 \left(\frac{b-a}{n} \right)^2.$$

Ч-4.15

Случайные величины $\xi_1, \xi_2, \dots$ независимы, одинаково распределены, имеют математическое ожидание a и для некоторого ε , $0 < \varepsilon < 1$,

$$m_\varepsilon = M |\xi_1 - a|^{1+\varepsilon} < \infty.$$

Показать, что для любого $\delta > 0$

$$\lim_{n \rightarrow \infty} P \left\{ \left| \frac{\xi_1 + \dots + \xi_n}{n} - a \right| > \delta \right\} = 0.$$

Ч-4.16

Лемма Бореля-Кантелли. Пусть $A_1, A_2, \dots$ - события, заданные на одном вероятностном пространстве, и случайная величина v равна числу одновременно происходящих событий. Показать, что:

а) если $\sum_{n=1}^{\infty} P\{A_n\} < \infty$, то $P\{v < \infty\} = 1$,

б) если события $A_1, A_2, \dots$ попарно независимы и $\sum_{n=1}^{\infty} P\{A_n\} = \infty$, то $P\{v = \infty\} = 1$.

Ч-4.17

Пусть $A_1, A_2, \dots$ события, заданные на одном вероятностном пространстве, и случайная величина v равна числу одновременно происходящих событий. Показать, что если $P\{A_n\} \geq a > 0, n = 1, 2, \dots$ то $P\{v = \infty\} \geq a$.

Ч-4.18

Последовательность чисел c_n удовлетворяет условиям $0 \leq c_n \leq 1, \lim_{n \rightarrow \infty} c_n = 0$. Всегда ли существует такая последовательность событий $A_1, A_2, \dots$, что $P\{A_n\} = c_n$ и для числа v одновременно происходящих событий справедливо соотношение $P\{v < \infty\} = 1$?

Ч-4.21

Случайные величины $\xi_1, \xi_2, \dots$ независимы и имеют одно и то же показательное распределение с параметром λ :

$$P\{\xi_n \leq x\} = 1 - e^{-\lambda x}, \quad x \geq 0, \quad n = 1, 2, \dots$$

Пусть ε - произвольное число из интервала $(0, 1/\lambda)$, случайная величина v_ε равна числу одновременно происходящих событий $A_n^e = \left\{ \frac{\xi_n}{\ln n} > \frac{1}{\lambda} + \varepsilon \right\}$, а μ_ε - числу одновременно происходящих событий

$$B_n^\varepsilon = \left\{ \max_{\sqrt{n} < k < n} \frac{\xi_k}{\ln k} < \frac{1}{\lambda} - \varepsilon \right\}.$$

а) Показать, что $P\{v_e < \infty\} = P\{\mu_e < \infty\} = 1$ при любом $\varepsilon \in (0, 1/\lambda)$.

б) Вывести из результата п. а), что последовательность случайных величин $\xi_n = \frac{\xi_n}{\ln n}$, $n = 1, 2, \dots$, удовлетворяет условию

$$P\left\{\overline{\lim}_{n \rightarrow \infty} \xi_n = \frac{1}{\lambda}\right\} = 1.$$

Ч-4.22

Случайные величины $\xi_1, \xi_2, \dots$ имеют математическое ожидание a и дисперсию $\sigma^2 < \infty$ и все коррелированы:

$$M(\xi_i - a)(\xi_i - a) = 0 \text{ при любых } i \neq 1.$$

Доказать, что

$$p\left\{\lim_{n \rightarrow \infty} \frac{\xi_1 + \dots + \xi_{n^2}}{n^2} = a\right\} = 1.$$

ФР задача на нормальное распределение (!?!?!?)

Two random variables are given: $X \sim \mathcal{N}(0, 9)$ and $Y \sim \mathcal{N}(0, 4)$. $\text{Corr}(X, Y) = -1$. (a) (5 pts) Evaluate $P(X > 1)$ (b) (5 pts) Evaluate $P(2X + Y > 3)$ (c) (5 pts) Evaluate $P(2X + Y > 3 | X = 1)$ (d) (5 pts) Evaluate $P(|X| > 2)$ (e) (5 pts) Evaluate $P((Y + 1)^2 < 2)$

(такую важно решать, но пока плохо понимаю его, не уверен в решениях вообще.)

6.5.2 Limit Theorem Proof Problems**Ч-4.25**

Случайные величины $\xi_1, \xi_2, \dots, \xi_n$ независимы и имеют одно и то же геометрическое распределение с параметром $p, 0 < p < 1$:

$$P\{\xi_1 = k\} = (1-p)p^n, \quad k = 0, 1, 2, \dots$$

а) Показать, что $P\{\xi_1 + \dots + \xi_n = k\} = C_{k+n-1}^{n-1} (1-p)^n p^k$, $k = 0, 1, 2, \dots$

б) Найти

$$q_k^{(n)} = \lim_{m \rightarrow \infty} P\{\xi_1 + \dots + \xi_n - m = k | \xi_1 + \dots + \xi_n \geq m\}_1 = 0, 1, 2, \dots$$

Ч-4.26

(см. задачу п1.53). Первый ряд кинотеатра состоит из N кресел. Зрители один за другим выполняют этот ряд, причем каждый из них может с равной вероятностью занять любое из кресел, свободных в момент его прихода. Пусть $\tau_1(N)$ - порядковый номер первого зрителя, который сел в кресло, находящееся рядом с уже занятым креслом, $\tau_2(N)$ - порядковый номер первого зрителя, который сел в кресло, симметричное относительно середины ряда одному из занятых кресел. Найти законы распределения $\tau_1(N)$ и $\tau_2(N)$ и предельные при $N \rightarrow \infty$ распределения случайных величин $\frac{\tau_i(N)}{VN}, i = 1, 2$, т. е. функции $G_i(x) = \lim_{N \rightarrow \infty} P\left\{\frac{\tau_i(N)}{\sqrt{N}} \leq x\right\}$.

Ч-4.27

. Плотность $p_s(x)$ случайной величины ξ непрерывна и ограничена на отрезке $[a, b]$ и равна 0 вне $[a, b]$. Положим $\eta_n = \{n\xi\}$, где $\{x\}$ - дробная часть числа x . Найти

$$\lim_{n \rightarrow \infty} P\{\eta_n \leq x\}, \quad 0 \leq x \leq 1.$$

Ч-4.28

Показать, что утверждение предыдущей задачи справедливо для случайной величины ξ с кусочно-непрерывной плотностью.

Доказать, что при любом характере зависимости между ξ_n и η_n для каждой точки непрерывности $F(x)$ справедливы равенства:

- а) $\lim_{n \rightarrow \infty} P\{\eta_n + \xi_n \leq x\} = F(x)$;
 б) $\lim_{n \rightarrow \infty} P\{\eta_n(1 + \xi_n) \leq x\} = F(x)$.

Ч-4.34

Случайные величины $\xi_1, \xi_2, \dots$ независимы,

$$P\{\xi_i = 1\} = p, \quad P\{\xi_i = 0\} = 1 - p, \quad i = 1, 2, \dots$$

Положим $\xi_i^{(k)} = \xi_i(1 - \xi_{i+1}\xi_{i+2}\dots\xi_{i+k})$, $\eta_n^{(k)} = \xi_1^{(k)} + \dots + \xi_n^{(k)}$. В случае когда $0 < p < 1$, $\lim_{n \rightarrow \infty} p^{k(n)}\sqrt{n} = 0$ найти

$$\lim_{n \rightarrow \infty} P\left\{\frac{\eta_n^{(h(n))} - np}{\sqrt{np(1-p)}} \leq x\right\}.$$

Ч-4.35

Пусть $\xi_1, \xi_2, \dots$ и $\eta_1, \eta_2, \dots$ - такие последовательности случайных величин, что для некоторых чисел a, b для любого $\varepsilon > 0$

$$\lim_{n \rightarrow \infty} P\{|\xi_n - a| > \varepsilon\} = 0, \quad \lim_{n \rightarrow \infty} P\{|\eta_n - b| > \varepsilon\} = 0.$$

Доказать, что при любом характере зависимости между ξ_n и η_n для любого $\varepsilon > 0$:

- а) $\lim_{n \rightarrow \infty} P\{|\xi_n - a| < \varepsilon, |\eta_n - b| < \varepsilon\} = 1$;
 б) если функция $f(x, y)$ непрерывна в точке (a, b) , то для любого $\varepsilon > 0$

$$\lim_{n \rightarrow \infty} P\{|f(\xi_n, \eta_n) - f(a, b)| < \varepsilon\} = 1.$$

Ч-4.36

Пусть $\xi^{(i)} = (\xi_1^{(i)}, \xi_2^{(i)})$, $i = 1, 2, \dots$, - независимые одинаково распределенные векторы в R^2 ,

$$M\xi_1^{(i)} = a_1 > 0, \quad M\xi_2^{(i)} = a_2 > 0, \quad i = 1, 2, \dots$$

Доказать, что последовательность случайных величин

$$\eta_n = \frac{\xi_1^{(1)} + \xi_1^{(2)} + \dots + \xi_1^{(n)}}{\xi_2^{(1)} + \xi_2^{(2)} + \dots + \xi_2^{(n)}}, \quad n = 1, 2, \dots$$

сходится по распределению к a_1/a_2 , если либо а) $D_{\xi_1^{(i)}} < \infty$ для некоторого $i > 0$.

Ч-4.37

В последовательности случайных величин $\xi_1, \xi_2, \dots$ случайная величина ξ_n при любом $n = 1, 2, \dots$ имеет геометрическое распределение с параметром q_n , $q_n \rightarrow 0$ ($n \rightarrow \infty$). Найти предельное распределение случайных величин

$$\xi_n = q_n \xi_n \text{ при } n \rightarrow \infty.$$

Ч-4.39

Пусть $\xi_1, \xi_2, \dots$ - последовательность независимых одинаково распределенных случайных величин с математическим ожиданием $a > 0$ и дисперсией $\sigma^2 < \infty$, случайная величина v не зависит от $\xi_1, \xi_2, \dots$ и принимает целые положительные значения, $Mv = b$ и

$$\tau = \xi_1 + \xi_2 + \dots + \xi_v.$$

Найти $M\tau$ и предельное распределение случайной величины $\tau/(M\tau)$, когда распределение v изменяется так, что $Mv \rightarrow \infty$ и существует непрерывная функция распределения $F(x)$, удовлетворяющая условиям $F(0) = 0$,

$$\lim P \left\{ \frac{v}{Mv} \leq x \right\} = F(x), \quad 0 \leq x < \infty.$$

Ч-4.40

Время & безотказной работы прибора имеет математическое ожидание $a > 0$ и дисперсию $\sigma^2 < \infty$. Каждый отказ прибора независимо от предыстории и длин Периодов безотказной работы с вероятностью p является несущественным (прибор требует только наладки), а с вероятностью $q = 1 - p$ - существенным (прибор нужно ремонтировать). Пусть τ_q -время от начала работы прибора до его ремонта.

а) Найти $M\tau_q, D\tau_q$.

б) Найти предельное распределение случайной величины $\tau_q/M\tau_q$, когда параметры $a > 0$ и $\sigma^2 < \infty$ фиксированы, а $q \rightarrow 0$.

Ч-4.44

В бункер помещается не более $N = 150$ деталей. Ежеминутно в бункер поступает случайное число деталей, имеющее распределение Пуассона с параметром $\lambda = 2$; числа деталей, поступающих в бункер в непересекающиеся интервалы времени, независимы. Через каждый час все находящиеся в бункере детали перегружаются в тележку и отправляются на дальнейшую обработку. В патальный момент времени бункер свободен. Пользуясь предельными теоремами, указать приближенное значение вероятности того, что 3 время $T = 100$ часов не произойдет ни одного переполнения бункера.

Ч-4.45

Показать, что если последовательность случайных величин $\xi_1, \xi_2, \dots$ сходится к ξ с вероятностью 1 или в среднем квадратичном, то она сходится по вероятности и по распределению. Привести пример последовательности $\xi_1, \xi_2, \dots$, сходящейся к ξ по вероятности, но не сходящейся ни в среднем квадратичном, ни с вероятностью 1.

Ч-4.46

Последовательность случайных величин $\xi_1, \xi_2, \dots$ сходится к ξ по распределению. Показать, что можно задать на одном вероятностном пространстве случайные величины $\xi'_1, \xi'_2, \dots$ и ξ' так, что для всех x

$$P_1 \{ \xi'_1 \leq x \} = P \{ \xi \leq x \}, \quad P \{ \xi'_n \leq x \} = P \{ \xi_n \leq x \}, \\ n = 1, 2, \dots$$

и последовательность $\xi'_1, \xi'_2, \dots$ сходится (??) с вероятностью 1.

Ч-4.48

Последовательность $\xi_1, \xi_2, \dots$ случайных величин удовлетворяет условиям $P \{ \xi_{n+1} \geq \xi_n \} = 1$ для любого $n = 1, 2, \dots$

$$\lim_{n \rightarrow \infty} M\xi_n = a < \infty.$$

Доказать, что существует случайная величина ξ , для которой

$$P \left\{ \lim_{n \rightarrow \infty} \xi_n = \xi \right\} = 1.$$

Верно ли равенство $\xi = a$?

Ч-4.49

Последовательность точек $\xi_1, \xi_2, \dots$ на отрезке $[0, 1]$ строится по следующему правилу: ξ_1 имеет равномерное распределение на $[0, 1]$, и если заданы $\xi_1, \dots, \xi_{k-1}$ ($k \geq 2$) определены, то точка ξ_k имеет равномерное распределение на минимальном по длине из k отрезков, на которые $[0, 1]$ разбивается точками $\xi_1, \dots, \xi_{k-1}$.

а) Доказать, что существует случайная величина ξ , удовлетворяющая условию $P \{ \lim_{n \rightarrow \infty} \xi_n = \xi \} = 1$.

б) Найти $M\xi, D\xi$.

Ч-4.50

Последовательность точек $\xi_0, \xi_1, \xi_2, \dots$ на отрезке $[0, 1]$ строится по следующему правилу: $\xi_0 = 0, \xi_1$ имеет равномерное распределение на $\{0, 1\}$, а ξ_n при любом $n \geq 2$ имеет равномерное распределение на интервале, образованном ξ_{n-1} и ξ_{n-2} .

- Доказать, что существует случайная величина ξ , удовлетворяющая условию $P\{\lim_{n \rightarrow \infty} \xi_n = \xi\} = 1$.
- Найти $M\xi, D$.

Ч-4.51

*. Последовательность случайных величин $\xi, \xi_2, \dots$, последовательности чисел $\{a_n\}, \{b_n\}$ и случайная величина ξ удовлетворяют условиям

$$\lim_{n \rightarrow \infty} a_n = a, \quad |a| < \infty, \quad \lim_{n \rightarrow \infty} b_n = 0, \quad P\{|\xi| < \infty\} = 1$$

$$\lim_{n \rightarrow \infty} P\left\{\frac{\xi_n - a_n}{b_n} \leq x\right\} = P\{\xi \leq x\}$$

в каждой точке непрерывности функции $F(x) = P\{\xi \leq x\}$. Функция $g(x)$ непрерывно дифференцируема в окрестности точки a и $g'(a) \neq 0$. Доказать, что в каждой точке непрерывности $F(x)$

$$\lim_{n \rightarrow \infty} P\left\{\frac{g(\xi_n) - g(a_n)}{b_n g'(a_n)} \leq x\right\} = F(x).$$

Ч-4.52

Последовательность случайных величин $\xi, \xi_2, \dots$, последовательности чисел $\{a_n\}, \{b_n\}$ и случайная величина ξ удовлетворяют тем же условиям, что в задаче Ч-4.51. Функция $g(x)$ имеет $k \geq 2$ непрерывных производных в окрестности точки a

$$g'(a) = g''(a) = \dots = g^{(k-1)}(a) = 0, \quad g^{(k)}(a) \neq 0.$$

оказать, что

$$\lim_{n \rightarrow \infty} P\left\{\frac{g(\xi_n) - g(a_n)}{b_n^k g^{(k)}(a_n)} \leq x\right\} = P\{\xi^k \leq x\}$$

в каждой точке непрерывности функции $F_k(x) = P\{\xi^k \leq x\}$.

Ч-4.53

Случайная величина μ_n равна числу успехов в n независимых испытаниях Бернулли с вероятностью успеха $p, 0 < p < 1$

- Найти $\lim_{n \rightarrow \infty} P\left\{\frac{(\mu_n - np)^2}{np(1-p)} \leq x\right\}$.

6) Пусть $0 < z < 1, z \neq p$. Найти такие последовательности чисел $A_n(z)$ и $B_n(z), n = 1, 2, \dots$, для которых

$$\lim_{n \rightarrow \infty} P\left\{\frac{(\mu_n - ns)^2 - A_n(z)}{B_n(z)} \leq x\right\} = \Phi(x), \quad -\infty < x < \infty.$$

Ч-4.54

". Случайная величина μ_n равна числу успехов в n независимых испытаниях Бернулли с вероятностью успеха p ; функция $g(x)$ определяется соотношениями

$$g(x) = \begin{cases} x \ln x + (1-x) \ln(1-x), & 0 < x < 1, \\ 0, & x = 0 \text{ или } x = 1. \end{cases}$$

- При $p = 1/2$ найти

$$\lim_{n \rightarrow \infty} P\left\{n \left(g\left(\frac{\mu_n}{n}\right) + \ln 2\right) \leq x\right\} \leftarrow x \geq 0.$$

- При $0 < p < 1, p \neq 1/2$, указать такую последовательность $B_n(p), n = 1, 2, \dots$ что

$$\lim_{n \rightarrow \infty} P\left\{B_n(p) \left(g\left(\frac{\mu_n}{n}\right) - g(p)\right) \leq x\right\} = \Phi(x), \quad -\infty < x < \infty.$$

6.5.3 Problems About Characteristic and Producing Functions

В этом параграфе задачи п4.58-п4.83 связаны с вычислением производящих функций, моментов случайных величин и характеристических функций, в задачах п4.84-п4.99 рассматриваются различные свойства характеристических функций, а в задачах п4.100-п4.104 вычисляются характеристические функции специальных распределений и исследуются свойства многомерного нормального распределения.

Ч-4.58

Найти производящие функции целотисленных распределений:

- а) пуассоновского: $P\{\xi = k\} = \frac{\lambda^k}{k!} e^{-\lambda}, \quad k = 0, 1, 2, \dots, 0 < \lambda < \infty$
- б) геометрического: $P\{\xi = k\} = pq^k, k = 0, 1, 2, \dots; p, q > 0, p + q = 1;$
- в) биномиального: $P\{\xi = k\} = C_n^k p^k q^{n-k}, k = 0, 1, \dots, n; p, q > 0, p + q = 1.$

Ч-4.59

Найти производящую функцию $\varphi(z)$ числа ξ_n успехов в n независимых испытаниях, если вероятность успеха в каждом испытании равна p . Используя этот результат, найти формулы для $M\xi_n, M_n^{[k]} (k = 1, 2, \dots)$,

Ч-4.60

Производящая функция распределения случайной величины ξ , принимающей целые неотрицательные значения, равна $\varphi(z) = Mz^\xi$.

- а) Найти $M\xi, D\xi, M(\xi - M\xi)^3$.
- б) Найти производящие функции $\varphi_i(z) = Mz^{\xi_i}$ случайных величин $\xi_1 = \xi/2, \xi_2 = 2\xi, \xi_3 = -\xi$ и $\xi_4 = \xi_1 - \xi_2$ где ξ_1 и ξ_2 - независимые случайные величины, имеющие то же распределение, что ξ .
- в) Доказать, что при $b > a > 0$

$$M \frac{1}{\xi + a} = \int_0^1 z^{a-1} \varphi(z) dz$$

$$M \frac{1}{(\xi + a) \left(\frac{\xi}{b} + 1 \right)} = \int_0^1 z^{b-a-1} \int_0^1 u^{a-1} \varphi(u) du dz.$$

Ч-4.61

Пусть τ_1 - порядковый номер первого из испытаний схемы Бернулли (т. е. последовательности независимых испытаний), которое окончилось успехом, вероятность успеха в каждом испытании равна p , неудачи $-q = 1 - p$. Найти $M\tau_1, D\tau_1, Mz^{\tau_1}$.

Ч-4.62

В схеме Вернулли обозначим через θ_k порядковый номер испытания, в котором появился k -й успех; считая вероятность успеха в каждом испытании равной p , найти: 1) $M\theta_k, D\theta_k$; 2) $\varphi_{\theta_k}(z) = Mz^{\theta_k}$.

Ч-4.63

Закон распределения случайной величины ξ определяется формулой

$$P\{\xi = n\} = C_{n-1}^{m-1} p^m q^{n-m} \quad n = m, m+1, \dots,$$

где m -целое положительное число, $0 < p < 1, q = 1 - p$. Найти производящую функцию распределения суммой m независимых одинаково распределенных случайных величин.

Ч-4.66

Производящая функция совместного распределения величин ξ, ξ_2 равна

$$\varphi_{\xi_1, \xi_2}(z) = e^{\lambda(p_1 z_1 + p_2 z_2 + p_3 z_1 z_2 - 1)}$$

где $p_1 + p_2 + p_3 = 1, p_i \geq 0, i = 1, 2, 3$. Найти:

- а) одномерные распределения ξ и ξ_2 ;
- б) $M\xi_1, M\xi_2, D\xi_1, D\xi_2, \text{cov}(\xi_1, \xi_2)$. Являются ли величины ξ_1 и ξ_2 независимыми?

Ч-4.67

Целочисленная неотрицательная случайная величина ξ имеет производящую функцию $\varphi_{\xi}(z)$. Доказать, что если для целого $k \geq 1$

$$\lim_{u \uparrow 1} \frac{d^k}{dz^k} \varphi_{\xi}(z) \Big|_{z=u} = m_k < \infty$$

то $M\xi^{[k]} = M\xi(\xi-1)\dots(\xi-k+1) = m_k$.

Ч-4.68

Пусть $\xi_{n,i}$ - число появлений i -го исхода в n независимых испытаниях с N несовместными исходами и вероятностью p_j появления j -го исхода в каждом испытании ($j = 1, \dots, N$). Найти

а) $\varphi(z_1, \dots, z_N) = Mz_1^{n,1} \dots z_N^{n,N}$,

Ч-4.69

Пусть выполнены условия предыдущей задачи и $N = 4$,

$$\begin{aligned} \eta_{n,1} &= \xi_{n,1}, & \eta_{n,2} &= \xi_{n,2} + \xi_{n,3}, & \eta_{n,3} &= \xi_{n,4}, \\ \xi_{n,i} &= \xi_{n,1} + \xi_{n,i+1} \quad (i = 1, 2, 3). \end{aligned}$$

Найти производящие функции распределений векторов $(\eta_{n,1}, \eta_{n,2}, \eta_{n,3})$, $(\zeta_{n,1}, \zeta_{n,2}, \zeta_{n,3})$ и коэффициенты корреляции $\rho(\eta_{n,1}, \eta_{n,2})$, $\rho(\zeta_{n,1}, \zeta_{n,2})$.

Ч-4.73

Найти характеристические функции распределений:

- а) пуассоновского: $P\{\xi = m\} = \frac{\lambda^m}{m!} e^{-\lambda}$, $m = 0, 1, \dots$;
 б) биномиального: $P\{\xi = m\} = C_n^m p^m q^{n-m}$, $m = 0, 1, \dots, n$;
 в) показательного: $p_{\lambda}(x) = \alpha e^{-\alpha x}$, $x > 0$.

Ч-4.74

Случайные величины $\xi_1, \xi_2, \dots$ независимы и имеют показательное распределение с параметром α :

$$P\{\xi_i \leq x\} = 1 - e^{-\alpha x}, \quad x \geq 0, \quad i = 1, 2, \dots$$

Найти $M(\xi_1 + \dots + \xi_n)^n$ при любых значениях $k, n = 1, 2, \dots$

Ч-4.75

Случайные величины $\xi_1, \xi_2, \dots$ независимы, одинаково распределены и имеют характеристическую функцию $f(t) = Me^{it_1}$ (производящую функцию $\varphi(s) = Ms^{\xi_1}$), случайная величина v не зависит от $\xi_1, \xi_2, \dots$ принимает только целые неотрицательные значения и имеет производящую функцию $g(s) = Ms^2$. Найти характеристическую функцию (производящую функцию) случайной величины

$$\xi = \begin{cases} \xi_1 + \xi_2 + \dots + \xi_v & \text{при } v \geq 1 \\ 0 & \text{при } v = 0 \end{cases}$$

Ч-4.76

Производящая функция распределения случайной величины ξ_1 принимающей целые неотрицательные значения, равна $\varphi(z) = Mz^{\xi_1}$. Выяснить, при каких условиях следующие выражения являются производящими функциями вероятностных распределений, и если являются, то указать, как соответствующие распределения связаны с распределением ξ :

- а) $\psi(z) = \frac{\varphi'(z)}{\varphi'(1)^r}$ б) $\psi(z) = \frac{1-\varphi(z)}{\varphi'(1)(1-z)^s}$
 в) $\psi(z) = \frac{1-\alpha}{1-\alpha\varphi(z)}$ г) $\psi(z) = 1 - \frac{\ln \varphi(z)}{\ln \varphi(0)^s}$
 а) $\psi(z) = \sqrt{\frac{1-\alpha}{1-\alpha\varphi(z)}}$

Ч-4.80

Случайные величины $\xi_1, \xi_2, \dots, \xi_n$ независимы. Доказать, что для любого действительного λ , удовлетворяющего условиям

$$Me^{\lambda \xi_i} < \infty \quad t = 1, \dots, n$$

справедливо равенство

$$Me^{\lambda(\xi_1 + \dots + \xi_n)} = \prod_{j=1}^n Me^{\lambda \xi_j}.$$

Ч-4.81

Случайная величина принимает действительные значения. Доказать, что для любого действительного x

$$P\{\xi \geq x\} \leq \inf_{\lambda > 0} e^{-\lambda x} Me^{\lambda \xi}$$

$$P\{\xi \leq x\} \leq \inf_{\lambda > 0} e^{\lambda x} Me^{-\lambda \xi}$$

Ч-4.88

Случайная величина ξ имеет нормальное распределение с параметрами $(0, \sigma^2)$. Установив тождество

$$\frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{hx^2} e^{-x^2/2\sigma^2} dx = \frac{1}{\sqrt{1-2h\sigma^2}}, \quad h < \frac{1}{2\sigma^2},$$

$$\eta = \xi e^{h\xi^2}.$$

Ч-4.89

Функция $f(t)$, $-\infty < t < \infty$, обладает следующими свойствами:

- а) $f(0) = 1, f(t) \geq 0$ для любого t ; 6) $f(t) = f(-t)$ для любого t ;
- в) графиком $f(t)$, $0 \leq t < \infty$, является выпуклая вниз ломаная линия, состоящая из конечного числа звеньев. Является ли $f(t)$ характеристической функцией распределения вероятностей?

Ч-4.90

Функция $f(t)$, $-\infty < t < \infty$, обладает следующими свойствами: $f(0) = 1, f(t) \geq 0, f(t) = f(-t)$ для любого t , $f(t)$ непрерывна и $f''(t) \geq 0$ для любого $t \neq 0$. Является ли $f(t)$ характеристической функцией распределения вероятностей?

Ч-4.91

Пусть $a > 0$ и функция $f_1(t)$ имеет период $4a$,

$$f_1(t) = 1 - \frac{|t|}{a}, \quad |t| \leq 2a.$$

- а) Является ли $f_1(t)$ характеристической функцией?
- б) Является ли характеристической функцией функция $f_2(t) = |f_1(t)|$?

Ч-4.92

Случайные величины $\xi_1, \dots, \xi_k$ ($k > 1$) независимы и имеют функцию распределения $F(x)$, случайные величины $\eta_1, \dots, \eta_k$ независимы и имеют функцию распределения $G(x)$. Следует ли из условия

$$P\{\xi_1 + \dots + \xi_k \leq x\} = P\{\eta_1 + \dots + \eta_k \leq x\}$$

для каждого x совпадение функций распределения $F(x)$ и $G(x)$?

Ч-4.96

*. Характеристическая функция $f(t) = Me^{it\xi}$ удовлетворяет условию $|f''(0)| < \infty$. Показать, что

$$M\xi^2 = -f''(0).$$

Ч-4.97

Показать, что если характеристическая функция $f(t) = Me^{it?}$ дважды дифференцируема при $t = 0$ и $|f''(0)| < \infty$, то $|f'(t)| < \sqrt{|f''(0)|}$ при любом t .

Ч-4.9

8. Являются ли характеристическими функциями вероятностных распределений следующие функции:

- а) $e^{2(e^{it}-1)}$; б) $\cos(t^2)$; в) $\frac{\sin t}{t}(t \neq 0), 1 (t = 0)$; г) $\cos^2 t$; д) $e^{-2t}(t \geq 0), \frac{1}{1+t^2} (t \leq 0)$; е) $\cos(|t|^{2/3})$; ж) $e^{-|t|}$?

Ч-4.99

Найти математическое ожидание и дисперсию случайной величины ξ , характеристическая функция $f(t) = Me^{it\xi}$ которой равна:

- а) $\frac{1}{at} \sin at (a \neq 0)$, б) $\frac{4}{t^2} \cos t \sin^2 \frac{t}{2}$, в) $\frac{4}{t^2} e^{it} \sin^2 \frac{t}{2}$, г) $\sqrt{2} e^{-a|t|/2} \sin\left(\frac{\pi}{4} + \frac{a|t|}{\sqrt{2}}\right) (a > 0)$,

Ч-4.10

1. Исходя из формулы для плотности гамма-распределения (см. введения к гл. 3 и 4), проверить, что если случайные величины ξ_1 и ξ_2 независимы, ξ_1 имеет гамма-распределение с параметром α , а ξ_2 - гамма-распределение с параметром β , то $\xi_1 + \xi_2$ имеет гамма-распределение с параметром $\alpha + \beta$. Используя это свойство семейства гамма-распределений, найти характеристическую функцию $f_a(t)$ и формулы для моментов $m_a^{(k)}(k = 1, 2, \dots)$ k -го порядка гамма-распределения с параметром $\alpha > 0$.

Ч-4.10

2. Случайный вектор $(\xi_1, \dots, \xi_k)$ имеет многомерное нормальное распределение в R^k с вектором математических ожиданий $a = (a_1, \dots, a_k)$ и матрицей ковариаций B . Найти распределение случайной величины

$$\zeta = c_1 \xi_1 + \dots + c_k \xi_k,$$

где $c_1, \dots, c_k$ - действительные числа.

Ч-4.10

3. Случайный вектор $(\xi_1, \dots, \xi_m)$ имеет многомерное нормальное распределение в R^k с вектором математических ожиданий $a = (a_1, \dots, a_m)$ и матрицей ковариаций B ; матрица $C = c_{ij}$ ($i = 1, \dots, m; j = 1, \dots, k$) состоит из действительных чисел. Найти распределение случайного вектора $(\xi_1, \dots, \xi_m)$, где

$$\xi_i = c_{i1} \xi_1 + \dots + c_{ik} \xi_k, \quad i = 1, \dots, m.$$

Ч-4.104

Случайный вектор $\xi = (\xi_1, \dots, \xi_k)$ имеет многомерное нормальное распределение в R^k с вектором математических ожиданий $a = (a_1, \dots, a_n)$ и невырожденной матрицей ковариаций B . Показать, что существует такое ортогональное преобразование C пространства R^k в себя, что вектор $\eta = (\eta_1, \dots, \eta_k) = C\xi$ имеет многомерное нормальное распределение с независимыми компонентами.

МФТИ-Т.9

Найти характеристическую функцию распределения Лапласа, которое определяется плотностью

$$f_\xi(x) = \frac{1}{2} e^{-|x|}$$

$$\phi_\xi(t) = \int_{-\infty}^{+\infty} \exp(itx) f_\xi(x) dx = \frac{1}{2} \int_{-\infty}^{+\infty} \exp(itx - |x|) dx$$

$x < 0$

$$\int_{-\infty}^0 \exp(itx - |x|)d(x) = \int_{-\infty}^0 \exp(itx + x)d(x) = \frac{i}{i-t}$$

$$x \geq 0$$

$$\int_0^{+\infty} \exp(itx - |x|)d(x) = \int_0^{+\infty} \exp(itx - x)d(x) = \frac{i}{i+t}$$

$$\phi_{\xi}(t) = \frac{-1}{(-1-t^2)} = \frac{1}{1+t^2}$$

МФТИ-Т.10

Найти характеристическую функцию нормального распределения с параметрами (a, σ^2)

$$\begin{aligned} \phi_{\xi}(t) &= \int_{-\infty}^{+\infty} \exp(itx) \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-a)^2}{2\sigma^2}\right) d(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{+\infty} \exp\left(itx - \frac{(x-a)^2}{2\sigma^2}\right) d(x) = \\ &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(ita + \left(\frac{it\sqrt{2}\sigma}{2}\right)^2\right) \int_{-\infty}^{+\infty} \exp\left(-\frac{(x-a)^2}{2\sigma^2} + 2\frac{it\sqrt{2}\sigma}{2} \frac{(x-a)}{\sqrt{2}\sigma} - \left(\frac{it\sqrt{2}\sigma}{2}\right)^2\right) d(x) = \\ &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(ita + \left(\frac{it\sqrt{2}\sigma}{2}\right)^2\right) \int_{-\infty}^{+\infty} \exp\left(-\left(\frac{x-a}{\sqrt{2}\sigma} - \frac{it\sqrt{2}\sigma}{2}\right)^2\right) d(x) = \\ &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(ita + \left(\frac{it\sqrt{2}\sigma}{2}\right)^2\right) \int_{-\infty}^{+\infty} \exp\left(-\left(\frac{x-a-it\sigma^2}{\sqrt{2}\sigma}\right)^2\right) d(x) = \frac{1}{\sqrt{2\sigma^2}} \exp\left(ita - \frac{t^2\sigma^2}{2}\right) \end{aligned}$$

МФТИ-Т.11

Найти распределения, которым соответствуют следующие характеристические функции:

$$h_1(t) = \cos t; \quad h_2(t) = \frac{1}{2} + \frac{\cos t}{2} + i \frac{\sin t}{6}; \quad h_3(t) = \frac{1}{2 - e^{-it}}$$

для первой:

$$h_1(t) = \cos(t) = \frac{\exp(it) + \exp(-it)}{2} \rightarrow f_1(x) = \frac{1}{2}I(x=1) + \frac{1}{2}I(x=-1)$$

Для второй:

$$h_2(t) = \frac{1}{2} + \frac{\cos(t)}{2} + i \frac{\sin(t)}{6} = \frac{1}{2} \exp(it \cdot 0) + \frac{\exp(it) + \exp(-it)}{4} + i \frac{\exp(it) - \exp(-it)}{12}$$

поэтому

$$f_2(x) = \frac{1}{2}\delta(x) + \frac{1}{4}\delta(x-1) + \frac{1}{4}\delta(x+1) + i \frac{1}{12}\delta(x-1) - i \frac{1}{12}\delta(x+1)$$

для третьей:

$$h_3(t) = \frac{1}{2 - \exp(-it)} = \frac{\frac{1}{2}}{1 - \frac{\exp(-it)}{2}} = \sum_{n=0}^{+\infty} \left(\frac{\exp(itn)}{2^{n+1}} \right)$$

$$f_3(x) = \sum_{n=0}^{+\infty} \left(\frac{1}{2^{n+1}} I(x=n) \right)$$

МФТИ-Т.12 распределение, у которого хар функция $h(t) = e^{-t^2} \cos t$

Найти распределение, которому соответствует характеристическая функция $h(t) = e^{-t^2} \cos t$
(!?!? просто посчитать интеграл?? мб потом упрощу вычисления??!)

$$\begin{aligned}
 f(x) &= \int_{-\infty}^{+\infty} \exp(-t^2) \cos(t) \exp(-itx) d(t) = \\
 &= \int_{-\infty}^{+\infty} \frac{\exp(-t^2 - itx + it) + \exp(-t^2 - itx - it)}{2} d(t) = \\
 &= \int_{-\infty}^{+\infty} \frac{\exp(-t^2 - it(x-1)) + \exp(-t^2 - it(x+1))}{2} d(t) = \\
 &= \int_{-\infty}^{+\infty} \frac{\exp\left(-t^2 - 2t\frac{x-1}{2}i - \left(\frac{x-1}{2}i\right)^2 + \left(\frac{x-1}{2}i\right)^2\right) + \exp\left(-t^2 - 2t\frac{x+1}{2}i - \left(\frac{x+1}{2}i\right)^2 + \left(\frac{x+1}{2}i\right)^2\right)}{2} d(t) = \\
 &= \int_{-\infty}^{+\infty} \frac{\exp\left(-\left(t + \frac{x-1}{2}i\right)^2 + \left(\frac{x-1}{2}i\right)^2\right) + \exp\left(-\left(t + \frac{x+1}{2}i\right)^2 + \left(\frac{x+1}{2}i\right)^2\right)}{2} d(t) = \\
 &= \int_{-\infty}^{+\infty} \frac{\exp\left(-\left(t + \frac{x-1}{2}i\right)^2\right) \exp\left(\left(\frac{x-1}{2}i\right)^2\right) + \exp\left(-\left(t + \frac{x+1}{2}i\right)^2\right) \exp\left(\left(\frac{x+1}{2}i\right)^2\right)}{2} d(t) = \\
 &= \sqrt{\pi} \frac{\exp\left(-\left(\frac{x-1}{2}\right)^2\right) + \exp\left(-\left(\frac{x+1}{2}\right)^2\right)}{2}
 \end{aligned}$$

Вот такому распределению.

МФТИ-Т.13 Характеристическая ли $h(t) = \cos t^2$

Является ли функция $h(t) = \cos t^2$ характеристической?

$\cos(t^2)$ не является равномерно непрерывной на $\mathbb{R}$, а следовательно, и не может быть характеристической функцией.

Докажем это.

Нужно доказать, что

$$\exists \epsilon > 0 : \forall \delta > 0 \exists t_1, t_2 |t_2 - t_1| < \delta \Rightarrow |\cos(t_2^2) - \cos(t_1^2)| \geq \epsilon$$

Найдём решения уравнения $\cos(t^2) = 1$:

$$t^2 = 2\pi n, n \in \mathbb{N} \cup \{0\}$$

Найдём решения уравнения $\cos(t^2) = -1$:

$$t^2 = 2\pi n - \pi, n \in \mathbb{N}$$

Пусть $t_1 = \sqrt{2\pi n}, t_2 = \sqrt{2\pi n - \pi}$. МФТИ-Тогда $|\cos(t_2^2) - \cos(t_1^2)| = 2 = \epsilon$

$$|t_2 - t_1| = |\sqrt{2\pi n} - \sqrt{2\pi n - \pi}| = \frac{\pi}{\sqrt{2\pi n} + \sqrt{2\pi n - \pi}} < \sqrt{\frac{\pi}{n}} \rightarrow n > \frac{\pi}{\delta^2} \rightarrow$$

$$\exists \epsilon = 2 : \forall \delta > 0 \exists t_1 = \sqrt{2\pi n}, t_2 = \sqrt{2\pi n - \pi}, n > \frac{\pi}{\delta^2}, |t_2 - t_1| < \delta \Rightarrow |\cos(t_2^2) - \cos(t_1^2)| \geq \epsilon$$

Доказали.

МФТИ-Т.14 Предел среднего от хар функции (!??!)

Случайная величина ξ_λ распределена по закону Пуассона с параметром λ . Найти

$$\lim_{\lambda \rightarrow \infty} \mathbf{P} \left(\frac{\xi_\lambda - \lambda}{\sqrt{\lambda}} \leq x \right)$$

Рассмотрим, к чему стремится характеристическая функция случайной величины $\eta = \frac{\xi_\lambda - \lambda}{\sqrt{\lambda}}$.

$$\begin{aligned} \phi_\eta(t) &= \mathbb{E} \left(\exp \left(it \frac{\xi_\lambda - \lambda}{\sqrt{\lambda}} \right) \right) = \exp(-it\sqrt{\lambda}) \mathbb{E} \left(\exp \left(it \frac{\xi_\lambda}{\sqrt{\lambda}} \right) \right) = \\ &= \exp(-it\sqrt{\lambda}) \exp \left(\lambda \left(\exp \left(\frac{it}{\sqrt{\lambda}} \right) - 1 \right) \right) = \exp \left(-it\sqrt{\lambda} + \lambda \left(\exp \left(\frac{it}{\sqrt{\lambda}} \right) - 1 \right) \right) \end{aligned}$$

И предел

$$\begin{aligned} \lim_{\lambda \rightarrow +\infty} \phi_\eta(t) &= \lim_{\lambda \rightarrow +\infty} \exp \left(-it\sqrt{\lambda} + \lambda \left(\exp \left(\frac{it}{\sqrt{\lambda}} \right) - 1 \right) \right) = \\ &= \lim_{\lambda \rightarrow +\infty} \exp \left(-it\sqrt{\lambda} + \lambda \left(1 + \frac{it}{\sqrt{\lambda}} + \frac{1}{2} \left(\frac{it}{\sqrt{\lambda}} \right)^2 + o \left(\left(\frac{t}{\sqrt{\lambda}} \right)^2 \right) - 1 \right) \right) = \\ &= \lim_{\lambda \rightarrow +\infty} \exp \left(-it\sqrt{\lambda} + \lambda \left(\frac{it}{\sqrt{\lambda}} - \frac{1}{2} \frac{1}{\lambda} t^2 + o \left(\left(\frac{t}{\sqrt{\lambda}} \right)^2 \right) \right) \right) \\ &= \lim_{\lambda \rightarrow +\infty} \exp \left(-\frac{1}{2} t^2 + o \left(\left(\frac{t}{\sqrt{\lambda}} \right)^2 \right) \right) = \exp \left(-\frac{1}{2} t^2 \right) \\ \lim_{\lambda \rightarrow +\infty} \left(P \left(\frac{\xi_\lambda - \lambda}{\sqrt{\lambda}} \leq x \right) \right) &= \int_{-\infty}^x \exp \left(-\frac{1}{2} t^2 \right) d(t) = \Phi(x) \end{aligned}$$

МФТИ-Т.15

Используя результат предыдущей задачи, найти предел

$$\lim_{n \rightarrow \infty} e^{-n} \sum_{k=n}^n \frac{n^k}{k!}$$

Так как

$$\begin{aligned} \Phi(x) &= \lim_{\lambda \rightarrow +\infty} P \left(\frac{\xi_\lambda - \lambda}{\sqrt{\lambda}} \leq x \right) = \lim_{\lambda \rightarrow +\infty} P \left(\xi_\lambda \leq \sqrt{\lambda}x + \lambda \right) = \\ &= \lim_{\lambda \rightarrow +\infty} \left[\exp(-\lambda) \sum_{k=1}^{\lfloor \sqrt{\lambda}x + \lambda \rfloor} \left(\frac{\lambda^k}{(k)!} \right) \right] = \lim_{\lambda \rightarrow +\infty} \left[\exp(-\lambda) \left(\sum_{k=1}^{\lfloor \lambda \rfloor} \frac{\lambda^k}{(k)!} + \sum_{k=\lfloor \lambda \rfloor + 1}^{\lfloor \sqrt{\lambda}x + \lambda \rfloor} \frac{\lambda^k}{(k)!} \right) \right] \end{aligned}$$

Теперь рассмотрим случай $\lambda \in \mathbb{N}$:

$$\begin{aligned} \Phi(x) &= \lim_{n \rightarrow +\infty} \left(e^{-n} \left(\sum_{k=1}^n \left(\frac{n^k}{(k)!} \right) + \sum_{k=n+1}^{\lfloor \sqrt{n}x + n \rfloor} \frac{n^k}{(k)!} \right) \right) \\ \Phi(x) &= \lim_{n \rightarrow +\infty} \left[\exp(-n) \sum_{k=n+1}^{\lfloor \sqrt{n}x + n \rfloor} \left(\frac{n^k}{(k)!} \right) \right] + \lim_{n \rightarrow +\infty} \left[\exp(-n) \sum_{k=1}^n \left(\frac{n^k}{(k)!} \right) \right] \end{aligned}$$

Возьмём теперь предел $x \rightarrow +\infty$:

$$\Phi(+\infty) = 1 = \lim_{n \rightarrow +\infty} \left(\exp(-n) \sum_{k=1}^n \left(\frac{n^k}{(k)!} \right) \right)$$

6.5.4 Problems About Bonferroni Inequalities and Convergence to the Poisson Distribution (???)

МФТИ-Т.16.

Пусть положительные независимые случайные величины $\xi_{m,n}$ $m = 1, 2, \dots, n$ одинаково распределены с плотностью $\alpha_n e^{-\alpha_n x}$, $x > 0$, где $\alpha_n = \lambda n$ и $\lambda > 0$. Найти предельное при $n \rightarrow \infty$ распределение случайной величины $\xi_n = \sum_{m=1}^n \xi_{m,n}$

$$f(x) = \alpha_n \exp(-\alpha_n x) = \lambda n \exp(-\lambda n x)$$

Найдём характеристическую функцию случайной величины $\xi_{m,n}$:

$$\begin{aligned} \phi_{\xi_{m,n}}(t) &= \mathbb{E}(\exp(itx)) = \int_{\mathbb{R}^+} \lambda n \exp(-\lambda n x) \exp(itx) dx = \lambda n \int_{\mathbb{R}^+} \exp(-\lambda n x) \exp(itx) dx = \\ &= \int_{\mathbb{R}^+} \lambda n \exp(-\lambda n x) \exp(itx) dx = - \int_{\mathbb{R}^+} \frac{\partial}{\partial x} (\exp(-\lambda n x)) \exp(itx) dx = \\ &= it \int_{\mathbb{R}^+} \exp(-\lambda n x) \exp(itx) dx - (\exp(-\lambda n x) \exp(itx))|_{x=0}^{x \rightarrow +\infty} = it \int_{\mathbb{R}^+} \exp(-\lambda n x) \exp(itx) dx + 1 \rightarrow \\ &\rightarrow \phi_{\xi_{m,n}}(t) = \int_{\mathbb{R}^+} \lambda n \exp(-\lambda n x) \exp(itx) dx = \frac{\lambda n}{\lambda n - it} = 1 - \frac{it}{it - \lambda n} \end{aligned}$$

Найдём предел характеристических функций:

$$\begin{aligned} \phi(t) &= \lim_{n \rightarrow +\infty} \left(\phi_{\sum_{m=1}^n \xi_{m,n}}(t) \right) = \lim_{n \rightarrow +\infty} \left(\left(1 - \frac{1}{1 - \frac{\lambda n}{it}} \right)^n \right) = \lim_{n \rightarrow +\infty} \left(\exp \left(n \ln \left(1 - \frac{1}{1 - \frac{\lambda n}{it}} \right) \right) \right) \\ &= \lim_{n \rightarrow +\infty} \left(\exp \left(-\frac{1}{\frac{1}{n} - \frac{\lambda}{it}} \right) \right) = \exp \left(\frac{it}{\lambda} \right) \\ f(x) &= I \left(x = \frac{1}{\lambda} \right) \end{aligned}$$

6.5.4 Problems About Bonferroni Inequalities and Convergence to the Poisson Distribution (???)

(???)

Ч-4.105

Случайные величины $\xi_1^{(n)}, \xi_2^{(n)}, \dots, \xi_n^{(n)}$ независимы, $P\{\xi_k^{(n)} = 1\} = 1 - P\{\xi_k^{(n)} = 0\} = p_k(n)$, $k = 1, 2, \dots, n$. Пусть при $n \rightarrow \infty$ $\max_{1 \leq k \leq n} p_k(n) \rightarrow 0$, $p_1(n) + \dots + p_n(n) \rightarrow \lambda$, $0 < \lambda < \infty$. Используя метод производящих функций, доказать, что

$$\lim_{n \rightarrow \infty} P\{\xi_1^{(n)} + \dots + \xi_n^{(n)} = m\} = \frac{\lambda^m}{m!} e^{-\lambda}, \quad m = 0, 1, 2, \dots$$

ФР Вшэ

1. Стрельба по цели ведется до тех пор, пока суммарное число попаданий не достигнет N . Найдите математическое ожидание, дисперсию, а также функцию распределения вероятностей потребовавшегося числа выстрелов k , если выстрелы производятся независимо друг от друга и вероятность попадания в отдельной попытке равна p .

(!! тут вообще жесткие задачи, про них почти не думал, потом подумаю.)

Число выстрелов, требуемое чтобы получить N попаданий можно представить в виде суммы $k = k_1 + k_2 + \dots + k_N$, где k_n - это число выстрелов которое потребовалось для достижения n -го попадания, если отсчитывать от предыдущего попадания (или от начала стрельбы, если речь идет об $n = 1$). В силу отсутствия памяти о процесса Бернулли слагаемые статистически независимы. Каждое из них при этом

6.5.4 Problems About Bonferroni Inequalities and Convergence to the Poisson Distribution (???)

имеет геометрическую функцию распределения вероятностей, заданную формулой (4). Воспользовавшись результатами, изложенными на третьей лекции и уравнением (5) и (7), запишем для случая произвольной вероятности p успеха в отдельном испытании

$$\mu_k = N\mu_{k_1} = \frac{N}{p},$$
$$\sigma_k^2 = N\sigma_{k_1}^2 = N\frac{1-p}{p^2}.$$

Чтобы получить функцию распределения вероятностей $p(k)$ нужно умножить вероятность получить ровно $N-1$ попаданий в серии из первых $k-1$ выстрелов на вероятность попадания в k -м исходе, то есть

$$p(k) = pB(N-1, k-1, p) = \frac{(k-1)!}{(N-1)!(k-N)!} p^N (1-p)^{k-N}$$

Этот результат известен как распределение Паскаля.

ФР Вшэ

2. Найти среднее число бросков симметричной монетки необходимое чтобы каждая ее сторона выпала хотя бы два раза.

Подбросим монетку три раза. С вероятностью $1/4$ все три исхода будут одинаковы, и тогда среднее оставшееся число попыток, требуемое для накопления двух недостающих исходов, равно $2+2=4$. С вероятностью $3/4$ не все исходы первых трех попыток одинаковы, и тогда, чтобы получить недостающий исход второй раз, в среднем придется подбросить монетку еще 2 раза. В итоге получаем, что среднее число бросков симметричной монетки необходимое чтобы каждая ее сторона выпала хотя бы два раза равно

$$\frac{1}{4}(3+4) + \frac{3}{4}(3+2) = \frac{22}{4} = 5.5$$

ФР Вшэ

3. Симметричный игральный кубик подбрасывают раз за разом. Найти среднее число попыток $\langle k \rangle$, необходимое, чтобы каждая из шести граней выпала хотя бы один раз. Найти также дисперсию требуемого числа попыток.

Обозначим через $k(p)$ число попыток, требуемых, чтобы в серии испытаний Бернулли впервые выпала единица, если вероятность получить единицу в отдельном испытании равна p . На каждом шаге случайного эксперимента с кубиком пространство элементарных исходов можно разделить на две группы: те исходы, которые уже выпадали, и те, что еще не выпадали. Каждый раз, когда в очередном исходе будет выпадать еще не наблюдавшаяся грань кубика, будем добавлять ее к списку уже произошедших исходов, и, соответственно, уменьшать на единицу количество исходов в группе еще ни разу не выпадавших. Тогда, если в текущий момент первая группа состоит из m исходов (неважно каких именно), то в следующем броске с вероятностью $(6-m)/6$ выпадет еще не наблюдавшийся исход, а с вероятностью $m/6$ повторится один из уже нам встречавшихся. Из сказанного становится ясным, что случайную величину k , математическим ожиданием которой мы интересуемся, можно представить в виде суммы

$$k = 1 + k_1\left(\frac{5}{6}\right) + k_1\left(\frac{2}{3}\right) + k_1\left(\frac{1}{2}\right) + k_1\left(\frac{1}{3}\right) + k_1\left(\frac{1}{6}\right)$$

где $k_1(p)$ - это время появления первой единицы в серии испытаний Бернулли с вероятностью получить единицу в отдельном исходе равной p . Усредняя, получаем

$$\langle k \rangle = 1 + \left\langle k_1\left(\frac{5}{6}\right) \right\rangle + \left\langle k_1\left(\frac{2}{3}\right) \right\rangle + \left\langle k_1\left(\frac{1}{2}\right) \right\rangle + \left\langle k_1\left(\frac{1}{3}\right) \right\rangle + \left\langle k_1\left(\frac{1}{6}\right) \right\rangle = 14,7$$

где мы использовали Ур. (5).

ФР Вшэ

4. Сколько в среднем раз необходимо подбросить симметричную монетку, чтобы впервые встретить в последовательности исходов реализацию вида "орел, решка"?

Для соответствующей случайной величины справедливо

$$k = k_1^o + k_1^p,$$

а значит

$$\langle k \rangle = \langle k_1^o \rangle + \langle k_1^p \rangle = \frac{1}{p} + \frac{1}{1-p} = \frac{1}{p(1-p)}.$$

В частности, для симметричной монетки находим, что для появления пары "орел, решка" необходимо в среднем подкинуть монетку $\langle k \rangle = 4$ раз. Отметим, что среднее число попыток, необходимое для появления пары "орел,орел" равно 6.

ФР Вшэ

5. Сколько в среднем раз необходимо подбросить симметричную монетку, чтобы впервые встретить в последовательности исходов реализацию вида "десять орлов подряд"?

Для ответа на поставленный вопрос, рассмотрим схему Бернулли, где единица соответствует выпадению орла, а ноль - выпадению решки. Будем решать задачу в общем случае, когда вероятность выпадения единицы в отдельном испытании равна p , и речь идет об ожидании цепочки исходов из N единиц подряд. Количество попыток, требуемое для реализации первой такой цепочки обозначим через k . Каждая последовательность независимых повторений эксперимента Бернулли характеризуется некоторым количеством нулей, выпавших до первого появления единицы, и количеством единиц, выпавших подряд сразу вслед за этим. Если число идущих подряд единиц в первой встретившейся цепочке единиц меньше чем N , то можно считать, что процесс ожидания цепочки из N единиц подряд начинается заново в том смысле, что оставшееся число попыток k' необходимое для выпадения N единиц подряд имеет ту же статистику, что и k . Сказанное может быть записано в форме уравнения

$$k = k_1^0 + (N - 1) \cdot I(k_1^P > N - 1) + (k_1^P + k') I(k_1^P \leq N - 1)$$

где нами были введены следующие обозначения: k_1^0 - номер попытки в цепочке независимых испытаний Бернулли, когда впервые выпадает единица, k_1^P - номер попытки в цепочке независимых испытаний Бернулли, когда впервые выпадает ноль, k' - независимая реализация случайной величины k , $I(\dots)$ - индикаторная функция. Рекурсивное соотношение (43) представляет собой пример того, что называется уравнением обновления.

Произведем статистическое усреднение

$$\langle k \rangle = \langle k_1^0 \rangle + (N - 1) \cdot \langle I(k_1^P > N - 1) \rangle + \langle k_1^P I(k_1^P \leq N - 1) \rangle + \langle k' I(k_1^P \leq N - 1) \rangle.$$

Словом, мы приравняли друг другу математические ожидания левой и правой частей Ур. (43). В силу отсутствия памяти у процесса Бернулли, k_1^0 и k' это статистически независимые случайные величины, а значит $\langle k' I(k_1^P \leq N - 1) \rangle = \langle k' \rangle \langle I(k_1^P \leq N - 1) \rangle = \langle k \rangle \langle I(k_1^P \leq N - 1) \rangle$, и мы можем переписать

$$\langle k \rangle = \langle k_1^0 \rangle + (N - 1) \cdot \langle I(k_1^P > N - 1) \rangle + \langle k_1^P I(k_1^P \leq N - 1) \rangle + \langle k \rangle \langle I(k_1^P \leq N - 1) \rangle$$

Решая это уравнение относительно $\langle k \rangle$, находим

$$\langle k \rangle = N - 1 + \frac{\langle k_1^0 \rangle + \langle k_1^P I(k_1^P \leq N - 1) \rangle}{\langle I(k_1^P > N - 1) \rangle}$$

Остается лишь вычислить средние значения, входящие в правую часть полученного выражения.

Вспомним, что статистика случайных величин k_1^0 и k_1^P описывается геометрическими функциями распределения вероятностей, $P_0(k_1^0) = p(1-p)^{k_1^0-1}$ и $P_P(k_1^P) = (1-p)p^{k_1^P-1}$ (см Ур. (4)). Значит

$$\begin{aligned} \langle I(k_1^P > N - 1) \rangle &= \text{Prob}[k_1^P > N - 1] = \sum_{k_1^P=N}^{\infty} (1-p)p^{k_1^P-1} = p^{N-1}, \\ \langle k_1^P I(k_1^P \leq N - 1) \rangle &= \sum_{k_1^P=1}^{N-1} k_1^P (1-p)p^{k_1^P-1} = \frac{1 + (N-1)p^N - N^{N-1}}{1-p}. \end{aligned}$$

Подстановка этих формул в Ур. (46) дает финальный ответ к задаче

$$\langle k \rangle = \frac{1}{1-p} \left(\frac{1}{p^N} - 1 \right).$$

В частности, для симметричной монетки ($p = 1/2$) имеем $\langle k \rangle = 2^{N+1} - 2$. То есть, чтобы встретить десять орлов подряд придется подбросить монетку (в среднем) $2^{11} - 2 = 2046$ раз.

ФР Вшэ

6. Крупные метеориты падают на Землю в пуассоновские моменты времени с рейтом λ . Обозначим через t_n время ожидания падения n -го по счету метеорита, отсчитываемое от произвольно выбранного момента начала наблюдений. Какова функция плотности вероятности случайной величины t_n ?
(черновик от другого решателя)

$$\begin{aligned} P(T > t_n) &= P(X = 0) + P(X = 1) + \dots + P(X = n - 1) = \sum_{k=0}^{n-1} \frac{e^{-\lambda t} : (\lambda t)^k}{k!} \Rightarrow \\ \Rightarrow F_n \neq H &= P(T \leq t_n) = 1 - P(T \leq t) = 1 - \sum_{k=0}^{n-1} \frac{e^{-H} (\lambda t)^k}{k!} \end{aligned}$$

$$f_n(t) = F_n(t) = - \sum_{k=0}^{n-1} \frac{\lambda^k}{k!} \cdot \frac{d}{dt} (e^{-\lambda t} \cdot t^k) =$$

$$- \sum_{k=0}^{n-1} \frac{\lambda^k k}{k!} (-\lambda e^{-\lambda t} \cdot t^k + e^{-\lambda t} \cdot t^{k-1}) = \lambda e^{-\lambda t} \sum_{k=0}^{n-1} \frac{14k}{k!} - e^{-\lambda t} \cdot \sum_{k=1}^{n-1} \frac{(\lambda t)}{(k-1)!} =$$

(тут кривая копия, много не так взял матшникс)

$$= \lambda e^{-\lambda t} \sum_{k=0}^{n-1} \frac{(\Delta t)^k}{k!} - \sum_{k=0}^{n-2} \frac{(a+)^{k-1}}{k!} \Bigg) = \frac{\lambda^n \cdot t^{n-1} \cdot e^{-\lambda t}}{(n-1)!}$$

ФР Вшэ

7. В области полностью ионизированного водорода начинает происходить рекомбинация электронов с ионами, в результате чего газ постепенно деионизируется. Вероятность того, что электрон рекомбинируется за время dt равна $\rho(t)\beta dt$, где $\rho(t)$ это концентрация ионов в текущий момент времени t , а β - некоторый постоянный коэффициент, характеризующий акт рекомбинации. Известно, что концентрация ионов в начальный момент времени равна ρ_0 , а система в целом электронейтральна. Найдите плотность вероятностей для случайной величины T - времени, прошедшего до момента рекомбинации отдельно взятого электрона. Каково медианное значение времени рекомбинации T ?
(решения не давалось)

Ч-4.106

Случайная величина ξ принимает значения 1 и 0 с вероятностями p и $q = 1 - p$ соответственно, а случайная величина η имеет распределение Пуассона с параметром p : $P\{\eta = k\} = \frac{p^k}{k!} e^{-p}, k = 0, 1, 2, \dots$. Доказать, что ξ и η можно задать на одном вероятностном пространстве так, что $P\{\xi \neq \eta\} < p^2$.

Ч-4.107

Случайные величины $\xi_1, \xi_2, \dots, \xi_n$ независимы, $P\{\xi_i = 1\} = 1 - P\{\xi_i = 0\} = p_i, i = 1, \dots, n$. Используя результаты задачи п4.106, доказать, что

$$\frac{1}{2} \sum_{k=0}^{\infty} |P\{\xi_1 + \dots + \xi_n = k\} - \frac{(p_1 + \dots + p_n)^k}{k!} e^{-(p_1 + \dots + p_n)}| \leq \sum_{i=1}^n p_i^2.$$

Ч-4.109

Случайная величина ξ принимает только целые неотрицательные значения. Доказать, что если $m_k = M\xi^{[k]} = M\xi(\xi-1)\dots(\xi-k+1), k = 1, 2, \dots$ и для целого $d \geq 1$ величина $m_{2d-1} < \infty$, то

$$\sum_{k=1}^{2d} (-1)^{k-1} \frac{m_k}{k!} \leq P\{\xi > 0\} \leq \sum_{k=1}^{2d-1} (-1)^{k-1} \frac{m_k}{k!}.$$

Ч-4.110

Пусть выполнены условия задачи п4.109. Доказать, что

$$\sum_{k=n}^{n+2d-1} (-1)^{k-n} C_k^n \frac{m_k}{k!} \leq P\{\xi = n\} \leq \sum_{k=n}^{n+2d} (-1)^{k-n} C_k^n \frac{m_k}{k!}$$

если $n = 0, 1, \dots$ и $m_{n+2d-1} < \infty$.

Ч-4.111

Пусть выполнены условия задачи п4.109. Доказать, что при любом $n = 1, 2, \dots$

$$\sum_{k=n}^{n+2d-1} (-1)^{k-n} C_{k-1}^{n-1} \frac{m_k}{k!} \leq P\{\xi \geq n\} \leq \sum_{k=1}^{n+2d} (-1)^{k-n} C_{k-1}^{n-1} \frac{m_k}{k!}$$

если $m_{n+2d-1} < \infty$.

Ч-4.114

Пассажирский поезд состоит из N вагонов по s мест в каждом вагоне. В момент отправления в поезде находилось n пассажиров. Обозначим символом $\mu_i = \mu_i(n, N, s)$ число вагонов, в которых при отправлении поезда находилось ровно i пассажиров, $i = 0, 1, \dots, \varepsilon$. Предполагая, что все $n!C_{Ns}^n = (Ns)^{[n]}$ вариантов размещения пассажиров в поезде равновероятны, найти формулы для факториальных моментов величин $\mu_0, \mu_1, \dots, \mu_s$. Проанализировать поведение математических ожиданий μ_i при изменении n от 0 до Ns .

Ч-4.115

Пусть выполнены условия задачи Ч-4.114. Найти явное выражение для $P\{\mu_i(n, N, s) = k\}$.

Ч-4.116

Пусть выполнены условия задачи Ч-4.114. Показать, что если s и i фиксированы, а $n, N \rightarrow \infty$ так, что $M\mu_i(n, N, s) \rightarrow \lambda \in (0, \infty)$, то при любом $k = 1, 2, \dots$

$$M(\mu_i(n, N, s))^{[k]} \rightarrow \lambda^k.$$

Вывести отсюда, что тогда для любого $m = 0, 1, \dots$

$$P\{\mu_i(n, N, s) = m\} \rightarrow \frac{\lambda^m}{m!} e^{-\lambda}.$$

Ч-4.117

Пусть n частиц размещаются по N ячейкам, причем каждая частица независимо от остальных с одними и теми же вероятностями (равными $1/N$) может попасть в любую из N ячеек. Обозначим через $\mu_r(n, N)$ число ячеек, в которых находится ровно по r частиц. Найти формулу для $M\mu_r(n, N)$ и доказать, что если $r \Rightarrow 0, 1, 2, \dots$ фиксировано, а $n, N \rightarrow \infty$ так, что $M\mu_r(n, N) \rightarrow \lambda_r$, $0 < \lambda < \infty$, то

$$P\left\{\mu_r(n, N) = \{m\} \rightarrow \frac{\lambda^m}{m!} e^{-\lambda}, m = 0, 1, \dots\right.$$

Ч-4.118

Пусть частицы последовательно и независимо друг от друга размещаются по N ячейкам так, что i -я ($i = 1, 2, \dots$) частица попадает в $j - 1$ ($j = 1, \dots, N$) ячейку с вероятностью $1/N$. Положим

$$v_r(N) = \min\{n\mu_r(n, N) > 0\}, \quad r = 2, 3, \dots$$

137 где $\mu_r(n, N)$ - число ячеек, содержащих после размещения n частиц ровно по r частиц. Найти такую последовательность чисел $b_1, b_2, \dots$ что распределения случайных величин $v_r(N)/b_N$ при $N \rightarrow \infty$ и фиксированном r сходятся к невырожденному закону, и сам предельный закон.

Ч-4.119

Пусть n частиц независимо размещаются по N^2 ячейкам, расположенным в виде квадратной таблицы размером $N \times N$. Назовем ячейку свободной, если после размещения n частиц ни в нее, ни в ячейки, находящиеся с ней в одной строке или в одном столбце, не попало ни одной частицы. Найти предельное распределение числа $\chi(n, N)$ свободных ячеек, когда

$$n = N(\ln \lambda N + o(1)), \quad N \rightarrow \infty.$$

6.5.5 Problems About the Central Limit Theorem and the Method of Characteristic Functions (??!!)

(мб это супер нужное??)

(!!!! напишу в 1 часть указания про решения!!!! тут крутые указания должны быть!!!!)

Вероятность прихода более половины студентов на лекции (!?!?!!?)

Обзорную лекцию должны прослушать 100 студентов. Вероятность присутствовать на этой лекции для каждого студента равняется 0,7. Найти вероятность того, что на лекцию придет больше половины студентов.

(??? задача хорошая, но я слишком плохо разобрался, не уверен, что такое решение годное.)

Пусть X - случайная величина, равная числу студентов, пришедших на лекцию. Она равна сумме $n = 100$ независимых одинаково распределённых случайных величин, каждая из которых принимает значение 1 с вероятностью $p = 0,7$ (студент пришёл) и значение 0 с вероятностью $q = 1 - p = 0,3$ (студент не пришёл).

Используем центральную предельную теорему, согласно которой, при больших значениях n случайную величину $Y = (X - MX)/\sqrt{DX}$ можно приближённо считать нормально распределённой с параметрами 0 и 1. Согласно стандартным формулам, $MX = np = 70$, $DX = npq = 21$. Нас интересует вероятность события $X > 50$, что равносильно условию $X - MX > -20$, то есть $Y > -20/\sqrt{21}$. Вероятность этого события находим, интегрируя функцию плотности нормального распределения:

$$P\{Y > -20/\sqrt{21}\} = \frac{1}{\sqrt{2\pi}} \int_{-20/\sqrt{21}}^{\infty} e^{-x^2/2} dx \approx 0,9999936245$$

то есть эта вероятность практически равна единице.

Белан Вшэ (!!!!)

3. Пусть $X_1, X_2, \dots, X_N$ - набор независимых одинаково распределённых неотрицательных случайных величин. Найдите функцию плотности вероятности $P_N(g)$ среднего геометрического $G = \sqrt[N]{\prod_{i=1}^N X_i}$ в случае $N \gg 1$.

(нужно тут прокачаться ещё!!)

Обозначим через $\rho_N(y)$ функцию плотности вероятности случайной величины $Y = \ln G = \frac{1}{N} \sum_{i=1}^N \ln X_i$.

Понятно, что $P_N(G) = \rho_N(Y(G)) \left| \frac{dY(G)}{dG} \right|$, а значит, если мы вычислим $\rho_N(y)$, то и искомую функцию распределения $P_N(g)$ можно будет легко выписать. Заметим, что Y представляет собой среднее арифметическое суммы независимых и одинаково распределённых случайных чисел $Y_1, Y_2, \dots, Y_N$, где $Y_i = \ln X_i$. При этом Y обладает конечным средним значением $\mu_Y = \langle Y \rangle = \langle \ln X \rangle$ и дисперсией $\sigma_Y^2 = \langle Y^2 \rangle - \langle Y \rangle^2 = \langle \ln^2 X \rangle - \langle \ln X \rangle^2$ (данное утверждение справедливо независимо от того, конечны ли $\langle X \rangle$ и $\langle X^2 \rangle$). Таким образом, выполнены условия применимости центральной предельной теоремы, и при $N \gg 1$ величина Y имеет нормальное распределение

$$\rho_N(y) = \mathcal{N}\left(\mu_Y, \frac{\sigma_Y^2}{N}\right) = \frac{\sqrt{N}}{\sqrt{2\pi}\sigma_Y} \exp\left(-\frac{N(y - \mu_Y)^2}{2\sigma_Y^2}\right).$$

Тогда

$$P_N(g) = \rho_N(y(g)) \left| \frac{dy(g)}{dg} \right| = \frac{\sqrt{N}}{\sqrt{2\pi}\sigma_Y g} \exp\left(-\frac{N(\ln g - \mu_Y)^2}{2\sigma_Y^2}\right).$$

В пределе $N \rightarrow \infty$ получаем $P_\infty(g) = \delta(g - e^{\langle \ln X \rangle})$.

Ч-4.120

(см. п4.2).

Случайная величина η_n равна сумме очков, выпавших игри n независимых бросаниях симметричной игральной кости. Используя центральную предельную теорему, выбрать n так, чтобы

$$P\left\{\left|\frac{\mu_n}{n} - 3,5\right| \geq 0,1\right\} \leq 0,1.$$

Ч-4.121

Складывается 10^4 чисел, округленных с точностью до 10^{-m} . Предполагая, что ошибки округления независимы и равномерно распределены в интервале $(-0,5 \cdot 10^{-m}, 0,5 \cdot 10^{-m})$, пайти пределы, в которых с вероятностью, не меньшей 0,99, будет лежать суммарная ошибка.

Ч-4.122

Случайные величины $\xi_1, \xi_2, \dots$ независимы и одинаково распределены:

$$P\{\xi_i = z\} = P\left\{\xi_i = \frac{1}{2}\right\} = \frac{1}{2}, \quad z \neq 1, \quad i = 1, 2, \dots$$

а) Найти пределы математического ожидания и дисперсии $n \rightarrow \infty$.

б) Доказать, что при $n \rightarrow \infty$ распределение η_n сходится к логарифмически нормальному распределению (см. задачу 3.233). Найти параметры a, σ^2 предельного логарифмически нормального распределения.

в) Сравнить найденные в п. а) значения $\lim_{n \rightarrow \infty} M\eta_n$ и $\lim_{n \rightarrow \infty} D\eta_n$ с математическим ожиданием и дисперсией случайной величины η , имеющей логарифмически нормальное распределение из п. б).

Ч-4.123

*. Случайные величины $\xi_1, \xi_2, \dots$ независимы и одинаково распределены:

$$P\{\xi_i = 1, 25\} = P\{\xi_i = 0, 8\} = \frac{1}{2}, \quad i = 1, 2, \dots$$

и $\eta_n = \xi_1 \dots \xi_n$.

а) Найти $M\eta_{1000}, D\eta_{1000}, M \ln \eta_{1000}, D \ln \eta_{1000}$.

б) Пользуясь асимптотической нормальностью $\ln \eta_n$ при $n \rightarrow \infty$, найти приближенные значения $P\{\eta_{1000} \leq 0,001\}, P\{\eta_{1000} < 1\}, P\{\eta_{1000} \leq 1\}, P\{\eta_{1000} \leq 10^4\}$.

в) Пользуясь формулой Стирлинга, найти $P\{\eta_{1000}^4 < 1\}, P\{\eta_{1000} \leq 1\}$.

Ч-4.124

Случайные величины $\xi_1, \xi_2, \dots$ независимы и одинаково распределены:

$$P\{\xi_i = 1, 25\} = P\{\xi_i = 0, 75\} = \frac{1}{2}, \quad i = 1, 2, \dots$$

и $\eta_n = \xi_1 \dots \xi_n$.

а) Найти $M\eta_{1000}, D\eta_{1000}, M \ln \eta_{1000}, D \ln \eta_{1000}$.

б) Пользуясь асимптотической нормальностью $\ln \eta_n$ при $n \rightarrow \infty$, найти приближенные значения

$$\begin{aligned} &P\{\eta_{1000} \leq 10^{-20}\}, \\ &P\{\eta_{1000} < 1, 25^{501} \cdot 0, 75^{499} \approx 1, 6 \cdot 10^{-14}\} \\ &P\{\eta_{1000} \leq 1, 25^{501} \cdot 0, 75^{499}\}, \quad P\{\eta_{1000} \leq 10^{-7}\}. \end{aligned}$$

в) Пользуясь формулой Стирлинга, найти

$$\begin{aligned} &P\{\eta_{1000} < 1, 25^{501} \cdot 0, 75^{499}\} \\ &P\{\eta_{1000} \leq 1, 25^{501} \cdot 0, 75^{499}\} \end{aligned}$$

Ч-4.126

Случайная точка $\xi = (\xi_1, \dots, \xi_n) \in R^n$ имеет равномерное распределение на единичной сфере в R^n , т. е. на множестве точек $S^{n-1} \Rightarrow \{(x_1, \dots, x_n) \in R^n$:

Ч-4.127

Случайный вектор $\xi = (\xi_1, \dots, \xi_n) \in R^n$ имеет сферически симметричное нормальное распределение с нулевым вектором средних и единичной матрицей ковариаций. Равенство

$$\xi = \rho \varepsilon, \quad \rho = |\xi|, \quad \varepsilon = \frac{1}{|\xi|} \xi,$$

где $|\xi| = \sqrt{\xi_1^2 + \dots + \xi_n^2}$ - длина вектора ξ , дает представление ξ в виде произведения скалярной случайной величины ρ и вектора ε . Доказать, что ρ и ε независимы и что ε имеет равномерное распределение на единичной сфере в R^n (см. задачу

Ч-4.126

). Найти распределение, математическое ожидание и дисперсию ρ^2 .

Ч-4.128

Пусть вектор $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n) \in R^n$ имеет равномерное распределение на единичной сфере в R^n . Используя результаты задач 4.127 и 4.125, доказать, что при $n \rightarrow \infty$:

- распределение $\varepsilon_1 \sqrt{n}$ сходится к стандартному нормальному распределению,
- распределение вектора $(\varepsilon_1 \sqrt{n}, \dots, \varepsilon_k \sqrt{n}), k = \text{const}$, сходится к совместному распределению k независимых случайных величин, имеющих стандартное нормальное распределение.

Ч-4.129

Случайные величины $\xi_0, \xi_1, \xi_2, \dots$ независимы и имеют стандартное нормальное распределение. Распределение случайной величины

$$\tau_n = \frac{\xi_0}{\sqrt{(\xi_1^2 + \dots + \xi_n^2)/n}}$$

называется распределением Стьюдента с n степенями свободы. Найти

$$\lim_{n \rightarrow \infty} P\{\tau_n \leq x\} = \Phi(x) \quad -\infty < x < \infty.$$

Ч-4.130

Пусть $\xi_{n,1}$ обозначает число появлений t -го исхода в n независимых испытаниях с N несовместными исходами; вероятности появления этих исходов в каждом из испытаний равны $p_1, p_2, \dots, p_N \geq 0$ соответственно, $p_1 + \dots + p_N = 1$. Далее, пусть $a_1, \dots, a_N$ - действительные числа и

$$\eta_n = a_1 \xi_{n,1} + \dots + a_N \xi_{n,N}.$$

Найти $M\eta_n, D\eta_n$ и предельное распределение $\frac{\eta_n - M\eta_n}{\sqrt{D\eta_n}}$ при $n \rightarrow \infty$.

Ч-4.131

Случайные величины $\xi_1, \xi_2, \dots$ независимы и одинаково распределены; $\xi_a = a, D\xi_a \Rightarrow b^2, k = 1, 2, \dots$. Положим $\eta_k = \xi_n + \xi_{k+1} + \xi_{k+2}, k = 1, 2, \dots$. Найти:

- $M\eta_k, \text{cov}(\eta_k, \eta_c) (k \neq l), D\eta_k$; 6) $\lim_{n \rightarrow \infty} P\left\{\frac{\eta_1 + \dots + \eta_n - 3na}{b\sqrt{3n}} \leq x\right\}$. Пуассона с параметром λ . Найти

$$\lim_{\lambda \rightarrow \infty} P\left\{\frac{\xi_\lambda - \lambda}{\sqrt{\lambda}} \leq x\right\}.$$

Ч-4.134

Последовательность случайных величин $\xi_1, \xi_2, \dots$ такова, что при любом $n = 1, 2, \dots$

$$\xi_n = \xi_{n,1} + \xi_{n,2} + \dots + \xi_{n,n}$$

141 где $\xi_{n,1}, \dots, \xi_{n,n}$ - независимые случайные величины,

$$M\xi_{n,1} = 0, \quad i = 1, \dots, n,$$

$$s_n^2 = \sum_{k=1}^n D\xi_{n,k} < \infty, \quad \lim_{n \rightarrow \infty} \frac{1}{s_n^3} \sum_{k=1}^n M|\xi_{n,k}|^3 = 0.$$

Доказать, что распределения случайных величин ξ_n/s_n при $n \rightarrow \infty$ сходятся к стандартному нормальному распределению.

Ч-4.135

Случайные величины $\xi_1, \xi_2, \dots$ независимы,

$$P\{\xi_n = 1\} = \frac{1}{n^\alpha}, \quad P\{\xi_n = 0\} = \frac{n^\alpha - 1}{n^\alpha},$$

$$n = 1, 2, \dots, 0 < \alpha < 1.$$

Найти предельное распределение случайных величин

$$\xi_n = \frac{\xi_1 + \dots + \xi_n - M(\xi_1 + \dots + \xi_n)}{\sqrt{D(\xi_1 + \dots + \xi_n)}} \quad \text{при } n \rightarrow \infty.$$

Ч-4.138

Случайные величины $\xi_1^{(n)}, \dots, \xi_n^{(n)}$ независимы, $\mathbf{P} \{ \xi_i^{(n)} = 1 \} = p_i^{(n)}$

$$P \{ \xi_i^{(n)} = 0 \} = 1 - p_i^{(n)} \quad (1 \leq i \leq n)_S$$

и $\xi_n = \xi_1^{(n)} + \dots + \xi_n^{(n)}$.

а) Найти $M\xi_n, D\xi_n$.

б) Доказать, что если $D_n \rightarrow \infty$ при $n \rightarrow \infty$, то предельное при $n \rightarrow \infty$ распределение случайной величины $\frac{\xi_n - M\xi_n}{\sqrt{D\xi_n}}$ является стандартным нормальным.

в) Положим $m_n = f(p_1^{(n)}) + \dots + f(p_n^{(n)})$, где $f(x) = 0$ при $x \leq 1/2$ и $f(x) = 1$ при $x > 1/2$. Найти предельное при $n \rightarrow \infty$ распределение случайной величины $\xi_n - m_n$, если

$$\sum_{\substack{1 \leq i \leq n \\ p_i^{(n)} \leq 1/2}} p_i^{(n)} \rightarrow \lambda_{0v} \quad \sum_{\substack{1 \leq i \leq n \\ p_i^{(n)} > 1/2}} p_i^{(n)} \rightarrow \lambda_{1\varepsilon}$$

$$\max_{1 \leq i < n} \min [p_i^{(n)}, 1 - p_i^{(n)}] \rightarrow 0.$$

Ч-4.139

Случайная величина ξ_n имеет распределение - дробно-дирежной производящей функцией $\varphi_n(z) \Rightarrow = \frac{a_n + (1-d_n-a_n)z}{1-d_n z^2}$, $0 \leq a_n \leq 1, 0 \leq d_n < 1$. Какие невырожденные законы распределения $F(x)$ могут быть предельными для последовательности $F_n(x) = \mathbf{P} \left\{ \frac{\xi_n}{A_n} \leq x \right\}$, где $A_n \rightarrow \infty$ - некоторые нормирующие константы? Какие условия надо наложить на последовательности a_n, d_n, A_n , чтобы $\lim_{n \rightarrow \infty} F_n(x) = F(x)$, $x > 0$?

Ч-4.144

Пусть случайная величина v_h та же, что в задаче п4.143. Найти предельное распределение случайной величины $v_k - k$, когда $N \rightarrow \infty, k^2/N \rightarrow \lambda, 0 < \lambda < \infty$.

Ч-4.145

Пусть выполнены условия задачи п4.143. Найти асимптотические формулы для Mv_m, Dv_m и предельное распределение $\frac{v_m - v_m}{\sqrt{Dv_m}}$, когда $N \rightarrow \infty$ и $\alpha_1 N \leq m \leq \alpha_2 N$ при некоторых $\alpha_1, \alpha_2, 0 < \alpha_1 < \alpha_2 < 1$.

Ч-4.146

Производящая функция $\varphi(z) = Mz^2$ случайной величины ξ является многочленом степени N , причем все корни $z_1, \dots, z_N$ уравнения $\varphi(z) = 0$ вещественны. Доказать, что распределение ξ совпадает с распределением суммы $\zeta = \zeta_1 + \dots + \zeta_N$, где случайные величины $\zeta_1, \dots, \zeta_N$ независимы и

$$\mathbf{P} \{ \zeta_i = 1 \} = p_i, \quad \mathbf{P} \{ \zeta_i = 0 \} = 1 - p_i, \quad i = 1, \dots, N.$$

Найти значения $p_1, \dots, p_N$.

Ч-4.147

*. Производящая функция $\varphi(z) = Mz^2$ случайной величины ξ является многочленом степени N , причем все корни $z_1, \dots, z_N$ уравнения $\varphi(z) = 0$ имеют неположительные вещественные части: $\operatorname{Re} z_i \leq 0, i = 1, \dots, N$. Доказать, что распределение ξ совпадает с распределением суммы $\xi = \xi_1 + \dots + \xi_M$, где случайные величины $\xi_1, \dots, \xi_M$ независимы и $\mathbf{P} \{ \zeta_i = 2 \} = r_i, \quad \mathbf{P} \{ \zeta_i = 1 \} = p_i, \quad \mathbf{P} \{ \zeta_i = 0 \} = 1 - p_i - r_i, \quad i = 1, \dots, M$. Найти значения $M, p_i, r_i (i = 1, \dots, M)$.

Ч-4.149

Доказать, что если выполнены условия задачи п4.148, но $Dg_n \rightarrow \lambda < \infty$ при $n \rightarrow \infty$ и, кроме того,

$$\lim_{n \rightarrow \infty} \max_{1 \leq i \leq n} z_{n,i} = -\infty,$$

где $z_{n,1}, \dots, z_{n,n}$ - корни уравнения $\varphi_n(z) = 0$, то распределение ξ_n при $n \rightarrow \infty$ сходится к распределению Пуассона с параметром λ .

Ч-4.150

*. Пусть выполнены условия задач п4.143 и $f_{n,N}(z) = z^{\mu_0(n,N)}$ - производящая функция числа пустых ячеек. Найти рекуррентное уравнение, связывающее многочлены $f_{n,N}(z)$ и $f_{n+1,N}(z)$, и вывести из него, что все $N - 1$ корней уравнения $f_{n,N}(z) = 0$ вещественны при любом $n \geq 1$.

Ч-4.151

Пусть выполнены условия задачи п4.143 и $\mu_0(n, N)$ - число ячеек, оставшихся пустыми после размещения n частиц. Доказать, что если $n, N \rightarrow \infty$ так, что $M\mu_0(n, N) \rightarrow \infty$ и $N - M\mu_0(n, N) \rightarrow \infty$, то предельное распределение случайной величины

$$\frac{\mu_0(n, N) - M_0(n, N)}{\sqrt{D\mu_0(n, N)}}$$

является стандартным нормальным.

Ч-4.152

Случайные величины $\xi_1, \xi_2, \dots$ независимы и имеют одно и то же распределение с характеристической функцией $f(t)$. Доказать, что если при некоторых $C > 0, 0 < \alpha \leq 2$,

$$f(t) = 1 - C|t|^\alpha(1 + o(1)), \quad t \rightarrow 0,$$

то при $n \rightarrow \infty$ существует предельное распределение случайных величин

$$\zeta_n = \frac{\xi_1 + \xi_2 + \dots + \xi_n}{n^{1/\alpha}}.$$

Найти его характеристическую функцию.

Ч-4.157

Случайные величины ξ_1, ξ_2, ξ_3 имеют распределение Коши с параметром a ; при этом ξ_1 и ξ_2 независимы, а $P\{\xi_3 = \xi_1\} = 1$. Сравнить характеристические функции и функции распределения суммы $\xi_1 + \xi_2$ независимых и суммы $\xi_1 + \xi_3$ зависимых одинаково распределенных случайных величин.

Ч-4.158

Используя результат задачи п4.155, вычислить

$$I_{a,b}(x) = \int_{-\infty}^{\infty} \frac{du}{(a^2 + u^2)(b^2 + (x - u)^2)}.$$

6.6 Questions to Test Knowledge

6.6.1 General Questions to Test Understanding

(???)

6.6.2 Typical Technical Questions to Test Understanding

(???)

6.6.3 Questions to Test Understanding of Details

(???)

7 Other Problems

7.1 Problems About Elementary Random Processes

7.1.1 Problems About the Simplest Random Processes

(!!! тут могут быть очень крутые задачи.)

ЛШ-9 Возврат назад при случайном блуждании

Рассмотрим частицу, совершающую случайное блуждание по n -мерной кубической решетке. Движение начинается при $t = 0$ из начала координат. Попадая на очередном шаге в какой-то из узлов решетки, частица на следующем шаге с одинаковой вероятностью переходит в любой из $2n$ соседних узлов. Пусть $p(t, \mathbf{x})$ - вероятность того, что частица через t шагов оказалась в узле $\mathbf{x} = (x_1, \dots, x_n)$ решетки. (В этой

задаче время t является дискретным.) В теории вероятностей рассматривают производящую функцию

$$G(z, \mathbf{q}) = \sum_{\mathbf{x}, t} z^t e^{i\mathbf{q}\mathbf{x}} p(t, \mathbf{x}) \quad (t \geq 0, |z| < 1).$$

Свойства этой величины во многом аналогичны свойствам функции Грина. Покажите, что

$$G(z, \mathbf{q}) = \frac{1}{1 - zW(\mathbf{q})}, \quad W(\mathbf{q}) = \frac{1}{n} (\cos q_1 + \dots + \cos q_n).$$

Рассмотрим производящую функцию $\tilde{G}(z, \mathbf{q})$ для блужданий, начинающихся из начала координат, но ни разу не возвращающихся туда на последующих шагах. Величина $\tilde{G}(z, \mathbf{q})$ по своим свойствам похожа на функцию Грина частицы в поле отталкивающего центра (см. задачи 11, 12 и 13). В частности, для нее можно написать такое же интегральное уравнение (напоминающее уравнение Дайсона). Выразите $\tilde{G}(z, \mathbf{q})$ через $G(z, \mathbf{q})$. Найдите вероятность P того, что частица никогда не возвращается в начало координат. Покажите, что

$$P^{-1} = \int G(1, \mathbf{q}) \frac{d^n \mathbf{q}}{(2\pi)^n}.$$

Здесь интеграл берется по зоне Бриллюэна, т. е. по периоду обратной решетки. Вероятность возврата (2.17) имеет нетривиальную зависимость от размерности решетки n . Покажите, что

- а) $P = 0$ при $n \leq 2$;
- б) $0 < P < 1$ при $n > 2$;
- в) $P \rightarrow 1$ при $n \gg 2$.

Поскольку перечисленные свойства чувствительны только к поведению $G(1, \mathbf{q})$ при малых $\mathbf{q}$, т.е. на больших масштабах, они имеют место для произвольного диффузионного движения, а не только для блуждания по решетке. Типичная диффузионная траектория имеет бесконечно много возвратов при $n \leq 2$, и конечное число при $n > 2$.

Solution

(!!! в теории укажу это, пока теория слабее каталога!!!)

Рассмотрим вероятность $p(t, \mathbf{x})$ блуждающей частицы оказаться в узле $\mathbf{x}$ на шаге t . Вероятности на шаге t и $t + 1$ связаны простым соотношением:

$$p(t + 1, \mathbf{x}) = \frac{1}{2n} \sum_{|\mathbf{x}' - \mathbf{x}|=1} p(t, \mathbf{x}')$$

(сумма берется по $2n$ соседям узла $\mathbf{x}$). Перейдем к фурье-образу $p(t, \mathbf{q}) = \sum_{\mathbf{x}} e^{i\mathbf{q}\mathbf{x}} p(t, \mathbf{x})$. Соотношение (2.82) дает

$$p(t + 1, \mathbf{q}) = W(\mathbf{q})p(t, \mathbf{q}), \quad W(\mathbf{q}) = \frac{1}{n} (\cos q_1 + \dots + \cos q_n).$$

Находим производящую функцию

$$G(z, \mathbf{q}) = \sum_{t \geq 0} z^t p(t, \mathbf{q}) = \frac{1}{1 - zW(\mathbf{q})}.$$

Вероятности $p(t, \mathbf{x})$ выражаются через $G(z, \mathbf{q})$ так:

$$p(t, \mathbf{x}) = \frac{-i}{(2\pi)^{n+1}} \oint \frac{dz}{z^{t+1}} \int \frac{d^n q e^{i\mathbf{q}\mathbf{x}}}{1 - zW(\mathbf{q})},$$

где интеграл по $d^n q$ берется по области $-\pi < q_i < \pi (i = 1, \dots, n)$, а интеграл по z - по любому контуру, охватывающему точку $z = 0$.

Теперь, чтобы найти вероятность блужданий без возврата в начало координат, немного модифицируем правила игры. Предположим, что частица блуждает случайно, как и раньше, но, как только она попадает в начало координат, ее «удаляют с поля». В этом случае связь между $p(t+1, \mathbf{x})$ и $p(t, \mathbf{x})$ будет такая:

$$p(t+1, \mathbf{x}) = \begin{cases} \frac{1}{2^n} \sum_{|\mathbf{x}' - \mathbf{x}|=1} p(t, \mathbf{x}') & \text{при } \mathbf{x} \neq 0, \\ 0 & \text{при } \mathbf{x} = 0. \end{cases}$$

Переходя к фурье-образу $p(t, \mathbf{q})$, получаем

$$p(t+1, \mathbf{q}) = W(\mathbf{q})p(t, \mathbf{q}) - \int W(\mathbf{k})p(t, \mathbf{k}) \frac{d^n k}{(2\pi)^n}.$$

Производящая функция $\tilde{G}(z, \mathbf{q}) = \sum_{t \geq 0} z^t p(t, \mathbf{q})$ получается умножением (2.86) на z^{t+1} и суммированием по $t \geq 0$. Находим

$$\tilde{G}(z, \mathbf{q})(1 - zW(\mathbf{q})) = 1 - z \int W(\mathbf{k})\tilde{G}(z, \mathbf{k}) \frac{d^n k}{(2\pi)^n}$$

Это уравнение можно записать в таком виде:

$$\tilde{G}(z, \mathbf{q}) = G(z, \mathbf{q}) + G(z, \mathbf{q}) \int \Sigma(z, \mathbf{k})\tilde{G}(z, \mathbf{k}) \frac{d^n k}{(2\pi)^n},$$

где $\Sigma(z, \mathbf{k}) = G^{-1}(z, \mathbf{k}) - 1$. (Заметим, что по форме уравнение (2.89) напоминает уравнение Дайсона (4.9).) Ищем решение (2.89) в виде

$$\tilde{G}(z, \mathbf{q}) = \lambda(z)G(z, \mathbf{q})$$

где $\lambda(z)$ - некоторая функция z . Подставляя (2.90) в (2.89), находим

$$\lambda(z) = 1 - \lambda(z) \int (G(z, \mathbf{q}) - 1) \frac{d^n q}{(2\pi)^n},$$

откуда

$$\lambda^{-1}(z) = \int G(z, \mathbf{q}) \frac{d^n q}{(2\pi)^n}.$$

Рассмотрим вероятность P_t того, что частица за t шагов ни разу не вернулась в начало координат. Производящая функция $F(z) = \sum_{t \geq 0} z^t P_t$, очевидно, есть

$$F(z) = \tilde{G}(z, \mathbf{q} = 0) = \frac{\lambda(z)}{1 - z}.$$

Находим вероятность

$$P_t = \frac{1}{2\pi i} \oint_{|z|=r} \frac{dz}{z^{t+1}} \frac{\lambda(z)}{1 - z},$$

где $\lambda(z)$ дается (2.92), а радиус контура интегрирования $r < 1$. Для нахождения вероятности того, что частица никогда не вернется в начало координат, необходимо перейти в (2.94) к пределу $t \rightarrow \infty$. Это удобно сделать, рассмотрев

$$P_{a,t} = A_t (P_0 + aP_1 + \dots + a^{t-1}P_{t-1}) = \frac{A_t}{2\pi i} \oint_{|z|=r} \frac{(1 - a^t z^{-t})\lambda(z)}{(1 - z)(z - a)} dz$$

где a - параметр, который $0 < a < r$, а $A_t = (1 + a + \dots + a^{t-1})^{-1} = (1 - a^t)/(1 - a)$. Переходя в (2.95) к пределу $t \rightarrow \infty$, получаем

$$P_a^* = \frac{1 - a}{2\pi i} \oint_{|z|=r} \frac{\lambda(z) dz}{(1 - z)(z - a)},$$

Рассмотрим контурный интеграл (2.96). Функция $\lambda(z)$ аналитична внутри единичного круга, поэтому внутри контура $|z| = r$ подинтегральное выражение в (2.96) имеет один простой полюс $z = a$. Вычет в

точке $z = a$ есть $\lambda(a)$, и поэтому $P_a^* = \lambda(a)$. Искомая вероятность невозвращения, согласно (2.95), есть $\lim_{a \rightarrow 1} P_a^*$. Следовательно,

$$P = \lambda(a)_{a \rightarrow 1} = \left[\int G(1, \mathbf{q}) \frac{d^n q}{(2\pi)^n} \right]^{-1}.$$

При $n \leq 2$ интеграл в (2.97) расходится на малых q , поэтому вероятность невозвращения $P = 0$.

В другом пределе, при $n \rightarrow \infty$, величина суммы косинусов в $W(\mathbf{q})$ порядка $\sqrt{n}$, согласно закону больших чисел. (Это справедливо не для всех $\mathbf{q}$, а только для «типичных», но для оценки интеграла в (2.97) этого достаточно.) Заменяя в $G(1, q)$ величину $1 - W(\mathbf{q})$ на 1, находим, что при больших n вероятность невозвращения $P \rightarrow 1$.

Ч-5.1

Случайные величины $\xi_t (t = 0, \pm 1, \pm 2, \dots)$ независимы, $M\xi_t = a$, $D\xi_t = \sigma^2$. По действительным числам $c_0, c_1, \dots, c_h$, удовлетворяющим условию $c_0 + c_1 + \dots + c_k = 1$, Игостроен случайный процесс

$$\eta_1 = c_0 \xi_t + c_1 \xi_{t-1} + \dots + c_k \xi_{t-k}, \quad t = 0, \pm 1, \pm 2, \dots$$

Найти $M\eta_t, D\eta_t, \text{cov}(\eta_t, \eta_s)$. Показать, что $\text{cov}(\eta_t, \eta_s) = R(|s - t|)$, т. е. зависит только от $|s - t|$.

Ч-5.2

. Пусть выполнены условия задачи н5.1.

а) Доказать, что последовательность случайных величин

$$\zeta_n = \frac{1}{n} (\eta_1 + \dots + \eta_n)$$

при $n \rightarrow \infty$ сходится по вероятности к $a = M \xi_t$.

б) Сходится ли ζ_n к a с вероятностью 1 при $n \rightarrow \infty$?

Ч-5.3

Пусть выполнены условия задачи н5.1 и, кроме того, $a = 0, M\xi_t^4 = c < \infty$.

а) Доказать, что для любого целого $s \geq 0$ последовательность случайных величин

$$\zeta_n^{(s)} = \frac{1}{n} (\eta_1 \eta_{1+s} + \eta_2 \eta_{2+s} + \dots + \eta_n \eta_{n+s})$$

при $n \rightarrow \infty$ сходится по вероятности к $R(s)$.

б) Сходится ли $\zeta_n^{(s)}$ к $R(s)$ с вероятностью 1 при $n \rightarrow \infty$?

Ч-5.4

. Случайный процесс $\xi_t, t = 0, \pm 1, \pm 2, \dots$, определяется формулой

$$\xi_1 = \alpha \sin(At + \beta) + \varepsilon_t,$$

где $\alpha, \beta, \{\varepsilon_t\}_{t=-\infty}^{\infty}$ - независимые случайные величины, $M\alpha = 0, D\alpha = 1, \beta$ имеет равномерное распределение на отрезке $[-\pi, \pi]$, $M\varepsilon_t = 0, D\varepsilon_t = \sigma^2$. Найти $M\xi_t, D\xi_t, \text{cov}(\xi_t, \xi_s)$.

Ч-5.5

Случайный процесс $\xi_1, -\infty < t < \infty$, задан формулой

$$\xi_1 = \sin(t + \pi\eta_1) + \sin \pi(t + \eta_2),$$

где η_1 и η_2 независимы и имеют равномерное распределение на отрезке $[-1, 1]$. Найти $\xi_t, D\xi_t, \text{cov}(\xi_1, \xi_s)$.

Ч-5.6

Случайный процесс $\xi_t, -\infty < t < \infty$, задан формулой

$$\xi_t = \xi_1 \sin t + \xi_2 \sin \pi t$$

где ξ_1 и ξ_2 независимы и имеют равномерное распределение на отрезке $[-1, 1]$. Найти плотность распределения случайной величины $\eta = \sup_t \xi_t$.

Ч-5.7

Случайные величины $\alpha_1, \alpha_2, \dots$ независимы и одинаково распределены, $\alpha_1 = 0$, $D\alpha_1 = 1$. Случайные величины $\beta_1, \beta_2, \dots$ независимы, не зависят от $\alpha_1, \alpha_2, \dots$ и имеют равномерное распределение на отрезке $[-\pi, \pi]$. Последовательность случайных процессов $\eta_1(t), \eta_2(t), \dots$ определяется равенством

$$\eta_n(t) = \frac{1}{\sqrt{n}} \sum_{k=1}^n \alpha_k \cos(t + \beta_k).$$

- а) Найти предельное распределение $\eta_n(t)$ при $n \rightarrow \infty$, $t = \text{const}$.
 б) Доказать, что $\eta_n(t) = \theta_n \cos(t + \gamma_n)$. Найти предельное распределение вектора (γ_n, θ_n) при $n \rightarrow \infty$.

Ч-5.8

Случайный процесс $\varepsilon_t, 0 \leq t < \infty$, с вероятностью 1 имеет непрерывные траектории, $\varepsilon_t \equiv 0$, $D\varepsilon_t < \infty$, $\text{cov}(\varepsilon_t, \varepsilon_s) = R(|t - s|)$ при любых $t, s \geq 0$. Пусть η_t определяется случайная величина

$$\tau_N = \inf \{t \geq 0 : \xi_t \geq N\}.$$

Найти функцию распределения $F_N(x) = P\{\tau_N \leq x\}$ при $N > 0$ в следующих случаях:

- а) $\xi_t = \zeta t$, где $P\{\zeta \leq x\} = 1 - \frac{1}{(1+x)^k}$ ($x \geq 0, k > 0$);
 б) $\xi_t = t^2 - 2t\xi$, где $P\{\xi \leq x\} = 1 - e^{-\lambda x}$ ($x \geq 0$),
 в) $\xi_t = \zeta e^{at}$, где $a > 0$ и $P\{\zeta \leq x\} = 1 - e^{-\lambda x}$ ($x \geq 0$).

Ч-5.10

По последовательности $\xi_0, \xi_1, \dots$ независимых одинаково распределенных случайных величин построен случайный процесс $\eta_k = \xi_{k+1} - \xi_k, k = 0, 1, \dots$. Положим

$$\tau_1 = \inf \{k \geq 0 : \eta_k \leq 0\},$$

$$\tau_2 = \inf \{k \geq 0 : \eta_k < 0\}.$$

Найти $M\tau_1$ и τ_2 в следующих трех случаях:

- а) $P\{\xi_0 = 1\} = P\{\xi_0 = 2\} = 1/2$,
 б) $P\{\xi_0 = 1\} = P\{\xi_0 = 2\} = \dots = P\{\xi_0 = d\} = 1/d$,
 в) ξ_0 имеет равномерное распределение на отрезке $[0, 1]$

Ч-5.11

Найти распределение корня τ случайного уравнения $\xi x + \eta = 0$, где ξ и η - независимые случайные величины, имеющие стандартное нормальное распределение.

Ч-5.12

По последовательности $\xi_1, \xi_2, \dots$ независимых случайных величин, имеющих стандартное нормальное распределение, строится траектория случайного процесса $\xi_t (t \geq 0)$: $\xi_0 = 0, \xi_k = \xi_1 + \dots + \xi_n$ при любом $k \geq 1$, а при $t \in (k, k+1), k = 0, 1, \dots$, графиком траектории ξ_1 является отрезок, соединяющий точки (k, ξ_k) и $(k+1, \xi_{k+1})$, т. е.

$$\xi_t = (k+1-t)\xi_k + (t-k)\xi_{k+1} \text{ при } k < t < k+1.$$

Найти математическое ожидание числа μ_h нулей процесса ξ_1 на полуинтервале $[0, k)$ ($k \geq 1$ -целое) и асимптотику $M\mu_k$ при $k \rightarrow \infty$.

Ч-5.13

Случайный процесс $\xi_n(x), -\infty < x < \infty$, определяется равенством

$$\xi_n(x) = \xi_0 + \xi_1 x + \xi_2 x^2 + \dots + \xi_n x^n,$$

где случайная величина ξ_0 не зависит от случайного вектора $(\xi_1, \dots, \xi_n)$ и имеет непрерывную функцию распределения. Доказать, что при любых фиксированных x и a

$$P\{\xi_n(x) = a\} = 0.$$

Ч-5.14

Доказать, что если выполнены условия задачи п5.13, то с вероятностью 1 все действительные корни многочлена $\xi_n(x)$ - простые.

Ч-5.15

Случайный процесс $\xi_n(x)$, $-\infty < x < \infty$, определяется равенством

$$\xi_n(x) = \zeta_0 + \xi_1 x + \dots + \xi_n x^n,$$

где $\zeta_0, \xi_1, \dots, \xi_n$ - независимые случайные величины, имеющие нулевое математическое ожидание и единичную дисперсию. Найти $M\xi_n(x)$, $D\xi_n(x)$, $\text{cov}(\xi_n(x), \xi_n(y))$.

Ч-5.18

Случайные величины $\xi_1, \xi_2, \dots$ независимы и имеют одно и то же распределение на множестве $\{-1, 0, 1, 2, \dots\}$ с производящей функцией $f(s) = M_s^{k_1}$. По случайному блужданию

$$\rho_0 = 0, \rho_n = \xi_1 + \dots + \xi_n \quad (n \geq 1)$$

построим случайные величины $\tau_1, \tau_2, \dots$:

$$\tau_k = \begin{cases} \min \{n : \rho_n = -k\}, & \text{если } \inf_{n \geq 0} \rho_n \leq -k, \\ \infty & \text{в противном случае.} \end{cases} \quad c > 0$$

а) Доказать, что производящие функции $\varphi_k(s) = M s^{\tau_k}$ связаны соотношениями $\varphi_k(s) = \varphi_1^k(s)$, а производящая функция $\varphi_1(s)$ является решением уравнения $s f(\varphi_1(s)) = 1$.

б) Найти функцию $\varphi_1(s)$, если

$$P\{\xi_1 = k\} = (1-p)p^{k+1}, \quad k = -1, 0, 1, \dots$$

в) Найти функцию $\varphi_1(s)$, если

$$P\{\xi_1 = 1\} = p, \quad P\{\xi_1 = -1\} = 1-p, \quad 0 < p < 1.$$

Ч-5.19

По случайному блужданию ρ_n , определенному в задаче п5.18, построим случайную величину $\mu = \inf_{n \geq 0} \rho_n$. а) Найти распределение μ .

б) Доказать, что $P\{\mu > -\infty\} = 1$ тогда и только тогда, когда $M\xi_1 > 0$ или $P\{\xi_1 = 0\} = 1$.

Ч-5.20

Показать, что если $\rho_0 = 0$, $\rho_n = \xi_1 + \dots + \xi_n$ - случайное блуждание, построенное по независимым одинаково распределенным случайным величинам $\xi_1, \xi_2, \dots$

$$P\{\xi_n = 1\} = P\{\xi_n = -1\} = 1/2, \quad n = 1, 2, \dots,$$

то справедливы следующие соотношения:

$$\text{а) } M\rho_n = 0, \quad D\rho_n = n, \quad P\{\rho_{2n+1} = 0\} = 0, P\{\rho_{2n} = 0\} = C_{2n}^n 2^{-2n} \sim \frac{1}{\sqrt{\pi n}}, \quad n \rightarrow \infty;$$

$$\text{б) } M|\rho_n| = \sum_{k=0}^{n-1} P\{\rho_k = 0\} \sim \sqrt{\frac{2n}{\pi}}, \quad n \rightarrow \infty.$$

Ч-5.21

Случайные векторы $\xi_1, \xi_2, \dots$ в d -мерном евклидовом пространстве независимы и при любом $k = 1, 2, \dots$

$$M\xi_k = 0, \quad M|\xi_k|^2 = \sigma_k^2 < \infty,$$

где $|\xi_k|$ - евклидова длина ξ_k . Построим случайное блуждание $\rho_0 = 0, \rho_n = \xi_1 + \dots + \xi_n (n \geq 1)$. Доказать, что

$$M\rho_n = 0, \quad M|\rho_n|^2 = \sigma_1^2 + \dots + \sigma_n^2.$$

Ч-5.22

Случайные векторы $\xi_n = (\xi_{n,1}, \dots, \xi_{n,d})$ ($k = -1, 2, \dots$) в d -мерном евклидовом пространстве независимы и имеют равномерное распределение на единичной сфере

$$S^{d-1} = \{(x_1, \dots, x_d) \in R^d : x_1^2 + \dots + x_d^2 = 1\}.$$

Пусть $\rho_n = \xi_1 + \dots + \xi_n$ и $|\rho_n|$ - евклидова длина вектора ρ_n . Найти $M|\rho_n|^2$, $D|\rho_n|^2$.

Ч-5.23

Пусть выполнены условия задачи п5.22. Доказать, что при любом $\alpha > 3/2$

$$P\left\{\lim_{n \rightarrow \infty} \frac{|\rho_n|}{n^\alpha} = 0\right\} = 1.$$

Ч-5.24

*. Случайные величины $\xi, \xi_2, \dots$ независимы, принимают значения 1 и -1 с вероятностями p и $q = 1 - p$ соответственно, и $S_0 = 0$, $S_n = \xi_1 + \dots + \xi_n$, $n = 1, 2, \dots$. Обозначим через v число таких n , что $S_n = 0$.

а) Показать, что если $p \neq q$, то $P\{v < \infty\} = 1$; б) показать, что если $p = q = 1/2$, то $P\{v = \infty\} = 1$.

Ч-5.25

Независимые случайные векторы $\xi_n = (\xi_{n,1}, \dots, \xi_{n,s}) \in R^s$, $n = 1, 2, \dots$, имеют независимые компоненты и $P\{\xi_{n,i} = 1\} = P\{\xi_{n,i} = -1\} = \frac{1}{2}$, $1 \leq i \leq s$, $n = 1, 2, \dots$. Положим $S_n = \xi_1 + \dots + \xi_n$, $n = 1, 2, \dots$, и обозначим через v число таких n , что $S_n = (0, \dots, 0)$. Показать, что $P\{v = \infty\} = 1$ при $s = 2$ и что $P\{v < \infty\} = 1$ при $s \geq 3$.

Ч-5.26

(То же, что в В. а. л. д. а.) Пусть случайные величины $\xi_1, \xi_2, \dots$ независимы, $\xi_i = a$, и $|\xi_i| \leq C < \infty$ ($i \geq 1$). Пусть $B_1 \subset R^1, B_2 \subset R^2, B_3 \subset R^3, \dots$ - произвольная последовательность измеримых множеств, и случайная величина v определяется равенством

$$v = \inf\{n : (\xi_1, \dots, \xi_n) \notin B_n\}$$

($v = \infty$, если $(\xi_1, \dots, \xi_n) \in B_n$ при любом $n < \infty$). Используя равенство

$$\sum_{i=1}^v \xi_i = \sum_{i=1}^{\infty} \xi_i \chi\{v \geq i\} = \sum_{i=1}^{\infty} \xi_i (1 - \chi\{v < i\})$$

(где $\chi\{A\}$ - индикатор события A), доказать, что если $Mv < \infty$, то

$$M \sum_{i=1}^v \xi_i = aMv.$$

Ч-5.27

Проверить справедливость установленного в задаче п5.26 равенства в следующих случаях:

а) $P\{\xi_i = 1\} = P\{\xi_i = -1\} = 1/2$, $v = \min\{n \geq 0 : \xi_1 + \dots + \xi_n = 1\}$ б) $P\{\xi_i = 1\} = P\{\xi_i = -2\} = 1/2$, $v = \min\{n \geq 0 : \xi_1 + \dots + \xi_n = 1\}$. Объяснить результаты.

Ч-5.28

Случайные величины $\xi_1, \xi_2, \dots$ независимы и одинаково распределены, $a = M\xi_1 > 0$ и $P\{|\xi_1| \leq C\} = 1$ при некотором $C < \infty$. Пусть $\rho_0 = 0$, $\rho_n = \xi_1 + \dots + \xi_n$ ($n \geq 1$) и $N_1 = \inf\{n > 0 : S_n \geq t\}$. Доказать, что при любом $t > 0$

$$t \leq aMN_t \leq t + C.$$

Ч-5.29

Случайные величины $\rho_1, \rho_2, \dots$ при любом целом $n \geq 1$ удовлетворяют условию

$$M\{\rho_{n+1} | \rho_1, \dots, \rho_n\} \geq \rho_n.$$

Доказать, что при любых целых $k, n \geq 1$

$$M\{\rho_{n+k} | \rho_1, \dots, \rho_n\} \geq \rho_n.$$

Ч-5.30

*. Неотрицательные случайные величины $\rho_1, \rho_2, \dots$ при любом целом $n \geq 1$ удовлетворяют условию

$$M\{\rho_{n+1} | \rho_1, \dots, \rho_n\} \geq \rho_n.$$

Доказать, что для любого $\delta > 0$ и любого целого $n \geq 1$

$$P\{\max\{\rho_1, \dots, \rho_n\} \geq \delta\} \leq \frac{M\rho_n}{\delta}.$$

Ч-5.31

Случайные величины $\xi_1, \xi_2, \dots$ независимы, $M\xi_i = 0, i = 1, 2, \dots$. Доказать, что для любого $\delta > 0$ и любого целого $n \geq 1$

$$P\left\{\max_{1 \leq k \leq n} |\xi_1 + \dots + \xi_k| \geq \delta\right\} \leq \frac{M|\xi_1 + \dots + \xi_n|}{\delta}.$$

МФТИ-Т.17

Население региона делится по некоторому социально-экономическому признаку на три подгруппы. Следующее поколение с вероятностями 0,4; 0,6 и 0,2, соответственно, остается в своей подгруппе, а если не остается, то с равными вероятностями переходит в Любую из остальных подгрупп. Найти:

а) распределение населения по данному социально-экономическому признаку в следующем поколении, если в настоящем поколении в 1-ой подгруппе было 20% населения, во 2-ой подгруппе – 30%, и в 3-ей подгруппе – 50%

Данный процесс описывается марковской цепью с матрицей

$$\begin{pmatrix} 0,4 & 0,3 & 0,3 \\ 0,2 & 0,6 & 0,2 \\ 0,4 & 0,4 & 0,2 \end{pmatrix}$$

По условию начальное распределение было (0,2 0,3 0,5). Тогда распределение в следующем поколении

$$(x \ y \ z) = (0,2 \ 0,3 \ 0,5) \begin{pmatrix} 0,4 & 0,3 & 0,3 \\ 0,2 & 0,6 & 0,2 \\ 0,4 & 0,4 & 0,2 \end{pmatrix} = (0,34 \ 0,44 \ 0,22)$$

Найти б) предельное распределение по данному признаку, которое не меняется при смене поколений.

$$\begin{pmatrix} 0,4 - \lambda & 0,3 & 0,3 \\ 0,2 & 0,6 - \lambda & 0,2 \\ 0,4 & 0,4 & 0,2 - \lambda \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} \rightarrow \lambda^3 - 1,2\lambda^2 + 0,18\lambda + 0,02 = 0$$

$$\lambda_1 = 1, \lambda_2 = \frac{1 + \sqrt{3}}{10}, \lambda_3 = \frac{1 - \sqrt{3}}{10}$$

$$\begin{pmatrix} 0,4 - \lambda_1 & 0,3 & 0,3 \\ 0,2 & 0,6 - \lambda_1 & 0,2 \\ 0,4 & 0,4 & 0,2 - \lambda_1 \end{pmatrix} \begin{pmatrix} x_1 \\ y_1 \\ z_1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

$$\begin{pmatrix} -2x_1 & y_1 & z_1 \\ x_1 & -2y_1 & z_1 \\ x_1 & y_1 & -2z_1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

$$\begin{pmatrix} x_1 \\ y_1 \\ z_1 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$$

Аналогично получаем:

$$\begin{pmatrix} x_2 \\ y_2 \\ z_2 \end{pmatrix} = \begin{pmatrix} \frac{3}{4}(1 + \sqrt{3}) \\ \frac{1}{2}(-2 - \sqrt{3}) \\ 1 \end{pmatrix}$$

$$\begin{pmatrix} x_3 \\ y_3 \\ z_3 \end{pmatrix} = \begin{pmatrix} \frac{3}{4}(1 - \sqrt{3}) \\ \frac{1}{2}(-2 + \sqrt{3}) \\ 1 \end{pmatrix}$$

Тогда диагонализированная матрица

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{1+\sqrt{3}}{10} & 0 \\ 0 & 0 & \frac{1-\sqrt{3}}{10} \end{pmatrix}$$

и

$$\begin{aligned} \begin{pmatrix} 0,4 & 0,3 & 0,3 \\ 0,2 & 0,6 & 0,2 \\ 0,4 & 0,4 & 0,2 \end{pmatrix}^n &= \begin{pmatrix} 1 & \frac{3}{4}(1 + \sqrt{3}) & \frac{3}{4}(1 - \sqrt{3}) \\ 1 & \frac{1}{2}(-2 - \sqrt{3}) & \frac{1}{2}(-2 + \sqrt{3}) \\ 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{1+\sqrt{3}}{10} & 0 \\ 0 & 0 & \frac{1-\sqrt{3}}{10} \end{pmatrix}^n \begin{pmatrix} 1 & \frac{3}{4}(1 + \sqrt{3}) & \frac{3}{4}(1 - \sqrt{3}) \\ 1 & \frac{1}{2}(-2 - \sqrt{3}) & \frac{1}{2}(-2 + \sqrt{3}) \\ 1 & 1 & 1 \end{pmatrix}^{-1} = \\ &= \begin{pmatrix} 1 & \frac{3}{4}(1 + \sqrt{3}) & \frac{3}{4}(1 - \sqrt{3}) \\ 1 & \frac{1}{2}(-2 - \sqrt{3}) & \frac{1}{2}(-2 + \sqrt{3}) \\ 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \left(\frac{1+\sqrt{3}}{10}\right)^n & 0 \\ 0 & 0 & \left(\frac{1-\sqrt{3}}{10}\right)^n \end{pmatrix} \begin{pmatrix} 1 & \frac{3}{4}(1 + \sqrt{3}) & \frac{3}{4}(1 - \sqrt{3}) \\ 1 & \frac{1}{2}(-2 - \sqrt{3}) & \frac{1}{2}(-2 + \sqrt{3}) \\ 1 & 1 & 1 \end{pmatrix}^{-1} = \\ &= \begin{pmatrix} 1 & \frac{3}{4}(1 + \sqrt{3}) & \frac{3}{4}(1 - \sqrt{3}) \\ 1 & \frac{1}{2}(-2 - \sqrt{3}) & \frac{1}{2}(-2 + \sqrt{3}) \\ 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \left(\frac{1+\sqrt{3}}{10}\right)^n & 0 \\ 0 & 0 & \left(\frac{1-\sqrt{3}}{10}\right)^n \end{pmatrix} \begin{pmatrix} \frac{4}{13} & \frac{6}{13} & \frac{3}{13} \\ -\frac{4(-2+\frac{\sqrt{3}}{2})}{13\sqrt{3}} & -\frac{4(\frac{1}{4}+\frac{3\sqrt{3}}{4})}{13\sqrt{3}} & -\frac{4(\frac{7}{4}-\frac{5\sqrt{3}}{4})}{13\sqrt{3}} \\ -\frac{4(2+\frac{\sqrt{3}}{2})}{13\sqrt{3}} & -\frac{4(-\frac{1}{4}+\frac{3\sqrt{3}}{4})}{13\sqrt{3}} & -\frac{4(-\frac{7}{4}-\frac{5\sqrt{3}}{4})}{13\sqrt{3}} \end{pmatrix} \\ &\lim_{n \rightarrow +\infty} \left(\begin{pmatrix} 0,4 & 0,3 & 0,3 \\ 0,2 & 0,6 & 0,2 \\ 0,4 & 0,4 & 0,2 \end{pmatrix}^n \right) = \\ &= \begin{pmatrix} 1 & \frac{3}{4}(1 + \sqrt{3}) & \frac{3}{4}(1 - \sqrt{3}) \\ 1 & \frac{1}{2}(-2 - \sqrt{3}) & \frac{1}{2}(-2 + \sqrt{3}) \\ 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \frac{4}{13} & \frac{6}{13} & \frac{3}{13} \\ -\frac{4(-2+\frac{\sqrt{3}}{2})}{13\sqrt{3}} & -\frac{4(\frac{1}{4}+\frac{3\sqrt{3}}{4})}{13\sqrt{3}} & -\frac{4(\frac{7}{4}-\frac{5\sqrt{3}}{4})}{13\sqrt{3}} \\ -\frac{4(2+\frac{\sqrt{3}}{2})}{13\sqrt{3}} & -\frac{4(-\frac{1}{4}+\frac{3\sqrt{3}}{4})}{13\sqrt{3}} & -\frac{4(-\frac{7}{4}-\frac{5\sqrt{3}}{4})}{13\sqrt{3}} \end{pmatrix} \\ &= \begin{pmatrix} \frac{4}{13} & \frac{6}{13} & \frac{3}{13} \\ \frac{4}{13} & \frac{6}{13} & \frac{3}{13} \\ \frac{4}{13} & \frac{6}{13} & \frac{3}{13} \end{pmatrix} \end{aligned}$$

МФТИ-Т.18

Пусть матрица вероятностей перехода за один шаг цепи Маркова с двумя состояниями и имеет вид

$$\begin{pmatrix} 1-\alpha & \alpha \\ \beta & 1-\beta \end{pmatrix}, \quad 0 \leq \alpha, \quad \beta \leq 1$$

Найти вероятности перехода за n шагов и предельные вероятности.

$$\begin{pmatrix} 1-\alpha & \alpha \\ \beta & 1-\beta \end{pmatrix} \begin{pmatrix} x_{\pm} \\ y_{\pm} \end{pmatrix} = \lambda_{\pm} \begin{pmatrix} x_{\pm} \\ y_{\pm} \end{pmatrix} \rightarrow (1-\alpha-\lambda)(1-\beta-\lambda)-\alpha\beta=0 \rightarrow$$

$$\lambda^2 - \lambda(2-\alpha-\beta) + (1-\alpha)(1-\beta) - \alpha\beta = 0$$

$$\begin{aligned}
 \lambda_{\pm} &= \frac{2 - \alpha - \beta \pm \sqrt{(2 - \alpha - \beta)^2 - 4((1 - \alpha)(1 - \beta) - \alpha\beta)}}{2} = \\
 &= \frac{2 - \alpha - \beta \pm \sqrt{(1 - \alpha)^2 - 2(1 - \alpha)(1 - \beta) + (1 - \beta)^2 + 4\alpha\beta}}{2} = \\
 &= \frac{2 - \alpha - \beta \pm \sqrt{(\alpha - \beta)^2 + 4\alpha\beta}}{2} = \frac{2 - \alpha - \beta \pm \sqrt{\alpha^2 + 2\alpha\beta + \beta^2}}{2} = \frac{2 - \alpha - \beta \pm (\alpha + \beta)}{2} \rightarrow \\
 &\begin{cases} \lambda_+ = 1 \\ \lambda_- = 1 - (\alpha + \beta) \end{cases}
 \end{aligned}$$

$$\begin{pmatrix} 1 - \alpha - \lambda_{\pm} & \alpha \\ \beta & 1 - \beta - \lambda_{\pm} \end{pmatrix} \begin{pmatrix} x_{\pm} \\ y_{\pm} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \rightarrow \begin{pmatrix} (1 - \alpha - \lambda_{\pm})x_{\pm} + \alpha y_{\pm} \\ \beta x_{\pm} + (1 - \beta - \lambda_{\pm})y_{\pm} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \rightarrow \\
 \begin{pmatrix} x_{\pm} \\ y_{\pm} \end{pmatrix} = \begin{pmatrix} \alpha \\ -1 + \alpha + \lambda_{\pm} \end{pmatrix}$$

При $\alpha = \beta = 0$ тривиальный случай.

$$\begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \\
 \lim_{n \rightarrow +\infty} \begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix}^n = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

При $\alpha, \beta \in (0, 1]$ вырождения нет. Тогда диагонализированная матрица

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 - (\alpha + \beta) \end{pmatrix}$$

и

$$\begin{aligned}
 \begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix}^n &= \begin{pmatrix} \alpha & \alpha \\ \alpha & -\beta \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 - (\alpha + \beta) \end{pmatrix}^n \begin{pmatrix} \alpha & \alpha \\ \alpha & -\beta \end{pmatrix}^{-1} = \\
 &= \frac{1}{\alpha(\alpha + \beta)} \begin{pmatrix} \alpha & \alpha \\ \alpha & -\beta \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & (1 - (\alpha + \beta))^n \end{pmatrix} \begin{pmatrix} \beta & \alpha \\ \alpha & -\alpha \end{pmatrix} = \\
 &= \frac{1}{\alpha + \beta} \begin{pmatrix} \alpha(1 - (\alpha + \beta))^n + \beta & \alpha(1 - (1 - (\alpha + \beta))^n) \\ \beta(1 - (1 - (\alpha + \beta))^n) & \alpha + \beta(1 - (\alpha + \beta))^n \end{pmatrix}
 \end{aligned}$$

При $\alpha = 1, \beta = 1$

$$\begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix}^n = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

При $\alpha \neq 1, \beta \neq 1$

$$\lim_{n \rightarrow +\infty} \left(\begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix}^n \right) = \frac{1}{\alpha + \beta} \begin{pmatrix} \beta & \alpha \\ \beta & \alpha \end{pmatrix}$$

Определение параметров, при котором процесс докритический, критический, надкритический (??)

Для производящей функции процесса Гальтона-Ватсона $f(z) = \frac{1}{m+1-mz}$, $m > 0$ выясним, при каких значениях параметра m процесс является: докритическим, критическим, надкритическим.

Найдём математическое ожидание z :

$$\mathbb{E}(\xi) = \frac{d(f)}{d(z)} \Big|_{z=1} = - \frac{-m}{(m+1-mz)^2} \Big|_{z=1} = m$$

В итоге:

$m < 1$ - процесс докритический,

$m = 1$ - процесс критический,

$m > 1$ - процесс надкритический.

Найти вероятность вырождения процесса в надкритическом случае.

$$\begin{aligned} s = f(s) = \frac{1}{m+1-ms} \rightarrow s(m+1-ms) = ms + s - ms^2 = 1 \rightarrow ms^2 - s(m+1) + 1 = 0 \rightarrow \\ \rightarrow ms(s-1) - (s-1) = 0 \rightarrow (s-1)(ms-1) = 0 \rightarrow \begin{cases} s = 1 \\ s = \frac{1}{m} \end{cases} \end{aligned}$$

Показать, что n -тая итерация производящей функции может быть представлена в виде

$$\begin{aligned} f^n(z) &= \frac{m^n - 1 - m(m^{n-1} - 1)z}{m^{n+1} - 1 - m(m^n - 1)z} \\ \text{при } m \neq 1 \text{ и} \\ f^n(z) &= \frac{n - (n-1)z}{n+1 - nz} \end{aligned}$$

Докажем утверждение по индукции.

База индукции следующая:

$$\begin{aligned} f^2(z) = f(f(z)) = f\left(\frac{1}{m+1-mz}\right) &= \frac{1}{m+1-m\frac{1}{m+1-mz}} = \frac{m+1-mz}{(m+1)(m+1-mz)-m} \\ &= \frac{m+1-mz}{(m+1)^2 - (m+1)mz - m} \end{aligned}$$

при $m \neq 1$ (иначе случай будет ниже)

$$f^2(z) = \frac{(m-1)(m+1-mz)}{(m-1)((m+1)^2 - (m+1)mz - m)} = \frac{m^2 - 1 - m(m-1)z}{m^3 + 1 - m(m^2 - 1)z}$$

Шаг индукции:

$$\begin{aligned} f^{n+1}(z) = f(f^n(z)) &= f\left(\frac{m^n - 1 - m(m^{n-1} - 1)z}{m^{n+1} - 1 - m(m^n - 1)z}\right) = \\ &= \frac{1}{m+1-m\frac{m^n - 1 - m(m^{n-1} - 1)z}{m^{n+1} - 1 - m(m^n - 1)z}} = \\ &= \frac{m^{n+1} - 1 - m(m^n - 1)z}{(m+1)(m^{n+1} - 1 - m(m^n - 1)z) - m(m^n - 1 - m(m^{n-1} - 1)z)} = \\ &= \frac{m^{n+1} - 1 - m(m^n - 1)z}{(m+1)(m^{n+1} - 1) - m(m^n - 1) - m(m+1)(m^n - 1)z + m^2(m^{n-1} - 1)z} = \\ &= \frac{m^{n+1} - 1 - m(m^n - 1)z}{m^{n+2} - 1 - m(m^{n+1} - 1)z} \end{aligned}$$

теперь при $m = 1$ база индукции:

$$f^2(z) = \frac{m+1-mz}{(m+1)^2 - (m+1)mz - m} = \frac{(m+1) - (m)z}{m(m+1) + 1 - (m+1)z} = \frac{2 - (2-1)z}{2 + 1 - 2z}$$

и шаг индукции:

$$\begin{aligned} f^{n+1}(z) = f(f^n(z)) &= f\left(\frac{n - (n-1)z}{n+1 - nz}\right) = \\ &= \frac{1}{2 - \frac{n - (n-1)z}{n+1 - nz}} = \\ &= \frac{n+1 - nz}{2(n+1 - nz) - n + (n-1)z} = \\ &= \frac{n+1 - nz}{n+2 - (n+1)z} \end{aligned}$$

МФТИ-Т.20

Пусть $(W_t, t \geq 0)$ - винеровский процесс. Для $u \in \mathbb{R}$ определяется $\tau_u = \inf \{t : W_t = u\}$ - момент первого достижения уровня u траекторией винеровского процесса. Найти плотность распределения случайной величины τ_u

Если этот процесс впервые пересёк уровень u в точке τ_u , траекторию, исходящую из точки (τ_u, u) можно рассматривать, как самостоятельный винеровский процесс.

Следовательно, вероятность того, что в точке t будет выполнено $W_t \geq u$ равна $\frac{1}{2}$ в силу симметрии.

Заметим с другой стороны, что

$$P(W_t \geq u | \tau_u \leq t) = \frac{P(W_t \geq u)}{P(\tau_u \leq t)} = \frac{1}{2}$$

$$P(\tau_u \leq t) = 2P(W_t \geq u) = 2 \int_u^{+\infty} \frac{1}{\sigma\sqrt{2\pi t}} \exp\left(-\frac{x^2}{2\sigma^2 t}\right) dx = \frac{2}{\sqrt{2\pi}} \int_{\frac{u}{\sqrt{t}\sigma}}^{+\infty} \exp\left(-\frac{x^2}{2}\right) dx$$

$$\frac{\partial}{\partial(t)} (P(\tau_u \leq t)) = \frac{\partial}{\partial(t)} \left(\frac{2}{\sqrt{2\pi}} \int_{\frac{u}{\sqrt{t}\sigma}}^{+\infty} \exp\left(-\frac{x^2}{2}\right) dx \right) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{u^2}{2\sigma^2 t}\right) \frac{u}{t^{\frac{3}{2}}\sigma}$$

МФТИ-Т.21

Пусть

$$M_t = \max_{0 \leq s \leq t} W_s, \quad t > 0$$

где $(W_t, t \geq 0)$ - винеровский процесс.

Найти плотность распределения случайной величины M_t

$$P(M_t \geq x) = P(\tau_x \leq t) = 2 \int_x^{+\infty} \frac{1}{\sigma\sqrt{2\pi t}} \exp\left(-\frac{y^2}{2\sigma^2 t}\right) dy$$

$$P(M_t < x) = 1 - P(M_t \geq x) = 1 - 2 \int_x^{+\infty} \frac{1}{\sigma\sqrt{2\pi t}} \exp\left(-\frac{y^2}{2\sigma^2 t}\right) dy$$

$$1 - \sqrt{\frac{2}{\pi}} \int_{\frac{x}{\sigma\sqrt{t}}}^{+\infty} \exp\left(-\frac{y^2}{2}\right) dy = \sqrt{\frac{2}{\pi}} \int_0^{\frac{x}{\sigma\sqrt{t}}} \exp\left(-\frac{y^2}{2}\right) dy$$

В итоге:

$$\frac{\partial}{\partial(x)} (P(M_t < x)) = \frac{\partial}{\partial(x)} \left(\sqrt{\frac{2}{\pi}} \int_0^{\frac{x}{\sigma\sqrt{t}}} \exp\left(-\frac{y^2}{2}\right) dy \right) = \sqrt{\frac{2}{\pi}} \frac{1}{\sigma\sqrt{t}} \exp\left(-\frac{x^2}{2\sigma^2 t}\right)$$

МФТИ-Т.22.

Пусть $\mathbb{X} = (X_1, \dots, X_n)$ - выборка из пуассоновского распределения с параметром λ .

Найти оценку наибольшего правдоподобия параметра λ

Функция правдоподобия

$$f(\lambda) = \prod_{i=1}^n (P_\lambda(X_i = k_i)) = \prod_{i=1}^n \left(\frac{\lambda^{k_i}}{(k_i)!} \exp(-\lambda) \right) = \prod_{i=1}^n \left(\frac{1}{(k_i)!} \right) \lambda^{\sum_{i=1}^n k_i} \exp(-n\lambda)$$

$$\ln(f(\lambda)) = - \sum_{i=1}^n \ln((k_i)!) + \sum_{i=1}^n (k_i) \ln(\lambda) - n\lambda$$

Найдём максимум этой функции.

$$\frac{\partial}{\partial(\lambda)} (\ln(f(\lambda))) = \sum_{i=1}^n (k_i) \frac{1}{\lambda} - n = 0 \rightarrow \lambda = \frac{1}{n} \sum_{i=1}^n (k_i)$$

убедимся в максимуме с помощью второй производной:

$$\frac{\partial^2}{\partial(\lambda)^2} (\ln(f(\lambda))) = - \sum_{i=1}^n (k_i) \frac{1}{\lambda^2} = - \sum_{i=1}^n (k_i) \left(\frac{1}{n} \sum_{i=1}^n (k_i) \right)^{-2} = -n^2 \left(\sum_{i=1}^n (k_i) \right)^{-1} < 0$$

МФТИ-Т.23.

Пусть $\mathbb{X} = (X_1, \dots, X_n)$ - выборка из нормального распределения $N(a, \sigma^2)$.

Найти оценки наибольшего правдоподобия параметров a и σ^2

Функция правдоподобия:

$$\begin{aligned} f(a, \sigma) &= \prod_{i=1}^n (P_{a, \sigma}(X_i)) = \prod_{i=1}^n \left(\left(\frac{1}{\sqrt{2\pi}\sigma} \right) \exp \left(-\frac{(X_i - a)^2}{2\sigma^2} \right) \right) = \\ &= \left(\frac{1}{\sqrt{2\pi}\sigma} \right)^n \exp \left(-\frac{1}{2\sigma^2} \sum_{i=1}^n ((X_i - a)^2) \right) - \\ \ln(f(a, \sigma)) &= -n(\ln(\sqrt{2\pi}) + \ln(\sigma)) - \frac{1}{2\sigma^2} \sum_{i=1}^n ((X_i - a)^2) \\ \ln(f(a, \sigma)) &= -n(\ln(\sqrt{2\pi}) + \ln(\sigma)) - \frac{1}{2\sigma^2} \sum_{i=1}^n ((X_i - a)^2) \end{aligned}$$

Максимум этой функции:

$$\frac{\partial}{\partial(\sigma)}(\ln(f(a, \sigma))) = -\frac{n}{\sigma} + \frac{1}{\sigma^3} \sum_{i=1}^n ((X_i - a)^2) = 0$$

$$\sigma^2 = \frac{1}{n} \sum_{i=1}^n ((X_i - a)^2)$$

$$\frac{\partial}{\partial(a)}(\ln(f(a, \sigma))) = \frac{1}{\sigma^2} \sum_{i=1}^n (X_i - a) = 0$$

$$a = \frac{1}{n} \sum_{i=1}^n (X_i)$$

$$\frac{\partial^2}{\partial(\sigma)^2}(\ln(f(a, \sigma))) = \frac{n}{\sigma^2} - \frac{3}{\sigma^4} \sum_{i=1}^n ((X_i - a)^2)$$

$$\frac{\partial^2}{\partial(a)^2}(\ln(f(a, \sigma))) = -\frac{n}{\sigma^2}$$

$$\frac{\partial^2}{\partial(a)\partial(\sigma)}(\ln(f(a, \sigma))) = -\frac{2}{\sigma^3} \sum_{i=1}^n (X_i - a)$$

Матрица вторых производных в данной точке

$$\begin{pmatrix} -\frac{2n}{\sigma^2} & 0 \\ 0 & -\frac{n}{\sigma^2} \end{pmatrix} = -\frac{n}{\sigma^2} \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} < 0 \rightarrow \text{это точка строгого максимума}$$

задача про билеты Птф

(давно хотел прорешать, потом подумаю)

7.1.2 Problems About Poisson Processes**Ч-5.33**

Найти $\text{cov}(\xi_+, \xi_{++\odot})$ при $t, s \geq 0$, если:

- а) ξ_t - пуассоновский процесс на $[0, \infty)$ с интенсивностью λ ,
- б) ξ_t - пуассоновский процесс на $[0, \infty)$ с интенсивностью $\lambda(t)$.

Ч-5.34

Пусть $\{\tau_k\}_{k=1}^\infty$ ($0 < \tau_1 < \tau_2 < \dots$) - положения точек пуассоновского потока на $[0, \infty)$ с интенсивностью λ . Доказать, что при любом $T > 0$

$$\sum_{h=1}^{\infty} P\{\tau_k \leq T\} = \lambda T.$$

Ч-5.35

Пусть выполнены условия задачи п5.34. Найти $M\tau_k$, $D\tau_k$ и плотность $p_k(x)$ распределения τ_k . 160

Ч-5.36

Пусть выполнены условия задачи. п5.34. Найти:

- а) плотность $p(x_1, \dots, x_k)$ совместного распределения случайных величин $\tau_1, \tau_2, \dots, \tau_k$,
 б) $P\{\tau_1 \leq x | \tau_2 > T\}$, $0 \leq x \leq \infty$, в) $P\{\tau_1 \leq x | \tau_1 \leq T < \tau_2\}$, $0 \leq x \leq T$.

Ч-5.37

Пусть выполнены условия задачи п5.34. Найти условную плотность $p_T(x_1, \dots, x_k)$ совместного распределения случайных величин $\tau_1, \dots, \tau_h$ при условии, что $\tau_k \leq T < \tau_{k+1}$. Сравнить $p_T(x_1, \dots, x_k)$ с плотностью $q_T(x_1, \dots, x_h)$ совместного распределения членов варпадионного ряда $\xi_{(1)} \leq \xi_{(2)} \leq \dots \leq \xi_{(h)}$, построенного по независимым случайным величинам $\xi_1, \dots, \xi_k$, имеющим равномерное распределение на отрезке $[0, T]$.

Ч-5.43

Моменты поступления требований в систему массового обслуживания образуют пуассоновский поток на $(-\infty, \infty)$ с интенсивностью λ . Каждое требование независимо от остальных находится в едстве случайное ние числа ξ_1 требований в сиееме в момент t .

Ч-5.44

Для системы массового обслуживания, описанной в задаче п5.43, найти функцию $f(t) = P\{\xi_1 = 0 | \xi_0 = 0\}$. Убедиться в том, что $f(t)$ монотонно возрастает по t .

Ч-5.45

Случайно расположенные на плоскости точки образуют пуассоновское поле с интенсивностью λ , т. е. число точек в любой квадратуемой области S имеет пуассоновское распределение с математическим ожиданием $\lambda|S|$, где $|S|$ площадь области S , и числа точек в непересекающихся областях независимы. Пусть $\rho_1 \leq \rho_2 \leq \dots$ - упорядоченные по возрастанию расстояния от начала координат до точек этого поля.

- а) Как можно описать последовательность $\{\rho_k\}_{k=1}^\infty$?
 б) Найти плотность $p_n(x)$ распределения ρ_n , точную Формулу для ρ_n и асимптотику $M\rho_n$ при $n \rightarrow \infty$.

Ч-5.46

Случайно расположенные в пространстве R^3 точки образуют пуассоновское поле с интенсивностью λ , т. е. число точек в любой измеримой области V имеет пуассоновское распределение с математическим ожиданием $\lambda|V|$, где $|V|$ - объем области V , а числа точек в непересекающихся областях независимы. Ответить на те же вопросы, что в задаче 5.45.

Ч-5.47

Автобусы прибывают на остановку в случайные моменты времени $\tau_1 < \tau_2 < \dots$. Случайные величины $\tau_1, \tau_2, \dots$ независимы и τ_n имеет равномерное распределение в интервале $(n-1, n)$.

- а) Найти плотность $p(x)$ распределения интервала $\delta_n = \tau_{n+1} - \tau_n$ между автобусами.
 б) Найти плотность $p_k(x_1, x_2)$ распределения двумерного вектора $(\delta_n, \delta_{n+\lambda})$ при $k = 1, 2, \dots$
 в) Моменты ξ и ξ прихода двух пассажиров на остановку независимы и имеют равномерное распределение на интервале $(0, 1)$. Найти вероятность того, что они поедут на одном и том же автобусе.
 г) Моменты ξ_1 и ξ_2 прихода двух пассажиров на остановку независимы, ξ_1 имеет равномерное распределение на интервале $(0, 1)$, а ξ_2 - на интервале $(1, 2)$. Найти вероятность того, что эти пассажиры поедут на одном автобусе.
 д) Пассажир приходит на остановку в случайный момент ξ , имеющий равномерное распределение на интервале $(0, 1)$. Найти математическое ожидание, дисперсию и плотность $q(x)$ распределения времени ожидания автобуса.

7.1.3 Problems About Markov Chains (!!!)

(займусь теорвером - в первый же день по ним тренироваться буду)


Белан-ВШЭл5.у1.(3 мин) Последовательность бросаний кубика и монеты

Допустим, берется симметричный кубик и бросается раз за разом до первого выпадения грани с шестью очками. После того как шесть очков выпали, берется симметричная монетка и подбрасывается до тех пор, пока не выпадет первый орел. После того как орел выпал, снова берут кубик и подбрасывают его до первого появления шести очков. И так далее. Вы должны угадать, кубиком или монеткой будет выполнен седьмой бросок в этой игре?

Solution

Решаем самым обычным уравнением для марковских цепей, просто 1 формулой нужно воспользоваться.

Введем цепь Маркова с двумя состояниями и с дискретным временем. Будем говорить, что процесс находится в состоянии 1 в момент времени n , если $n + 1$ -й бросок в описанной игре осуществляется кубиком, и в состоянии 2, если $n + 1$ -й бросок выполняется монеткой. Схема наша как всегда:



Мы имеем $p = 1/6, q = 1/2$. Вероятности нахождения процесса в различных состояниях на n -м шаге даются выведенными формулами из возведения матрицы перехода в n -ю степень а также положить $\pi_1(0) = 1$ и $\pi_2(0) = 0$:

$$\pi_1(n) = \frac{q}{p+q} + \frac{p}{p+q}(1-p-q)^n = \frac{3}{4} + \frac{1}{4 \cdot 3^n},$$

$$\pi_2(n) = \frac{p}{p+q} - \frac{p}{p+q}(1-p-q)^n = \frac{1}{4} - \frac{1}{4 \cdot 3^n}.$$

Итак, $\pi_1(6) > \pi_2(6)$, а значит с большей вероятностью седьмой бросок будет выполняться кубиком.

Задача в 2 строчки решилась цепью Маркова. А иначе и не решить ее вообще! (????? подумаю больше про разные методы для нее.)

Белан-ВШЭл5.у2.(3 мин) Переходы между тремя равновероятными состояниями

Марковский процесс с непрерывным временем может находиться в одном из трех состояний, 1, 2 и 3, причем вероятность перехода между любыми двумя из них равна λ . Найдите вероятности, $\pi_1(t), \pi_2(t)$ и $\pi_3(t)$, найти систему в различных состояниях в произвольный момент времени $t > 0$, если известно, что в нулевой момент времени система находилась в состоянии 1.

Solution

Пишем управляющее уравнение, элементарно его решаем.

Запишем управляющее уравнение на вероятности нахождения марковского процесса в различных состояниях

$$\begin{aligned}\dot{\pi}_1(t) &= -2\lambda\pi_1(t) + \lambda\pi_2(t) + \lambda\pi_3(t), \\ \dot{\pi}_2(t) &= \lambda\pi_1(t) - 2\lambda\pi_2(t) + \lambda\pi_3(t) \\ \dot{\pi}_3(t) &= \lambda\pi_1(t) + \lambda\pi_2(t) - 2\lambda\pi_3(t).\end{aligned}$$

Легко понять, что в силу начального условия ($\pi_1(0) = 1, \pi_2(0) = \pi_3(0) = 0$) и симметрии вероятностей переходов, вероятности нахождения процесса во втором и третьем состояниях одинаковы, то есть $\pi_2(t) = \pi_3(t) = (1 - \pi_1(t))/2$. Отсюда получаем замкнутое уравнение на эволюцию вероятности $\pi_1(t)$

$$\dot{\pi}_1(t) = -3\lambda\pi_1 + \lambda$$

решение которого (с учетом начальных условий) имеет вид

$$\pi_1(t) = \frac{1}{3} + \frac{2}{3}e^{-3\lambda t}$$

и, следовательно,

$$\pi_2(t) = \pi_3(t) = \frac{1}{3} - \frac{1}{3}e^{-3\lambda t}$$

Белан-ВШЭл5.у3. Занятость телефонной линии

Телефонная линия может находиться в одном из двух состояний: (1) линия свободна и (2) линия занята. Длительности интервалов времени между разговорами на линии и длительности самих разговоров статистически независимы друг от друга и имеют экспоненциальные функции распределения, $p_1(\tau) = \lambda e^{-\lambda\tau}$ и $p_2(\tau) = \mu e^{-\mu\tau}$, соответственно.

а) Предположим, что в начальный момент времени линия свободна и мы наблюдаем за ее стохастической динамикой в течении времени T . Обозначим через $r_2(T)$ долю времени, которое линия была занята в течении окна наблюдения. Вычислите математическое ожидание и дисперсию случайной величины $r_2(T)$.

б) Известно, что в нулевой момент времени и спустя время T линия свободна. Найдите вероятность того, что линия занята в промежуточные моменты времени $0 < t < T$.

Solution a)

Рассмотрим стохастическую переменную

$$S(t) = \begin{cases} 1, & \text{если линия свободна в момент } t, \\ 2, & \text{если линия занята в момент } t. \end{cases}$$

Вероятности $\pi_1(t) \equiv \Pr[S(t) = 1]$ и $\pi_2(t) \equiv \Pr[S(t) = 2]$ подчиняются системе уравнений (21) и (22) с начальным условием $\pi_1(0) = 1, \pi_2(0) = 0$. Случайная величина $r_2(T)$, статистикой которой мы интересуемся, может быть представлена как

$$r_2(T) = \frac{\int_0^T I(S(t) = 2) dt}{T}$$

Тогда

$$\langle r_2(T) \rangle = \frac{\int_0^T \langle I(S(t) = 2) \rangle dt}{T} = \frac{\int_0^T \Pr[S(t) = 2 | S(0) = 1] dt}{T},$$

где $\Pr[S(t) = 2 | S(0) = 1]$ это вероятность обнаружить процесс в состоянии 2 в момент времени t при условии, что он находился в состоянии 1 в момент времени 0. Очевидно, $\Pr[S(t) = 2 | S(0) = 1]$ дается решением $\pi_2(t)$ (см формулу (29)) с начальными условиями $\pi_1(0) = 1, \pi_2(0) = 0$. То есть

$$\Pr[S(t) = 2 | S(0) = 1] = \frac{\lambda}{\lambda + \mu} (1 - \exp[-(\lambda + \mu)t])$$

Подставляя Ур. (43) в Ур. (42) и интегрируя, находим

$$\langle r_2(T) \rangle = \frac{\lambda}{\lambda + \mu} - \frac{\lambda}{(\lambda + \mu)^2} \frac{1 - \exp[-(\lambda + \mu)T]}{T}.$$

Итак, в пределе бесконечно большого окна наблюдения, $T \rightarrow \infty$, математическое ожидание доли времени, проведенного процессом в состоянии 2, равно статистически равновесной вероятности π_2^{st} найти процесс в этом состоянии.

Solution б)

Искомая вероятность может быть вычислена при помощи формулы Байеса. Действительно, с учетом марковского свойства вероятность $\Pr[S(t) = 2, S(T) = 1 | S(0) = 1]$ можно представить как

$$\begin{aligned} \Pr[S(t) = 2, S(T) = 1 | S(0) = 1] &= \Pr[S(T) = 1 | S(t) = 2, S(0) = 1] \Pr[S(t) = 2 | S(0) = 1] = \\ &= \Pr[S(T) = 1 | S(t) = 2] \Pr[S(t) = 2 | S(0) = 1] \end{aligned}$$

или как

$$\Pr[S(t) = 2, S(T) = 1 | S(0) = 1] = \Pr[S(t) = 2 | S(0) = 1, S(T) = 1] \Pr[S(T) = 1 | S(0) = 1]$$

Приравнявая эти два представления, получаем

$$\Pr[S(t) = 2 | S(0) = 1, S(T) = 1] = \frac{\Pr[S(T) = 1 | S(t) = 2] \Pr[S(t) = 2 | S(0) = 1]}{\Pr[S(T) = 1 | S(0) = 1]},$$

где

$$\Pr[S(T) = 1 | S(t) = 2] = \frac{\mu}{\lambda + \mu} \left(1 - e^{-(\lambda + \mu)(T-t)} \right),$$

$$\Pr[S(t) = 2 | S(0) = 1] = \frac{\lambda}{\lambda + \mu} (1 - \exp[-(\lambda + \mu)t]),$$

$$\Pr[S(T) = 1 | S(0) = 1] = \frac{\mu}{\lambda + \mu} + \frac{\lambda}{\lambda + \mu} e^{-(\lambda + \mu)T}.$$

Белан-ВШЭ-л3.38 О принципе больших отклонений (??)

Пусть $X_1, X_2, \dots, X_N$ - набор независимых случайных величин, каждая из которых имеет функцию распределения $p(x) = (1-q)\delta(x) + q\delta(x-1)$. Используя теорему Гартнера-Эллиса, выясните, удовлетворяет ли функция плотности вероятности $P_N(a)$ среднего арифметического $A = \frac{1}{N} \sum_{i=1}^N X_i$ принципу больших отклонений. Если да, то найдите соответствующую функцию Крамера $I(a)$.

(?? хз, о чем задача??)

Вычисляем производящую функцию кумулянтов

$$\lambda(\xi) = \langle e^{\xi X} \rangle = \ln \left(\frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{+\infty} \exp \left(-\frac{(x-\mu)^2}{2\sigma^2} + \xi x \right) dx \right) = \xi\mu + \frac{\sigma^2 \xi^2}{2}$$

Видим, что она дважды дифференцируема. Тогда согласно теореме Гартнера-Эллиса распределение $P_N(a)$ удовлетворяет принципу больших отклонений с функцией Крамера

$$I(a) = \sup_{\xi \in \mathbb{R}} \{ \xi a - \lambda(\xi) \} = \xi_* a - \lambda(\xi_*) = \frac{(a - \mu)^2}{2\sigma^2}$$

где $\xi_* = \frac{a - \mu}{\sigma^2}$ - точка максимума функции $\xi a - \lambda(\xi)$. Значит

$$P_N(a) \propto e^{-NI(a)} = e^{-\frac{N(a-\mu)^2}{2\sigma^2}}$$

Белан-ВШЭ-л1.дз-9 Энтропии выпадения орлов (???)

Пусть X это случайная величина, которая равна 1, если в эксперименте с подкидыванием монетки выпадает орел, и равна 0 в случае решки. Далее, Y это независимая от X случайная величина, которая равна 1, если в эксперименте с подкидыванием другой монетки выпадает орел, и равна 0 в случае решки. Наконец, Z это случайная величина равная 1, если в результате подкидывания двух вышеупомянутых монеток выпал хотя бы один орел, и ноль в любом другом случае. Найдите (а) частные энтропии, $S(X)$, $S(Y)$ и $S(Z)$, (б) совместные энтропии, $S(X, Y)$, $S(X, Z)$ и $S(Y, Z)$, а также (в) взаимные информации, $I(X, Y)$, $I(X, Z)$ и $I(Y, Z)$

(?? хз, о чем задача??)

Solution

ФР дисперсия 4х подряд орлов

Задача 10 (12 баллов) Монетку подбрасывают до тех пор пока в череде исходов не встретится последовательность из четырех подряд идущих орлов. Найти дисперсию требуемого числа попыток.

(решение другого решателя, тут какая-то копия, оно крупнейшее какое-то. не копирую, потому что все равно не уверен в этом. жесть, а не задача.)

ФР пассажиры на автобусной стоянке

Задача 12 (10 баллов) Пассажиры приходят на автобусную остановку независимо друг от друга со средней частотой λ пассажиров 6 минут. Прибытие автобусов это пуассоновские события со средней частотой μ автобусов в минуту. Будем предполагать, что вместимость автобуса достаточно велика, чтобы всякий раз принять всех скопившихся на остановке пассажиров. Кроме того, будем пренебрегать временем, необходимым для посадки пассажиров в автобус.

(а) Вычислить стационарную вероятность p_n найти на остановке ровно n пассажиров. Каковы условия существования статистически-стационарного режима?

(б) Вычислить среднее количество людей, ожидающих автобус на остановке, в статистически-стационарном режиме.

(6) Каково среднее число пассажиров, которые уедут с остановки в промежутке времени с 14:00 до 16:00, если $\lambda = 2 \text{ мин}^{-1}$ и $\mu = 0.2 \text{ мин}^{-1}$?

(элементарная задача, школой решается)

$$\pi_n(t + \varepsilon t) = (\lambda, \varepsilon_t) \pi_{n-1} + (1 - h\delta t) \pi_n + (1 - \mu dt) \pi_t$$

$$\pi_n = -\alpha \pi_n - \mu \pi_n + h \pi_{n-1}$$

$$6 \text{ в стац случае } \dot{\pi} = 0$$

$$\pi_n = \frac{A}{x + \mu} \pi_{n=1}$$

$$\text{аАа } \pi_0 : \pi_0 = \mu(1 - \pi_0) - \lambda \pi_0 \rightarrow \pi_0 = \frac{\mu}{R + M}$$

$$\pi_h = \left(\frac{h}{L\pi H} \right)^h \frac{\mu}{L + \mu} = \frac{\mu}{L} \left(\frac{\lambda}{L + \mu} \right)^{h+1}$$

$$\langle h \rangle = \mu = \sum_h n \pi_n = \frac{\mu}{\alpha + \mu} \sum_h n \left(\frac{\alpha}{L + \mu} \right)^n = \frac{A}{A}$$

Ч-5.49

Цепь Маркова ξ_t имеет множество состояний $\{-6, -5, \dots, 0, 1, \dots, 6\}$. Переходные вероятности $p_{ij} = P\{\xi_{t+1} = j | \xi_t = i\}$ при $i \neq 0$ определяются соотношениями $p_{ij} = \begin{cases} 1, & \text{если } j = i + 1 \leq 0 \\ 0 & \text{в остальных случаях.} \end{cases}$ или $j = i - 1 \geq 0$. Провести классификацию цепи Маркова и множества ее состояний, если:

а) $p_{0,6} = 1, p_{0,i} = 0 \quad (i \neq 6)$,

б) $p_{0,6} = p_{0,-6} = 1/2, p_{0,1} = 0 \quad (|i| \neq 6)$, в) $p_{0,6} = p_{0,-5} = 1/2, p_{0,4} = 0 \quad (i \neq -5, i \neq 6)$.

Ч-5.50

Матрица вероятностей перехода цепи Маркова имеет вид

$$P = \begin{pmatrix} 0,1 & 0,5 & 0,4 \\ 0,6 & 0,2 & 0,2 \\ 0,3 & 0,4 & 0,3 \end{pmatrix}.$$

Распределение по состояниям в момент времени $t = 0$ определяется вектором $(0, 7; 0, 2; 0, 1)$. Найти: 1) распределение по состояниям в момент $t = 2$; 2) вероятность того, что в моменты $t = 0, 1, 2, 3$ состояниями цепи будут соответственно 1, 3, 3, 2; 3) стационарное распределение.

Ч-5.51

Пусть ξ_t - номер состояния в цепи Маркова в момент времени t ; $P(\xi_0 = 1) = 1$, и матрица вероятностей перехода имеет вид

$$\begin{pmatrix} 3/7 & 3/7 & 1/7 \\ 1/11 & 2/11 & 8/11 \\ 1/11 & 4/11 & 6/11 \end{pmatrix}$$

положим

$$\eta_t = \begin{cases} 1, & \text{если } \xi_t = 1, \\ 2, & \text{если } \xi_t \neq 1. \end{cases}$$

Показать, что последовательность η_t является цепью Маркова. Найти ее матрицу вероятностей перехода.

Ч-5.52

Случайные величины $\xi_i, i = 1, 2, \dots$, независимы и

$$P(\xi_i = 1) = p, \quad P(\xi_i = -1) = 1 - p = q.$$

Положим $\eta_0 = 0; \eta_{t+1} = \eta_t + \xi_{t+1}$. Является ли последовательность η_t цепью Маркова? Найти $P(\eta_t = m), \quad m = 0, 1, 2, \dots$

Ч-5.53

Пусть $\xi_i, i = 1, 2, \dots$, - независимые случайные величины, $P(\xi_i = 1) = 1 - P(\xi_i = -1) = p$ являются ли цепями Маркова последовательности случайных величин:

а) $\eta_t = \xi_t \xi_{t+1}$ б) $\eta_t = \xi_1 \xi_2 \dots \xi_t$;

в) $\eta_t = \varphi(\xi_t, \xi_{t+1})$, где $\varphi(-1, -1) = 1, \varphi(-1, 1) = 2, \varphi(1, -1) = 3, \varphi(1, 1) = 4$?

Для цепей Маркова найти вероятности перехода за один шаг.

Ч-5.54

В N ячейках последовательно независимо друг от друга равновероятно размещают частицы. Пусть $\mu_0(n)$ - число ячеек, оставшихся пустыми после размещения n частиц. Показать, что последовательность $\mu_0(n), n = 1, 2, \dots$, является цепью Маркова. Найти вероятности перехода.

Ч-5.55

В N ячейках независимо друг от друга размещаются комплекты, состоящие из m частиц. Частицы одного комплекта размещаются в ячейках по одной, и все возможные выборы m мест из N имеют одинаковые вероятности. Обозначим через $\mu_0(n)$ число ячеек, оставшихся пустыми после размещения n комплектов. Показать, что последовательность $\mu_0(n), n = 1, 2, \dots$ является цепью Маркова. Найти вероятности перехода.

Ч-5.56

Урна вначале содержит N белых шаров. За 1 шаг из урны по схеме случайного равновероятного выбора вынимают 1 шар и заменяют его новым, который независимо от предыстории процесса - является черным с вероятностью p и белым с вероятностью $q = 1 - p$. Обозначим через ξ_n число белых шаров в урне после n -го шага.

а) Найти $p_{ij} = P\{\xi_{n+1} = j | \xi_n = i\}, i, j \in \{0, 1, \dots, N\}$. Является ли последовательность ξ_n цепью Маркова?

б) Найти $\lim_{n \rightarrow \infty} P\{\xi_n = k\} = \pi_k, \quad k \in \{0, 1, \dots, N\}$.

Ч-5.57

В бесконечной последовательности пронумерованных шаров каждый шар независимо от остальных является черным с вероятностью p и белым с вероятностью $q = 1 - p$. Будем считать, что в момент времени $n \in \{0, 1, \dots\}$ урна содержит шары с номерами $n + 1, n + 2, \dots, n + N$, и обозначим через ξ_n число белых шаров в урне в момент n . Ответить на те же вопросы, что в задаче 5.56.

Ч-5.58

Матрица $P = \|p_{ij}\|_{i,j=1}^N$ с неотрицательными элементами $p_{i2} + \dots + p_{iN} = p_{1i} + p_{2i} + \dots + p_{Ni} = 1$ для любого $i = 1, 2, \dots, N$. Показать, что если цепь Маркова ξ с состояниями $1, \dots, N$ имеет дважды стохастическую матрицу вероятностей перехода $P \rightsquigarrow \|p_{ij}\|_{i,j=1}^N$, то равномерно распределение на множестве $\{1, \dots, N\}$ является стационарным для цепи, ξ_t^n .

Ч-5.59

Пусть $\zeta_0, \zeta_1, \dots$ цепь Маркова с множеством состояний $\{1, 2, 3\}$, матрицей вероятностей перехода $\|p_{ij}\|$ и стационарным распределением π_f . Показать, что если $p_{11} = p_{22} = p_{33} = 0$ и $\pi_1 = \pi_2 = \pi_3 = 1/3$, то $p_{12} = p_{28} = p_{31}$ и $p_{13} = p_{32}$.

Ч-5.60

Пусть $\eta_1; \eta_2, \dots$ - последовательность независимых одинаково распределенных случайных величин, и $f(x, y)$ - функция, принимающая значения в множестве $\{1, \dots, N\}$. Является ли последовательность случайных величин

$$\xi_0 (P(\xi_0 = k) = p_k^{(0)}), \quad \xi_{t+1} = f(\xi_t, \eta_{t+1})_i \\ t = 0, 1, 2, \dots \times$$

цепью Маркова?

Ч-5.61

Пусть $\eta_0, \eta_1, \eta_2, \dots$ последовательность независимых случайных величин, имеющих равномерное распределение на отрезке $[0, 1]$. Показать, что для любого начального распределения $p^{(0)} = (p_1^{(0)}, \dots, p_N^{(0)})$ и для любой матрицы переходных вероятностей $P = \|p_{ij}\|_{i,j=1}^N$ можно указать такие функции $f(x, y)$ и $g(y)$ ($x \in \{1, \dots, N\}$, $y \in [0, 1]$), что последовательность $\varepsilon_0 g(\eta_0)$, $\xi_{t+1} = f(\xi_t, \eta_{t+1})$, $t = 0, 1, \dots$ будет цепью Маркова с начальным распределением $p^{(0)}$ и матрицей вероятностей перехода P .

Ч-5.62

Игральная кость последовательно перекладывается с одной грани равномерно на любую из четырех соседних независимо от предыдущего. К какому пределу стремится при $t \rightarrow \infty$ вероятность того, что при t -м перекладывании кость окажется на грани «6», если сначала она падалась в этом же положении?

Ч-5.63 Свойства цепи Маркова

Цепь Маркова ξ имеет два состояния: 1 и 2 - и матрицу вероятностей перехода $\begin{pmatrix} p_{11} & 1-p_{11} \\ 0 & 1 \end{pmatrix}$.

Пусть $\tau \equiv \min \{t \geq 1 : \xi_t = 2\}$.

- Найти условное распределение τ при условии $\xi_0 = 1$.
- Как изменится ответ, если матрица вероятностей перехода будет равна

$$\begin{pmatrix} p_{11} & 1-p_{11} \\ 1-p_{22} & p_{22} \end{pmatrix}, \quad 0 < p_{22} < 1?$$

Ч-5.64

Матрица вероятностей перехода $\|p_{ij}\|$ цепи Маркова ξ с $n+1$ состояниями имеет вид: $p_{il} = 1 - \alpha$, $l = 1, 2, \dots, n+1$; $p_{ij} = \alpha/n$, $i \neq j$. В процессе, начавшемся из состояния k , $k \neq n+1$, обозначим через τ_n момент первого попадания в состояние $n+1$. Подобрать последовательность A_n ($A_n \rightarrow \infty$ при $n \rightarrow \infty$) так, чтобы существовал предел

$$\lim_{n \rightarrow \infty} P(\tau_n \leq A_n < x)$$

и найти его.

Ч-5.65

Матрица вероятностей перехода цепи Маркова $\xi_n^{(\varepsilon)}$ с множеством состояний $\{1, 2, 3\}$ имеет вид

$$\begin{pmatrix} 1-p & p & 0 \\ 1-p-\varepsilon & p & \varepsilon \\ 0 & \varepsilon^2 & 1-\varepsilon^2 \end{pmatrix}$$

где $0 < p < 1 - \varepsilon \leq 1$. Предполагая, что $\xi_0^{(\varepsilon)} = 1$, рассмотрим случайную величину $\tau_1(\varepsilon) = \min \{n \geq 1 : \xi_n^{(\varepsilon)} = 1\}$ - время возвращения в состояние 1.

- Доказать, что при любом $t = 1, 2, \dots$

$$\lim_{\varepsilon \rightarrow 0} P\{\tau_1(\varepsilon) = t\} = P\{\tau_1(0) = t\}.$$

- Показать, что $M\tau_1(0) < \infty$, но $\lim_{\varepsilon \rightarrow 0} M\tau_1(\varepsilon) = \infty$.
- Найти производящую функцию $\varphi_\varepsilon(s) = M s^{\tau_1(\varepsilon)}$.

Ч-5.66

Цепь Маркова ξ_n имеет начальное состояние $\xi_0 = 0$ и переходные вероятности

$$P\{\xi_{n+1} = k+1 | \xi_n = k\} = p, \quad P\{\xi_{n+1} = k | \xi_n = k\} = 1-p,$$

где $k, n = 0, 1, \dots$ и $0 < p < 1$. Найти распределение ξ_n , математическое ожидание и дисперсию ξ_n . 166

Ч-5.67

По цепи Мартова ξ_n , описанной в задаче

Ч-5.66

, построим последовательность случайных величин

$$\tau_0 = 0, \quad \tau_k = \min\{n : \xi_n = k\}, \quad k = 1, 2, \dots$$

- а) Доказать, что $\{\tau_k\}_{k=0}^\infty$ — тоже цепь Маркова. Найти ее переходные вероятности.
 б) Найти $M\tau_k$, D_{τ_k} , $\varphi_k(s) = Ms^{\tau_k}$.

Ч-5.68

Цепь Маркова ξ_n имеет начальное состояние $\xi_0 = 0$ и переходные вероятности

$$P\{\xi_{n+1} = k+1 | \xi_n = k\} = p, \quad P\{\xi_{n+1} = 0 | \xi_n = k\} = 1-p,$$

где $k, n \in \{0, 1, \dots\}$ и $0 < p < 1$. Положим

$$\tau_0 = 0, \quad \tau_k = \min\{n \geq 1 : \xi_n = k\}, \quad k = 1, 2, \dots$$

- а) Найти производящую функцию $\varphi_k(s) = Ms^{\tau_k}$ и $M\tau_k$.
 б) Найти предельное распределение случайной величины $\eta_h = \frac{\tau_h}{M\tau_h}$ при $k \rightarrow \infty$, пользуясь результатом п. а) и методом характеристических функций.

Ч-5.69

Переходные вероятности цепи Маркова ξ_n при любом $n = 0, 1, \dots$ определяются равенствами

$$\begin{aligned} P\{\xi_{n+1} = 1 | \xi_n = 0\} &= p, \quad P\{\xi_{n+1} = 0 | \xi_n = 0\} = 1-p, \\ P\{\xi_{n+1} = k+1 | \xi_n = k\} &= p, \\ P\{\xi_{n+1} = k-1 | \xi_n = k\} &= 1-p, \quad k = 1, 2, \dots \end{aligned}$$

Пусть $\tau_{i,j}$ — время перехода цепи ξ_n из состояния i в состояние j :

$$P\{\tau_{i,j} = t\} = P\{\min\{n \geq 1 : \xi_n = j\} = t | \xi_0 = i\}, \quad t = 1, 2, \dots$$

- а) Найти производящую функцию $\varphi_{0,1}(s) = Ms^{\tau_{0,1}}$ и $M\tau_{0,1}$.
 б) Найти рекуррентные формулы, связывающие производящие функции $\psi_{k,k+1}(s) = Ms^{\tau_{k,k+1}}$ ($k = 0, 1, \dots$).
 в) Используя результаты п. а) и б), составить и решить рекуррентные уравнения для $\tau_{k,k+1}$, $k = 0, 1, \dots$.
 Найти асимптотические формулы для $M\tau_{0,k}$ при $k \rightarrow \infty$.

Ч-5.70

Пусть $\xi_1, \xi_2, \dots$ — независимые одинаково распределенные случайные величины, $P\{\xi_1 = 1\} = P\{\xi_1 = -1\} = 1/2$; $z_0 = 0$, $z_n = z_{n-1} + \xi_n$, $k = 1, 2, \dots$. Положим $\tau_N = \min\{n \geq 1 : |z_n| = N\}$. Найти $M\tau_1, M\tau_2, M\tau_3$.

Ч-5.71

Пусть последовательность $z_0, z_1, \dots$ определена как в задаче 5.70. Показать, что $M\tau_N = N^2$.

Ч-5.77

Матрица вероятностей перехода P и распределение q по состояниям в момент $t = 0$ цепи Маркова ξ , имеют вид

$$P = \begin{pmatrix} z/12 & 2/12 & 1/12 & 3/12 & 1/12 & 2/12 \\ 1/12 & 1/12 & 8/12 & 1/12 & 4/12 & 2/12 \\ 0 & 0 & 3/4 & 1/4 & 0 & 0 \\ 0 & 0 & 1/2 & 1/2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/3 & 2/3 \\ 0 & 0 & 0 & 0 & 2/3 & 1/3 \end{pmatrix},$$

$$q = (1/2, 1/2, 0, 0, 0, 0).$$

Найти:

- несущественные состояния;
- математическое ожидание времени τ до выхода из множества несущественных состояний;
- вероятности $p_i^{(\alpha)}, p_i^{(\beta)}$ попадания в классы состояний $\alpha = \{3, 4\}, \beta = \{5, 6\}$, если начальным состоянием является $i \in \{1, 2\}$;
- предельное при $t \rightarrow \infty$ распределение по состояниям, т. е. величины $\pi_k = \lim_{t \rightarrow \infty} P(\xi_t = k)$.

Ч-5.78

Цепь Маркова ξ имеет начальное состояние $\xi_0 = 0$ и переходные вероятности

$$P\{\xi_{n+1} = k+1 | \xi_n = k\} = \frac{1}{a^k s}$$

$$P\{\xi_{n+1} = k | \xi_n = k\} = 1 - \frac{1}{a^k}$$

где $k, n \in \{0, 1, \dots\}$ и $a > 1$. Составить рекуррентные уравнения для величин $p_k^{(n)} = P\{\xi_n = k\}$ и с их помощью найти Ma^{ξ_n} и D^{ξ_n} .

Ч-5.79

Показать, что если все собственные числа $\lambda_1, \dots, \lambda_n$ матрицы $A = \|a_{ij}\|_{i,j=1}^n$ различны, то элементы матрицы $A^m = \|a_{ij}^{(m)}\|_{i,j=1}^n$ представляются в виде

$$a_{ij}^{(m)} = \sum_{h=1}^n c_{ij}^{(h)} \lambda_h^m, \quad m = 1, 2, \dots$$

Ч-5.80

Показать, что если матрица $A = \|a_{ij}\|_{i,j=1}^n$ имеет собственные числа $\lambda_1, \dots, \lambda_r$, причем кратность λ_s равна $s_i, s_1 + \dots + s_r = n$, то элементы матрицы $A^m = \|a_{ij}^{(m)}\|_{i,j=1}^n$ представляются в виде

$$a_{ij}^{(m)} = \sum_{k=1}^r \lambda_k^m \sum_{l=0}^{s_k-1} c_{ij}^{(k,l)} m^l, \quad m = 1, 2, \dots$$

Ч-5.81

Используя задачи 5.795.80, показать, что получение формулы для вероятности перехода за n шагов из состояния i в состояние j в цепи Маркова с N состояниями может быть сведено к нахождению собственных чисел матрицы $P = \|p_{ij}\|_{i,j=1}^N$ вероятностей перехода за один шаг, вычислению $P, P^2, \dots, P^{N-1}$ и решению системы линейных уравнений.

Ч-5.82

Матрица вероятностей перехода $P = \|p_{ij}\|$ цепи Маркова ξ , с состояний 1 и 2 определяется формулами

$$p_{11} = 1 - \alpha, p_{12} = \alpha, p_{21} = \beta, p_{22} = 1 - \beta.$$

Найти вероятности $p_{ij}(t)$ перехода за время t и стационарные вероятности π .

Ч-5.83

Для цепи Маркова ξ_t , рассмотренной в задаче н5.82, обозначим через $v_1(t)$ число попаданий в состояние 1 за время t . Показать, что для любого $j = 1, 2$ при $t \rightarrow \infty$

$$M\{v_1(t)|\xi_0 = j\} = \frac{\nu\beta}{\alpha + \beta}t(1 + o(1)),$$

$$D\{v_1(t)|\xi_0 = j\} = o(t^2).$$

Ч-5.89

Вероятность того, что атом радиоактивного элемента, существующий в момент t , распадется на интервале $(t, t+h)$, при $h \downarrow 0$ имеет вид $\alpha h + o(h)$ и не зависит от предыстории процесса. Найти:

- вероятность того, что время τ до распада атома будет больше t , $\{ \tau > t \} = 1/2$,
- плотность $p_n(t)$ распределения времени до распада хотя бы одного из n существующих в момент $t = 0$ и независимо ведущих себя атомов.

Ч-5.90

Пусть ξ_t - цепь Маркова с непрерывным временем и множеством состояний $\{1, 2\}$. Вероятности перехода $p_{ij}(h) = P(\xi_{t+h} = j | \xi_t = i)$ удовлетворяют условиям

$$p_{12}(h) = \alpha h + o(h), \quad p_{21}(h) = \beta h + o(h), \quad h \rightarrow 0.$$

Найти $p_{ij}(t)$, $t \geq 0$, $i, j \in \{1, 2\}$.

Ч-5.91

Пусть $p^{(\alpha, \beta)}(t) = \|p_{ij}^{(\alpha, \beta)}(t)\|$ - матрица вероятностей перехода за время t цепи Маркова ξ , описанной в задаче н5.90, а $A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}$ - стохастическая матрица. Доказать, что для любого фиксированного числа $u > 0$ цепь Маркова ξ_t удовлетворяющая условию

$$P^{(\alpha, \beta)}(u) = A,$$

существует тогда, и только тогда, когда $a_{11} + a_{22} > 1$.

Ч-5.92

Для цепи Маркова ξ_i , определенной в задаче н5.90, обозначая через $\tau_k(t)$ ($k = 1, 2$) суммарную длительность пребывания цепи в состоянии k за время t . Найти $m_i(t) = M\{\tau_1(t) | \xi_0 = i\}$ и главный член асимптотической формулы для $t \rightarrow \infty$ $b_i^2(t) = D\{\tau_1(t) | \xi_0 = i\}$ при $t \rightarrow \infty$

Ч-5.95

Цепь Маркова ξ с непрерывным временем имеет множество состояний $\{0, 1, \dots, N\}$, а вероятности $p_{ij}(h) = P\{\xi_{t+h} = j | \xi_t = i\}$ перехода из состояния i в состояние j за время h удовлетворяют при $h \downarrow 0$ условиям

$$p_{00}(h) = 1 - \alpha h + o(h), \quad p_{NN}(h) = 1 - \beta h + o(h),$$

$$p_{ii}(h) = 1 - (\alpha + \beta)h + o(h), \quad 1 \leq i \leq N-1,$$

$$p_{t-1,i}(h) = \alpha h + o(h),$$

$$p_{i,i-1}(h) = \beta h + o(h), \quad 1 \leq i \leq N,$$

$$p_{ij}(h) = o(h), \quad \text{если } |i - j| > 1,$$

а) Составить систему дифференциальных уравнений для $p_{ij}(t)$, $0 \leq i, j \leq N$, $t \geq 0$.

б) Найти $\pi_j^{(N)} = \lim_{t \rightarrow \infty} p_{ij}(t)$, $j = 0, 1, \dots, N$ как функции от $\theta = \alpha/\beta$. Показать что если $\theta = 1$, то $\pi_0^{(N)} = \dots = \pi_N^{(N)} = \frac{1}{N+1}$.

в) В случае $\theta < 1$ найти $\pi_j = \lim_{N \rightarrow \infty} \pi_j^{(N)}$, $j = 0, 1, \dots$, а в случае $\theta > 1$ найти $\pi_j^* = \lim_{N \rightarrow \infty} \pi_{N-1-j}^{(N)} = 0, 1, \dots$

landau-inst-summer-2025.6/8. Случайные блуждания

Пациенту прописали выпивать лекарства в дозировке четверть таблетки раз в день и выдали 10 таблеток в мешочке. Пациент раз в день достает из мешочка случайную таблетку. Если она целая, пациент ломает ее на 4 части, затем выпивает одну, а три четвертинки закидывает обратно в пакет. В противном случае, если он достает четвертинку, то сразу ее выпивает. Сколько в среднем останется целых таблеток через 8 дней?

Указание: задачу разрешается решать численно.

Solution

(???)

7.1.4 Problems About Markov Chains Applications

(!!!??? ну теорию доучить нужно, пока это вообще не продумал еще.)

Белан-ВШЭ-лб.3.4 Стоянка автомобилей

Автомобили приезжают на стоянку бесконечной вместимости независимо друг от друга в случайные моменты времени со средней частотой λ . Каждый автовладелец оставляет свой автомобиль на случайное время, имеющее экспоненциальную функцию плотности вероятности с параметром μ . Найдите среднее число занятых парковочных мест $\langle n(t) \rangle$ как функцию от времени t , если в нулевой момент на стоянке находилось m автомобилей.

Solution

Используя Ур. (48) и (49), можно вывести замкнутое дифференциальное уравнение на эволюцию статистического момента $\langle n(t) \rangle = \sum_{n=0}^{+\infty} n \pi_n(t)$

$$\begin{aligned} \frac{d\langle n(t) \rangle}{dt} &= \sum_{n=0}^{+\infty} n \dot{\pi}_n(t) = \sum_{n=1}^{+\infty} n [\lambda \pi_{n-1}(t) - (\lambda + n\mu) \pi_n(t) + (n+1)\mu \pi_{n+1}(t)] = (62) \\ &= \sum_{n=1}^{+\infty} [\lambda(n-1)\pi_{n-1}(t) + \lambda \pi_{n-1}(t) - \lambda n \pi_n(t) - \mu n^2 \pi_n(t) + \mu(n+1)^2 \pi_{n+1}(t) - \\ &\quad - \mu(n+1)\pi_{n+1}(t)] = \lambda \langle n(t) \rangle + \lambda - \lambda \langle n(t) \rangle - \mu \langle n^2(t) \rangle + \mu \langle n^2(t) \rangle - \mu \pi_1(t) - \\ &\quad - \mu \langle n(t) \rangle + \mu \pi_1(t) = \lambda - \mu \langle n(t) \rangle \end{aligned}$$

решение которого с начальным условием $\langle n(0) \rangle = m$ имеет вид

$$\langle n(t) \rangle = m e^{-\mu t} + \frac{\lambda}{\mu} (1 - e^{-\mu t}).$$

За время порядка μ^{-1} этот ответ релаксирует к стационарному значению $\frac{\lambda}{\mu} (1 - e^{-\mu t})$.

Белан-ВШЭ-лб.3.5 Запросы на сервер (????)

5 Запросы поступают на сервер независимо друг от друга со средней частотой λ . Лист ожидания сервера имеет конечную вместимость N и каждый новый запрос добавляется в очередь, если в момент его прибытия в листе ожидания есть свободное место. Если же лист ожидания оказывается полностью заполненным в момент поступления очередного запроса, то этот запрос отражается. Сервер обслуживает запросы по одному в порядке их поступления. Время обработки одного запроса это случайная величина, имеющая экспоненциальную функцию распределения с параметром μ .

а) Вычислите вероятность p_n обнаружить ровно n задач на сервере в статистически-стационарном режиме его работы.

б) Вычислите среднее число задач в листе ожидания.

(?? не решал)

Белан-ВШЭ-лб.3.6 Автобус и пассажиры (!!?)

Пассажиры приходят на автобусную остановку независимо друг от друга со средней частотой λ пассажиров в минуту. Прибытие автобусов это пуассоновские события со средней частотой μ автобусов в минуту. Будем предполагать, что вместимость автобуса достаточно велика, чтобы всякий раз принять всех скопившихся на остановке пассажиров. Кроме того, будем пренебрегать временем, необходимым для посадки пассажиров в автобус.

(а) Вычислить стационарную вероятность p_n найти на остановке ровно n пассажиров. Каковы условия существования статистически-стационарного режима?

(б) Вычислить среднее количество людей, ожидающих автобус на остановке, в статистически-стационарном режиме.

(в) Каково среднее число пассажиров, которые уедут с остановки за период времени длительностью T в статистически-стационарном режиме?

Solution

(?!?! решал ее, элементарно делается, потом допишу)

Белан-ВШЭ-лб.3.7 Очереди к кассам в магазине

В магазине работают 2 кассы самообслуживания. Обслуживание одного покупателя на любой из касс занимает случайный промежуток времени, имеющий экспоненциальное распределение с параметром μ . Покупатели подходят к кассам в пуассоновские моменты времени со средней частотой λ , выстраиваясь в случае необходимости в единую очередь. Очередной покупатель проходит на ту кассу, которая освобождается первой. Предполагая, что система находится в статистически стационарном режиме, найдите

(а) вероятность того, покупателю, пришедшему на кассы в случайный момент времени, не придется стоять в очереди;

(б) среднее число людей, ожидающих своей очереди.

Будем обозначать через $\pi_n(t)$ вероятность того, что в момент времени t возле касс самообслуживания находятся ровно n покупателей. Эти вероятности подчиняются следующему управляющему уравнению

$$\begin{aligned}\frac{d\pi_0(t)}{dt} &= -\lambda\pi_0(t) + \mu\pi_1(t) \\ \frac{d\pi_1(t)}{dt} &= \lambda\pi_0(t) - (\lambda + \mu)\pi_1(t) + 2\mu\pi_2(t) \\ \frac{d\pi_n(t)}{dt} &= \lambda\pi_{n-1}(t) - (\lambda + 2\mu)\pi_n(t) + 2\mu\pi_{n+1}(t), \quad \text{при } n \geq 2\end{aligned}$$

Отсюда легко находим, что в статистически стационарном режиме $\pi_n^{st} = 2 \left(\frac{\lambda}{2\mu}\right)^n \pi_0^{st}$ для $n \geq 1$. С учетом условия нормировки, $\sum_{n=0}^{+\infty} \pi_n^{st} = 1$, окончательно имеем

$$\begin{aligned}\pi_0^{st} &= \frac{1 - \frac{\lambda}{2\mu}}{1 + \frac{\lambda}{2\mu}} \\ \pi_n^{st} &= 2 \frac{1 - \frac{\lambda}{2\mu}}{1 + \frac{\lambda}{2\mu}} \left(\frac{\lambda}{2\mu}\right)^n, \quad \text{при } n \geq 1\end{aligned}$$

Видим, что стационарное решение существует при условии $\lambda < 2\mu$.

а) Вероятность того, покупателю, пришедшему на кассы в случайный момент времени, не придется стоять в очереди очевидно равна $\pi_0^{st} + \pi_1^{st}$.

б) Среднее число людей, ожидающих своей очереди равно $\sum_{n=3}^{+\infty} (n-2)\pi_n^{st}$.

Белан-ВШЭ-лб.3.8 Очереди на прививки

Пациенты приходят в пункт вакцинации без предварительной записи и независимо друг от друга со средней частотой λ . Каждый пациент, дождавшись своей очереди, сперва попадает в кабинет терапевта, где проводит случайное время, имеющее экспоненциальную функцию распределения с параметром μ_1 , а затем направляется в прививочный кабинет для прохождения процедуры вакцинации, которая длится экспоненциально распределенный с параметром μ_2 случайный промежуток времени.

(а) При каких условиях в такой системе существует статистически стационарное состояние? Предполагая, что система находится в статистически стационарном режиме, определите (б) среднее число пациентов, ожидающих своей очереди в кабинет терапевта;

- (в) среднее число пациентов, ожидающих своей очереди в прививочный кабинет.
 (г) вероятность p_k обнаружить в пункте вакцинации ровно k человек, прийдя туда в случайный момент времени.

Solution

(крутая задача, но не смотрел пока что.)

а) Описанная в условии задачи система массового обслуживания представляет собой тандем из двух последовательных М/М/1 систем. Очевидно, что происходящее во второй системе не оказывает влияния на стохастическую динамику первой системы. Поэтому статистически-стационарная частная функция распределения вероятностей числа пациентов в первой системе имеет вид

$$\pi_n^{(1)} = \left(1 - \frac{\lambda}{\mu_1}\right) \left(\frac{\lambda}{\mu_1}\right)^n$$

а само условие существования стационара дается неравенством $\lambda < \mu_1$. Далее, очевидно, что в статистически стационарном режиме, интенсивность потока пациентов, прибывающих в пункт вакцинации, равна интенсивности потока пациентов, покидающих кабинет терапевта. Словом, на вход второй системы тандема подается поток интенсивности λ . Значит частная функция распределения вероятностей числа пациентов во второй системе дается выражением

$$\pi_m^{(2)} = \left(1 - \frac{\lambda}{\mu_2}\right) \left(\frac{\lambda}{\mu_2}\right)^m$$

при условии, что $\lambda < \mu_2$. Итого получаем, что статистический стационар имеет место, если $\lambda < \min(\mu_1, \mu_2)$. б) Среднее число пациентов, ожидающих своей очереди в кабинет терапевта, равно

$$\sum_{n=1}^{\infty} (n-1) \pi_n^{(1)} = -1 + \pi_0^{(1)} + \left(1 - \frac{\lambda}{\mu_1}\right) \sum_{n=0}^{+\infty} n \left(\frac{\lambda}{\mu_1}\right)^n = \frac{\lambda^2}{\mu_1(\mu_1 - \lambda)}$$

(в) Среднее число пациентов, ожидающих своей очереди в прививочный кабинет: $\frac{\lambda^2}{\mu_2(\mu_2 - \lambda)}$

г) Вероятность p_k обнаружить в пункте вакцинации ровно k человек, прийдя туда в случайный момент времени, может быть записана как

$$p_k = \Pr[n + m = k] = \sum_{n=0}^k \pi_{n, k-n}^{st}$$

где $\pi_{n,m}^{st}$ - совместная функция распределения количеств пациентов в обеих системах. Последняя факторизуется в произведение частных вероятностей

$$\pi_{n,m}^{st} = \pi_n^{(1)} \pi_m^{(2)} = \left(1 - \frac{\lambda}{\mu_1}\right) \left(\frac{\lambda}{\mu_1}\right)^n \left(1 - \frac{\lambda}{\mu_2}\right) \left(\frac{\lambda}{\mu_2}\right)^m$$

в чем легко убедиться подстановкой этого анзаца в систему управляющих уравнений тандема

$$\begin{aligned} \dot{\pi}_{0,0} &= -\lambda \pi_{0,0} + \mu_2 \pi_{0,1} \\ \dot{\pi}_{n,0} &= \lambda \pi_{n-1,0} - (\lambda + \mu_1) \pi_{n,0} + \mu_2 \pi_{n,1}, \quad \text{при } n \geq 1 \\ \dot{\pi}_{0,m} &= \mu_2 \pi_{0,m+1} - (\lambda + \mu_2) \pi_{0,m} + \mu_1 \pi_{1,m-1}, \quad \text{при } m \geq 1, \\ \dot{\pi}_{n,m} &= -(\lambda + \mu_1 + \mu_2) \pi_{n,m} + \mu_1 \pi_{n+1,m-1} + \mu_2 \pi_{n,m+1} + \lambda \pi_{n-1,m}, \quad \text{при } n, m \geq 1 \end{aligned} \quad (78)$$

(780)

Значит

$$p_k = \left(1 - \frac{\lambda}{\mu_1}\right) \left(1 - \frac{\lambda}{\mu_2}\right) \left(\frac{\lambda}{\mu_2}\right)^k \sum_{n=0}^k \left(\frac{\mu_2}{\mu_1}\right)^n = \left(1 - \frac{\lambda}{\mu_1}\right) \left(1 - \frac{\lambda}{\mu_2}\right) \left(\frac{\lambda}{\mu_2}\right)^k \frac{\left(\frac{\mu_2}{\mu_1}\right)^{k+1} - 1}{\frac{\mu_2}{\mu_1} - 1}$$

7.2 Problems of Elementary Mathematical Statistics (????)

(ничего пока не шарю, нужно потом когда-то разобраться.)
 (пока мало что тут есть.)

7.2.1 Problems About the Parameters of Mathematical Statistics

Ч-6.1

Пусть $x_1, x_2, \dots, x_n$ - случайная выборка с $x_k = a$, $Dx_k = \sigma^2$, $M(x_k - a)^4 < \infty$, $k = 1, \dots, n$. Найти математическое ожидание величины

$$s^2 = \frac{1}{n-1} \sum_{k=1}^n (x_k - \bar{x})^2 \quad \left(\bar{x} = \frac{1}{n} \sum_{k=1}^n x_k \right).$$

Является ли s^2 состоятельной оценкой σ^2 ?

Ч-6.2

Найти математическое ожидание и дисперсию эмпирического момента $m_r = \frac{x_1^r + \dots + x_n^r}{n}$ независимой выборки, соответствующей случайной величине ξ с $M\xi^k = a_k$, $1 \leq k \leq 2r$

Ч-6.3

По неоднородной выборке $x_1, \dots, x_{n_1}$ где $x_k, k = 1, \dots, n$, независимы, $x_n = a$, $Dx_k = \sigma_k^2$ (σ_k известны), найти несмещенную линейную относительно x_n ($k = 1, \dots, n$) оценку а параметра a , которая имеет наименьшую возможную дисперсию.

Ч-6.4

*. Пусть $x_{i1}, \dots, x_{in_i}$ ($i = 1, 2, \dots, I$) - независимые нормально распределенные величины с параметрами (a, σ_1^2)

$$\bar{x}_i = \frac{1}{n_i} \sum_{k=1}^{n_i} x_{ik_i} \quad s_i^2 = \frac{1}{n_i - 1} \sum_{k=1}^{n_i} (x_{ik} - \bar{x}_i)^2.$$

Является ли оценка

$$\bar{x} = \sum_{i=1}^I c_i \bar{x}_{iv} \quad c_i = \frac{s_i^2}{s_1^2 + \dots + s_I^2}$$

несмещенной оценкой параметра a ?

Ч-6.5

Пусть $x_1, \dots, x_n$ - независимые одинаково распределенные случайные величины с $Dx_1 > 0$, $Mx_1^4 < \infty$. Положим

$$\bar{x} = \frac{1}{n} \sum_{k=1}^n x_{k\epsilon} \quad s^2 = \frac{1}{n-1} \sum_{k=1}^n (x_k - \bar{x})^2.$$

Найти

$$\lim_{n \rightarrow \infty} P\left(\frac{\bar{x} - Mx_1}{s/\sqrt{n}} \leq x\right)$$

Ч-6.6

Пусть $(x_1, y_1), \dots, (x_n, y_n)$ - независимая выборка, соответствующая случайному вектору. (ξ, η) , Т. θ_0 $P(x_i \leq x, y_1 \leq y) = P(\xi \leq x, \eta \leq y)$. Показать, что величина

$$m = \frac{1}{n-1} \sum_{k=1}^n (x_k - \bar{x})(y_k - \bar{y})_i$$

где

$$\bar{x} = \frac{1}{n} \sum_{k=1}^n x_{k1} \quad \bar{y} = \frac{1}{n} \sum_{k=1}^n y_{ks}$$

является несмещенной и состоятельной оценкой $\text{cov}(\xi, \eta)$.

Ч-6.7

Пусть μ_n - число успехов в n испытаниях схемы Бернулли с вероятностью успеха p в каждом испытании. Найти оценку наибольшего правдоподобия p параметра p . Доказать ее несмещенность и состоятельность.

Ч-6.8

Пусть μ_n - число успехов в n испытаниях схемы Бернулли с вероятностью успеха p в каждом испытании. Построить для p асимптотически доверительный интервал с доверительной вероятностью $1 - 2\alpha$.

Ч-6.9

Используя таблицу случайных чисел, получить реализацию выборки $x_1, \dots, x_n$ где x_k равномерно распределены; $n = 50$). Найти:

а) вариационный ряд $x_{(1)} \leq x_{(2)} \leq \dots \leq x_{(n)}$;
 б) эмпирическую функцию распределения (построить ее график и график теоретической функции распределения);

в) $\bar{x} = \frac{1}{n} (x_1 + \dots + x_n)$ (сравнить с x_k); г) $s^2 = \frac{1}{n-1} \sum_{k=1}^n (x_k - \bar{x})^2$ (сравнить с $D x_k$).

Ч-6.10

Используя таблицу нормально распределенных случайных чисел, получить реализацию выборки $x_1, \dots, x_n$, где x_k имеет нормальное распределение с параметрами $a = Mx_k = 0,5, \sigma^2 = Dx_k = 1; n = 50$. Найти: а) вариационный ряд, б) эмпирическую функцию распределения, в) $\bar{x}, r, s^2$ (см. задачу 6.9).

Ч-6.11

По выборке, полученной в задаче 6.10, построить доверительный интервал для a (считая a и b неизвестными) с доверительной вероятностью 0,95.

Ч-6.12

Используя метод наибольшего правдоподобия, найти по выборке $x_1, \dots, x_n$, где $P(x_k = m) = \frac{\lambda^m}{m!} e^{-\lambda}, m = 0, 1, \dots$, оценку λ параметра λ . Будет ли эта оценка несмещенной и состоятельной? Найти $M\lambda^*, D\lambda^*$.

Ч-6.13

Пусть $x_1, x_2, \dots, x_n$ - выборка, соответствующая показательному распределению с параметром λ . Найти оценку максимального правдоподобия λ для λ . Вычислить $M\frac{1}{\lambda^*}, D\frac{1}{\lambda^*}$.

Ч-6.14

Для оценки параметров a, b, c имеются три независимые выборки $a_1, \dots, a_n; b_1, \dots, b_n; c_1, \dots, c_n$. Известно, что $c = a + b$ и величины a_i, b_i, c_i распределены нормально с $M a_i = a, M b_i = b, M c_i = c$. Дисперсии $D a_i = \sigma_a^2, D b_i = \sigma_b^2, D c_i = \sigma_c^2$ известны. Найти:

а) оценки наибольшего правдоподобия a^*, b^* , с параметров a, b, c используя для каждого параметра только соответствующую ему выборку, а также найти $a^*, M b, M c^*, D a^*, D b^*, D c^*$;

б) оценки наибольшего правдоподобия a^{**}, b^{**}, c^{**} , используя сразу три выборки и связь $c = a + b$, и также найти $M a^{**}, M b^{**}, M c^{**}, D a^{**}, D b^{**}, D c^{**}$.

Ч-6.15

Пусть $y^{(k)} = (y_1^{(k)}, y_2^{(k)})$, $k = 1, \dots, n$, - независимые нормально распределенные случайные векторы, $M y_i^{(k)} = a_i, D y_i^{(k)} = \sigma_i^2, i = 1, 2, \text{cov}(y_1^{(k)}, y_2^{(k)}) = \rho \sigma_1 \sigma_2$. Параметры σ_1, σ_2, ρ известны. Найти:

а) оценку максимального правдоподобия (a_1^*, a_2^*) параметров (a_1, a_2) по выборке $y^{(1)}, \dots, y^{(n)}$;

б) оценку максимального правдоподобия a_i^{**} параметра a_i по выборке $y_i^{(1)}, \dots, y_i^{(n)}$; в) $M a_i^*, D a_i^*, M a_i^{**}, D a_i^{**}$.

Ч-6.16

*. Пусть $y^{(k)} = (y_1^{(k)}, y_2^{(k)})$, $k = 1, \dots, n$, - независимые нормально распределенные случайные векторы, $My_1^{(k)} = a_1$, $My_2^{(k)} = a_2$, $D_i^{(k)} = 1$ ($i = 1, 2$), $\text{cov}(y_1^{(k)}, y_2^{(k)}) = \rho_k$ ($k = 1, \dots, n$). Параметры ρ_k известны. Найти:

- в) оценки максимального правдоподобия a_1^*, a_2 параметров a_1, a_2 по выборке $y^{(1)}, \dots, y^{(n)}$;
 б) оценку максимального правдоподобия a_1^{**} параметра a_1 по выборке $y_1^{(1)}, \dots, y_1^{(n)}$

Ч-6.19

Для величины $A = \alpha + \beta a + \gamma b$ получены оценки

$$A_1^* = \alpha + \beta z_1 + \gamma z_3, \quad A_2^* = \alpha + \beta z_2 + \gamma z_3.$$

где α, β, γ - известные постоянные, z_1, z_2, z_3 - независимые оценки неизвестных параметров: $Mz_3 = b$, $Mz_1 = Mz_2 = a$; $Dz_i = \sigma_i^2$, $i = 1, 2, 3$. Подобрать постоянные c_1, c_2 так, чтобы оценка $A^* = c_1 A_1^* + c_2 A_2^*$ была несмещенной и имела среди несмещенных оценок наименьшую дисперсию.

Ч-6.20

Пусть z_1, z_2, z_3 - несмещенные оценки параметра a ; $Dz_i = 1$ ($i = 1, 2, 3$), $\text{cov}(z_1, z_2) = \rho$, $\text{cov}(z_i, z_3) = 0$ ($i = 1, 2$). Найти несмещенную линейную оценку $z = c_1 z_1 + c_2 z_2 + c_3 z_3$ параметра a с наименьшей возможной дисперсией и дисперсию этой оценки. Рассмотреть случаи: а) $|\rho| < 1$, б) $\rho = -1$, в) $\rho = 1$.

Ч-6.21

Функция $y = Ax$ измерена в точках $x_1, \dots, x_n$. Пусть результаты измерений являются реализацией независимых случайных величин $\tilde{y}_1, \tilde{y}_2, \dots, \tilde{y}_n$, у которых $M\tilde{y}_i = y_i = Ax_i$, $D\tilde{y}_i = \sigma^2 < \infty$ ($i = 1, 2, \dots, n$).

Найти: сумму квадратов, т. е. минимизируя по A выражение $I(A) = \sum_{i=1}^n (\tilde{y}_i - Ax_i)^2$

б) MA_n^*, DA_n^* Доказать, что оценка A_n состоятельна, если $X_n^2 = \sum_{i=1}^n x_i^2 \rightarrow \infty$ при $n \rightarrow \infty$. Пусть результаты измерений являются реализацией n независимых двумерных случайных векторов $(\tilde{x}_i, \tilde{y}_i)$. ($i = 1, \dots, n$), у которых координаты

6.2*. Проведено n измерений значений функций $y =$

измерений являются реализацией n независимых двумерных случайных векторов $(\tilde{x}_i, \tilde{y}_i)$. ($i = 1, \dots, n$), у которых координаты

независимы, $M\tilde{x}_i = x_i$, $M\tilde{y}_i = y_i = Ax_i$, $D\tilde{x}_i = \sigma^2$, $D\tilde{y}_i = \sigma^2$ Найти:

а) оценку A_n параметра A , минимизируя по A и $x_1, \dots, x_n$ выражение

$$I(A, x_1, \dots, x_n) = \sum_{i=1}^n (\tilde{y}_i - Ax_i)^2 + \sum_{i=1}^n (\tilde{x}_i - x_i)^2.$$

б) Показать, что оценка A_n состоятельна, если при

$$X_n^2 = \sum_{i=1}^n x_i^2 \rightarrow \infty, \quad n/X_n^2 \rightarrow 0.$$

Ч-6.23

Функция $y = Ax$ измерена n_1 раз в точке $x_1, \dots, x_{n_1}$ раз в точке x_k . Пусть результаты измерений y_{ij} ($j = 1, \dots, n_i, i = 1, \dots, k$) некоррелированы и имеют вид

$$y_{ij} = Ax_i + \delta_{ij},$$

где $M\delta_{ij} = 0$, $D\delta_{ij} = \sigma^2$.

а) Найти оценку A параметра A , используя метод наименьших квадратов, т. е. минимизируя по A выражение

$$I(A) = \sum_{i=1}^k \sum_{j=1}^{n_i} (y_{ij} - Ax_i)^2.$$

б) Найти MA и DA .

Ч-6.24

В предыдущей задаче обозначим

$$\bar{y}_i = \frac{1}{n_i} \sum_{j=1}^{n_i} y_{ij}, \quad i = 1, \dots, k.$$

Подобрать $c_1, \dots, c_k$ так, чтобы оценка

$$A^* = c_1 \bar{y}_1 + c_2 \bar{y}_2 + \dots + c_k \bar{y}_k$$

была несмещенной и имела наименьшую дисперсию. Найти DA^* при наилучшем выборе $c_1, c_2, \dots, c_k$.

Ч-6.25

Из урны, содержащей N белых и черных шаров, производится выборка объема n с возвращением. Пусть μ_n - число белых шаров в выборке, а M - неизвестное число белых шаров в урне. Для оценки величины $p = M/N$ используется статистика $p_n^* = \mu_n/n$. Найти Mr_n, Dp_n .

Ч-6.26

Из урны, содержащей N белых и черных шаров, производится выборка объема n без возвращения. Пусть μ_n - число белых шаров в выборке, а M - неизвестное начальное число белых шаров в урне. Для оценки величины $p = M/N$ используется статистика $p_n^{**} = \mu_n/n$. Найти Mr_n^{**}, Dp_n^{**} .

Ч-6.27

Для сравнения точности оценок p_n^*, p_n^{**} , определенных в задачах 6.25, 6.26, найти $\lim_{n \rightarrow \infty} (Dp_n^{**}/Dp_n^*)$ в следующих случаях: а) $n \rightarrow \infty, \frac{n}{N} \rightarrow \gamma$ ($0 < \gamma < \infty$); б) $n \rightarrow \infty, n/N \rightarrow 0$.

Ч-6.28

Из урны, содержащей неизвестное число шаров N (шары завуамерены), производится выборка объема n с возвращением. Для оценки числа N используется величина $1/\eta_n$, где

$$\eta_n = \frac{1}{n(n-1)} \sum_{k=1}^N \xi_k (\xi_k - 1),$$

ξ_k - число появлений шара с номером k в выборке. Найти $M\eta_n$ и асимптотическую формулу для $D\eta_n$ при $n, N \rightarrow \infty, n/N \rightarrow \alpha \in (0, \infty)$.

Ч-6.33

Чтобы оценить ширину кольца, образованного двумя окружностями с общим центром, измерялись их радиусы R и r , т. е. была получена выборка $y_1, \dots, y_n, x_1, \dots, x_n$, образованная независимыми случайными величинами, которые имеют нормальные распределения с $My_i = R, \quad Mx_i = r \quad (R > r), \quad Dx_i = Dy_i = \sigma^2, i = 1, 2, \dots, n$. Пусть $\bar{x} = (x_1 + \dots + x_n)/n$ и $\bar{y} = (y_1 + \dots + y_n)/n$. Для оценок $\bar{a} = \bar{y} - \bar{x}, a^* = \max(0, \bar{y} - \bar{x})$ ширины кольца $a = R - r$ найти: $\bar{a}, Ma^*, \sqrt{M(\bar{a} - a)^2}, \sqrt{M(a^* - a)^2}$. Вычислить эти величины при $R = 1001$ м., $r = 1000$ м., $\sigma = 10$ м., $n = 200$.

Ч-6.34

Независимые наблюдения $\theta_1^*, \dots, \theta_n$ имеют не известные математические ожидания $\theta_i^* = \theta_i$ и известные дисперсии $D\theta_i^* = \sigma_i^2, \quad i = 1, \dots, n$. Для оценки линейной комбинации $I = c_1\theta_1 + \dots + c_n\theta_n$ с заданными $c_1, \dots, c_n$ используются статистики

$$I_N = c_1\theta_1^* + \dots + c_N\theta_N^* \quad (1 \leq N \leq n).$$

а) Доказать, что I_n - несмещенная, а I_N при $N < n$ - смещенная оценка I .

б) Найти $M(I_N - I)^2$. При каких условиях среднеквадратическое отклонение смещенной оценки I_{n-1} меньше среднеквадратического отклонения несмещенной оценки I_n ?

Ч-6.35

Используя критерий χ^2 , проверить гипотезу о том, что выборка, полученная в задаче нб.9, соответствует равномерному распределению на отрезке $[0, 1]$. уровень значимости $\alpha = 0,05$.

Ч-6.36

Используя критерий χ^2 , проверить гипотезу о том, что выборка, полученная в задаче нб.10н, соответствует нормальному распределению; параметры a и σ считать неизвестными. Уровень значимости положить равным 0,05. Соотношения для различения по выборке $x_1, \dots, x_n$ гипотез $H_1 : x_h$ распределены нормально с параметрами (a_1, σ^2) , $H_2 : x_k$ распределены нормально с параметрами (a_2, σ^2) .

Ч-6.40

*. Пусть $x_1, \dots, x_n$ - выборка. Гипотеза H_1 состоит в том, что x_i равномерно распределены на интервале $(0, 1)$ и независимы, а гипотеза H_2 - в том, что x_i независимы и имеют непрерывно дифференцируемую плотность распределения $g(x); g(x) \neq 1$ при $0 \leq x \leq 1$, $g(x) = 0$ при $x \notin [0, 1]$. Разобьем $[0, 1]$ на N полуинтервалов $[0, \frac{1}{N}), [\frac{1}{N}, \frac{2}{N}), \dots, [\frac{N-1}{N}, 1)$ и обозначим через μ_a число полуинтервалов, в которые не попало ни одно из наблюдений x_i . При гипотезах H_1 и H_2 найти $a_j = \lim (M_0/N), j = 1, 2$, когда $n/N \rightarrow \gamma \in (0, \infty), n, N \rightarrow \infty$.

Ч-6.41

Пусть $0 < \tau_1 < \tau_2 < \dots$ положения точек пуассоновского потока с неизвестной интенсивностью λ . Найти: а) оценку максимального правдоподобия λ_n параметра λ , построенную по $\tau_1, \dots, \tau_n$; б) $M\lambda_n^*, D\lambda_n$.

Ч-6.42

Пусть $\tau_{12}(t)$ - число переходов в цепи Маркова, определенной в задаче нб.82, из состояния 1 в состояние 2 за время t .

а) Найти $\lim_{t \rightarrow \infty} M \frac{\tau_{12}(t)}{t} \times \lim_{t \rightarrow \infty} D \frac{\tau_{12}(t)}{t}$.

б) Является ли величина $\tau_{12}(t)/t$ состоятельной оценкой параметра $\gamma = \frac{\alpha\beta}{\alpha+\beta}$ при $t \rightarrow \infty$?

7.3 Problems Arising from the Application of Probability Theory

(пока избры, но вообще тут много можно интересного порешать)
(наверное, крутые очень задачи, просто не до них абсолютно. доработаю базу - буду тренироваться, если захочу.)

7.3.1 Problems of Probability Theory in Physics

(тоже такая часть каталога.)

Избры-1 Теорема Вика

На лекции было получено, что

$$\left\langle \exp \left(-ip \frac{\xi}{N} \right) \right\rangle \approx \exp \left(-\frac{p^2 \langle \xi^2 \rangle}{2N^2} \right)$$

1. Для начала найдите моменты случайной величины, т.е. $\langle \xi^n \rangle = M_n$ 2. Рассмотрите усреднение в левой части выражения для случайной величины ξ , распределенной по Гауссу с нулевым средним.

Избры-2

Пусть случайная величина ξ нормально распределена с нулевым матожиданием и дисперсией $\langle \xi^2 \rangle = \sigma^2$. Как распределена случайная величина ξ^2 ? А как распределена ξ^n ?

Избры-3

В этой задаче попробуем проверить корректность нашего понимания вероятности. Пусть мы измеряем величину ξ и провели N измерений, которые дали нам результаты $\{\xi_1, \dots, \xi_N\}$. По этим результатам мы построили плотность вероятности: $\rho_e(\xi) = \frac{1}{N} \sum_{i=1}^N \delta(\xi - \xi_i)$. Мы будем считать, что все ξ_i - независимые случайные величины с плотностью вероятности $\rho(\xi)$. Мы будем обозначать чертой - усреднение по ρ_e т.е. $\overline{f(\xi)} = \int_{-\infty}^{\infty} f(\xi) \rho_e(\xi) d\xi$. Треугольными скобками будем обозначать усреднение по реализации, то есть $\langle \delta(\xi - \xi_i) \rangle = \rho(\xi)$. - Покажите, что функция, полученная с помощью эксперимента в среднем совпадает с истинной функцией распределения: $\langle \rho_e(\xi) \rangle = \rho(\xi)$. Таким образом $\langle \overline{f(\xi)} \rangle = \langle f(\xi) \rangle$ - Вычислим теперь "дисперсию" для экспериментальной функции распределения:

$$\langle \langle \rho_e(\xi') \rho_e(\xi'') \rangle \rangle = \langle \rho_e(\xi') \rho_e(\xi'') \rangle - \langle \rho_e(\xi') \rangle \langle \rho_e(\xi'') \rangle$$

- Предыдущий пункт помогает нам понять во сколько в среднем мы отклоняемся при вычислении различных экспериментальных величин. Выразите следующую величину через $\langle f^2(\xi) \rangle$:

$$\langle \{ \overline{f(\xi)} \}^2 \rangle - \overline{\langle f(\xi) \rangle^2}$$

Избры-4

В этой задаче нам предстоит разобраться с условиями применимости ЦПТ. Нас будет интересовать случайная величина $\xi = \frac{1}{N} \sum_{i=1}^N \xi_i$. Все ξ_i - независимые случайные величины. Нас интересовала производящая функция этой величины:

$$f(p) = \langle e^{-ip\xi} \rangle = \left\langle \prod_i e^{-i\frac{p}{N}\xi_i} \right\rangle$$

Ключевым местом в доказательстве было предположение о малости величины в экспоненте. Попробуем понять применимость этого приближения. Для этого разложим величину $\langle e^{-ip\xi} \rangle$ до 4-го порядка по $1/N$ и выполним усреднение. Наша цель - представить результат этого усреднения в виде $f(p) = \exp(G(-ip/N))$. Где $G(x) = \sum_k \Delta_k \frac{x^k}{k!}$

- Выразите первые 4 коэффициента Δ_k через моменты величины $\xi - M_l = \langle \xi^l \rangle$
- Получите условие на то, что плотность вероятности можно восстановить используя метод перевала
- Воспользуйтесь этим условием для задачи, разобранный на лекции и получите правильное условие применимости ЦПТ

8 Problems from Applications

8.1 Problems From Applications in Physics

8.1.1 Problems From Applications in Statistical Physics

Б-3.1.10 Two Level Systems

(!!! write here the combinatorical part, question exactly about it. remind the formulas for the probability!)
Термодинамическая система представляет собой ансамбль из $N \gg 1$ невзаимодействующих двухуровневых систем с одинаковой энергией возбуждения, равной ε . Кратность основного уровня энергии равна g_1 , а возбужденного уровня - g_2 . Найти энтропию системы:

- а) с помощью статистического определения энтропии;
- б) через распределение Гиббса.

Определить теплоемкость системы и показать, что область отрицательных температур отвечает термодинамически неустойчивым состояниям системы с инверсной заселенностью.

Состояние аморфных веществ (стекол) при низких температурах описывают, как состояние с макроскопически большим набором различных двухуровневых систем, у которых энергия возбуждения ε представляет случайную величину с плотностью распределения $\rho(\varepsilon)$, нормированной на число

двухуровневых систем в единице объема $\int_0^\infty \rho(\varepsilon) d\varepsilon = N$. Предполагая, что плотность распределения двухуровневых систем $\rho(\varepsilon)$ остается конечной и непрерывной при $\varepsilon = 0$, найти поведение теплоемкости аморфных веществ в пределе низких температур.

Solution a)

Суть решения в том, что

Если уровни энергии не вырождены, то число состояний с энергией $E = n\varepsilon$ есть число сочетаний из N по n :

$$\Gamma(E = n\varepsilon) = C_N^n = \frac{N!}{(N-n)!n!}.$$

Если нижний уровень g_1 -кратно вырожденный, а верхний g_2 -кратно вырожденный, то число состояний с энергией $E = n\varepsilon$ будет равно

$$\Gamma(E = n\varepsilon) = g_1^{N-n} g_2^n C_N^n.$$

Соответственно, энтропия системы дается выражением

$$S(E) = \ln \Gamma(E = n\varepsilon) = (N-n) \ln g_1 + n \ln g_2 + \ln \frac{N!}{(N-n)!n!}.$$

Part IV

Exotic Topics of Probability Theory in a Nutshell

9 Exotic Theoretical Topics in a Nutshell

9.0.1 Quantum Probability Theory (???)

(???? write some main ideas...)

10 Applications in a Nutshell

10.1 Applications in Mathematics

(I'll write later)

10.2 Applications in Physics

10.2.1 Quantum Mechanics

(мб Амосова для этого почитаю)

10.2.2 Thermodynamics

(тут всё вроде не так сложно)

10.2.3 QFT

(это затрагивали, но мне не зашло, ибо там всё сложно. функция грина, всё с ней связанное)

10.2.4 Physical Kinetics

(тут больше про случпроцы, в них сперва эту связь разберу, потом здесь укажу её)

Part V

Other Topics

11 Random Processes: Basics

Приведем основы описания стохастических систем и покажем, как можно добавлять случайные процессы в любые физические модели.

11.1 Basics of the Theory of Random Processes in a Nutshell

Разберем вкратце переход теории вероятности в теорию случайных процессов.
(по Жуковскому мфти)

11.1.1 Typical Random Processes

(где-то уже писал ведь их, наверное, в случпроцах, там разберу, вставлю сюда)

ветвящиеся процессы Гальтона-Ватсона

В настоящей главе, а также в третьей и четвертых главах мы поговорим о важнейших примерах случайных процессов.

В этой главе речь пойдет о ветвящихся процессах.

Модель ветвящихся процессов была впервые предложена Ф. Гальтоном и Г. Ватсоном в 1873 году в связи с анализом вырождения аристократических фамилий Великобритании.

Физическая модель: в дискретные моменты времени частицы распадаются на случайное количество таких же частиц. Число потомков каждой частицы имеет одно и то же распределение.

Формальная математическая модель выглядит следующим образом.

Математическая модель: пусть ξ - случайная величина со значениями в $\mathbb{Z}_+$, а $\{\xi_k^{(n)}, k, n \in \mathbb{N}\}$ - набор независимых случайных величин с тем же распределением, что и ξ . Положим

$$X_0 = 1, X_n = \sum_{k=1}^{X_{n-1}} \xi_k^{(n)}, n \geq 1$$

Определение

Процесс $(X_n, n \in \mathbb{Z}_+)$ называется ветвящимся процессом Гальтона-Ватсона с законом размножения частиц ξ .

Интерпретация величин модели совершенно ясна: X_n - это

число частиц в n -м поколении, а $\xi_k^{(n)}$ - число потомков k -й частицы из $(n-1)$ -го поколения.

Нас будут интересовать два вопроса относительно ветвящихся процессов. - Какова вероятность вырождения ветвящегося процесса $\mathbf{P}(\exists n : X_n = 0)$? - Каково распределение общего числа частиц в процессе? Ответы на эти вопросы можно найти с помощью метода производящих функций.

А Определение 2.2. Пусть ξ - случайная величина. Тогда ее производящей функцией называется функция

$$\varphi_\xi(z) = \mathbf{M}z^\xi, z \in \mathbb{R}$$

определенная для тех z , для которых конечно данное математическое ожидание. Выделим основные свойства производящих функций.

1. $\varphi_\xi(1) = 1$
2. $\varphi'_\xi(1) = \mathbf{M}\xi$
3. Если ξ и η независимы, то $\varphi_{\xi+\eta}(z) = \varphi_\xi(z) \cdot \varphi_\eta(z)$.

Если ξ принимает значения в $\mathbb{Z}_+$, то имеются дополнительные свойства:

4. $\varphi_\xi(z) = \sum_{k=0}^{\infty} z^k \mathbf{P}(\xi = k)$ - степенной ряд, сходящийся абсолютно и равномерно в области $\{|z| \leq 1\}$
5. $\varphi_\xi(z)$ непрерывно дифференцируема бесконечное число раз в области $\{|z| < 1\}$
6. $\varphi_\xi(0) = \mathbf{P}(\xi = 0)$
7. $\mathbf{P}(\xi = k) = \frac{1}{k!} \left(\frac{d^k}{dz^k} \varphi_\xi(z) \right) \Big|_{z=0}$

Пусть теперь $X = (X_n, n \in \mathbb{N})$ - это ветвящийся процесс Гальтона-Ватсона с законом размножения частиц ξ .

Лемма

2.1. Между производящими функциями элементов процесса X имеют место следующие рекуррентные соотношения:

1. $\varphi_{X_{n+1}}(z) = \varphi_{X_n}(\varphi_\xi(z))$

2. $\varphi_{X_n}(z) = \underbrace{\varphi_\xi(\varphi_\xi(\dots \varphi_\xi(z)))}_{n \text{ раз}};$

3. $\varphi_{X_{n+1}}(z) = \varphi_\xi(\varphi_{X_n}(z)).$

Данная лемма позволяет найти уравнение для вероятности вырождения. Обозначим через q вероятность вырождения, т.е.

$$q = \mathbf{P}(\text{процесс выродился}) = \mathbf{P}(\exists n : X_n = 0)$$

Тогда выполнено следующее утверждение. Лемма 2.2. Вероятность вырождения q ветвящегося процесса с законом размножения частиц ξ является решением уравнения

$$z = \varphi_\xi(z)$$

Уравнение (2) - это первая замечательная формула в теории ветвящихся процессов. Отметим, что

$$z = 1$$

всегда является решением (2). Однако при наличии других корней нам надо уметь отбирать среди них вероятность вырождения. Полную классификацию ситуаций дает

Теорема

о вероятности вырождения.

Теорема 2.1 (0 Вероятности вырождения).

Пусть случайная величина ξ не равна единице тождественно, т.е. $\mathbf{P}(\xi = 1) \neq 1$. Обозначим $\mu = \mathbf{M}\xi$. Если $\mu \leq 1$, то уравнение (2) имеет на $[0, 1]$ только одно решение: $z = 1$. В этом случае $q = 1$. Если $\mu > 1$, то уравнение (2) имеет единственное решение $z_0 \in [0, 1)$. В этом случае $q = z_0$.

Вывод: вероятность вырождения - это наименьший корень уравнения $z = \varphi_\xi(z)$ из отрезка $[0, 1]$.

Отметим, что

Теорема

имеет весьма жизненную интерпретацию: если среднее число потомков в процессе не больше единицы, то процесс обречен на вымирание; в противном случае есть ненулевая вероятность бесконечного существования популяции.

Общее число частиц. Обсудим теперь вопрос об общем

16

числе частиц в ветвящемся процессе. Обозначим общее число частиц в процессе до момента времени n включительно $Y_n = 1 + X_1 + \dots + X_n$. Тогда имеет место следующее рекуррентное соотношение между производящими функциями Y_n и Y_{n-1} , аналогичное соотношениям в лемме 2.1 для самого X_n :

$$\varphi_{Y_{n+1}}(z) = z\varphi_\xi(\varphi_{Y_n}(z))$$

Далее, для каждого $z \in [0, 1)$ существует предел

$$\rho(z) = \lim_{n \rightarrow +\infty} \varphi_{Y_n}(z)$$

который можно интерпретировать как производящую функцию общего числа частиц за все время существования процесса. Пусть Y - общее число частиц в ветвящемся процессе за все время (если процесс не выродился, то, конечно, $Y = +\infty$). Тогда

$$\rho(z) = \mathbf{M}(z^Y I(Y < +\infty)) = \sum_{k=0}^{+\infty} z^k \mathbf{P}(Y = k)$$

Найти $\rho(z)$ можно из второго замечательного уравнения в теории ветвящихся случайных процессов. Лемма 2.3. Функция $\rho(z)$ является решением уравнения

$$\rho(z) = z\varphi_\xi(\rho(z))$$

Понятие времени корреляции

Обозначим случайную бинарную переменную, кодирующую информацию о погодных условиях в n -й день наблюдений, через S_n . Здесь и далее мы будем считать, что случайный процесс S_n является статистически стационарным, то есть частные функции распределения случайных величин S_{n_1} и S_{n_2} одинаковы для любых моментов времени n_1 и n_2 . Можно предложить несколько формальных способов детектирования присутствия корреляций во временном ряду значений такого случайного процесса. - Корреляции в строке значений $s_1, s_2, s_3, \dots$ проявляются, к примеру, в том, что условная вероятность $\Pr[S_{n_2} = s_{n_2} | S_{n_1} = s_{n_1}]$, где $n_2 \geq n_1$, не равна частной вероятности $\Pr[S_2 = s_2]$. Другим способом увидеть статистическую зависимость значений случайной функции S_n в разные моменты времени является основан на анализе автокорреляционной функции, которая определена как

$$C(n_1, n_2) = \langle S_{n_1} S_{n_2} \rangle$$

Очевидно, что при наличии корреляций справедливо $C(n_1, n_2) \neq \langle S_{n_1} \rangle \langle S_{n_2} \rangle = \langle S_n^2 \rangle$. Наконец, вывод о присутствии статистических корреляций можно сделать из факта отличия от нуля взаимной информации $I(S_{n_1}, S_{n_2}) = \left\langle \log_2 \frac{P(S_{n_1}, S_{n_2})}{P_1(S_{n_1})P_2(S_{n_2})} \right\rangle$. Корреляции имеют тенденцию затухать с ростом интервала $n_2 - n_1$ между моментами наблюдений, так что в пределе $n_2 - n_1 \rightarrow +\infty$ переменные S_{n_1} и S_{n_2} становятся статистически

независимыми. Это, в частности, означает, что $\lim_{n_2-n_1 \rightarrow +\infty} C(n_1, n_2) = \lim_{n_2-n_1 \rightarrow +\infty} \langle S_{n_1} S_{n_2} \rangle = \langle S_{n_1} \rangle \langle S_{n_2} \rangle = \langle S_n \rangle^2 \leq \langle S_n^2 \rangle = C(n, n)$. В этой связи удобно ввести масштаб времени τ_c , называемый временем корреляции, на котором корреляционная функция выходит на свое асимптотическое значение. Анализ поведения отношения $\Pr[S_{n_2} = s_{n_2} | S_{n_1} = s_{n_1}] / \Pr[S_{n_2} = s_{n_2}]$ и взаимной информации $I(S_{n_1}, S_{n_2})$ обычно дает ту же оценку по порядку величины для характерного времени корреляции.

11.1.2 Independent Incremental Processes - Poisson Process

Теория по кому-то

Пуассоновский и винеровский процесс являются процессами с независимыми приращениями.

Так что изучим такие процессы.

Определение 3.1.

Процесс $(X_t, t \geq 0)$ называется процессом с независимыми приращениями, если для любого натурального n и любых $t_n > \dots > t_1 \geq 0$ случайные величины $X_{t_n} - X_{t_{n-1}}, \dots, X_{t_2} - X_{t_1}, X_{t_1}$ независимы в совокупности.

Отметим, что в случае дискретного времени процесс с независимыми приращениями является не чем иным, как случайным блужданием.

Тем самым процессы с независимыми приращениями естественно рассматривать как обобщение случайных блужданий для непрерывного времени.

В теории случайных процессов важнейшую роль играют вопросы существования процессов.

Дело в том, что совместные распределения случайных величин X_t не могут быть уж совсем произвольными, а должны быть в определенном роде согласованы.

Для процессов с независимыми приращениями критерий существования звучит следующим образом.

Теорема:

задано $Q_{s,t}$ -распределение вероятностей на $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ с характеристической функцией $\varphi_{s,t}$ пусть Q_0 -также некоторое распределение вероятностей на $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$. Тогда процесс $(X_t, t \geq 0)$ с независимыми приращениями и условиями

$$X_t - X_s \stackrel{d}{=} Q_{s,t} \quad \forall 0 \leq s < t \text{ и } X_0 \stackrel{d}{=} Q_0$$

существует $\Leftrightarrow \forall 0 \leq s < u < t, \forall \tau \in \mathbb{R}$

$$\varphi_{s,t}(\tau) = \varphi_{s,u}(\tau) \cdot \varphi_{u,t}(\tau)$$

Необходимость условия (4) почти очевидна из свойств характеристических функций, но оказывается, что его достаточно для установления существования процесса.

Первым процессом с независимыми приращениями нашего курса является пуассоновский процесс.

Определение пуассоновского случапроца

Случайный процесс $(N_t, t \geq 0)$ называется пуассоновским процессом интенсивности $\lambda > 0$, если выполнены три условия:

1. $N_0 = 0$ п.н.;
2. N_t имеет независимые приращения;
3. $N_t - N_s \sim \text{Pois}(\lambda(t-s))$ для любых неотрицательных $t > s$.
(ну и дальше куча теории Жуковского, потом мб пойму ее)

Суперпозиция потоков Пуассона

Важным свойством пуассоновского потока является его стабильность относительно сложения. Оказывается, суперпозиция двух независимых пуассоновских потоков с рейтами λ_1 и λ_2 дает пуассоновский поток с рейтом $\lambda_1 + \lambda_2$.

Для доказательства сложим количества событий, n_1 и n_2 , произошедших за один и тот же интервал времени t в двух пуассоновских потоках с рейтами λ_1 и λ_2 , соответственно. Характеристическая функция вероятностного распределения случайной величины $n = n_1 + n_2$ равна

$$\Phi(k) \equiv \langle e^{ik(n_1+n_2)} \rangle = \langle e^{ikn_1} \rangle \langle e^{ikn_2} \rangle = \varphi_1(k) \varphi_2(k),$$

где $\phi_1(k)$ и $\phi_2(k)$ - характеристические функции распределений Пуассона случайных величин n_1 и n_2 , то есть

$$\begin{aligned} \varphi_1(k) &\equiv \langle e^{ikn_1} \rangle = \exp(-\lambda_1 t) \sum_{n_1=0}^{+\infty} e^{ikn_1} \\ &= \exp(\lambda_1 t (e^{ik} - 1)) \\ \varphi_2(k) &\equiv \langle e^{ikn_2} \rangle = \exp(\lambda_2 t (e^{ik} - 1)) \end{aligned}$$

Таким образом, находим

$$\Phi(k) = \varphi_1(k) \varphi_2(k) = \exp((\lambda_1 + \lambda_2) t (e^{ik} - 1))$$

откуда следует что случайная величина n имеет пуассоновское распределение с рейтом $\lambda_1 + \lambda_2$.

Деление потока Пуассона

Аналогично, пуассоновский поток с рейтом λ можно разделить на два независимых пуассоновских подпотока с рейтами λ_1 и $\lambda_2 = \lambda - \lambda_1$. Деление исходного процесса производится посредством подбрасывания монетки с вероятностями орла и решки равными λ_1/λ и $1 - \lambda_1/\lambda$, случайный исход которого определяет, к какому из двух производных процессов будет принадлежать очередное пуассоновское событие. Проще всего указанное свойство можно доказать, показав их справедливость для процесса Бернулли с малой вероятностью единицы, и перейдя затем к континуальному пределу. Упражнение: Посетители прибывают в торговый комплекс в случайные моменты времени и независимо друг от друга со средней частотой λ . В каждом отдельном случае вероятность того, что очередной посетитель будет женского пола равна q . Найдите математическое ожидание и дисперсию числа женщин, прибывших в торговый комплекс за промежуток времени t .

Неоднородный поток Пуассона

Неоднородным потоком Пуассона называется поток событий, в котором вероятность $\lambda(t)dt$ наблюдать событие в промежутке времени $[t, t + dt]$ как и в однородном случае не зависит предшествующей истории наблюдений, но зависит от времени t . Такой процесс можно рассматривать как предельный случай последовательности независимых испытаний Бернулли, в которой вероятность выпадения единицы в очередной попытке зависит от номера этой попытки. К примеру, события распада радиоактивных атомов в образце, состоящим в начальный момент времени из большого их числа $N_0 \gg 1$, представляют собой неоднородный пуассоновский

Распределение числа событий

Пусть n - это число событий, произошедших за время t , отсчитываемое от произвольно выбранного момента начала наблюдений. Найдём функцию распределения вероятностей $p(n; t)$. В случае $n = 0$ получаем

$$P(0; t) = \Pr[qT_1 > t] = \int_t^{+\infty} dT_1 \frac{T_1}{\langle T_1 \rangle} P(T_1) \int_{t/T}^1 dq = \int_t^{+\infty} dT_1 \frac{T_1}{\langle T_1 \rangle} P(T_1) \left(1 - \frac{t}{T}\right)$$

где q - это отношение, в котором точка начала наблюдений делит межсобытийный интервал, на который она попала. Если $n > 0$, то искомое распределение можно записать как

$$\begin{aligned} P(n; t) &= \Pr[(qT_1 + T_2 + \dots + T_n < t) \cap (T_{n+1} > t - qT_1 - T_2 - \dots - T_n)] = (22) \\ &= \int_0^t d\tau \rho(\tau) \int_{t-\tau}^{+\infty} dT_{n+1} P(T_{n+1}) \end{aligned}$$

где $\rho(\tau)$ - плотность вероятности случайной величины $qT_1 + T_2 + \dots + T_n$.

Решение в Лаплас-представлении (!?????)

Из соотношения (23) с учетом теоремы о свертке следует простая связь между Лаплас-образами функциями $p(n; t)$, $\rho(\tau)$ и $P(T)$, а именно

$$\tilde{p}(n; s) = \frac{1 - \tilde{P}(s)}{s} \tilde{\rho}(s),$$

где

$$\begin{aligned} \tilde{\rho}(s) &= \langle e^{-s\tau} \rangle = \langle e^{-s(qT_1 + T_2 + \dots + T_n)} \rangle = \left\langle e^{-sqT_1} \prod_{i=2}^n e^{-sT_i} \right\rangle = \langle e^{-sqT_1} \rangle \prod_{i=2}^n \langle e^{-sT_i} \rangle \quad (25) \\ &= \tilde{P}^{n-1}(s) \int_0^{+\infty} dT_1 \frac{T_1}{\langle T_1 \rangle} P(T_1) \int_0^1 dq e^{-sqT_1} = \tilde{P}^{n-1}(s) \int_0^{+\infty} dT_1 P(T_1) \frac{1 - e^{-sT_1}}{\langle T_1 \rangle s} = (26) \\ &= \frac{\tilde{P}^{n-1}(s)(1 - \tilde{P}(s))}{\langle T \rangle s} \end{aligned}$$

Подставляя этот результат в Ур. (24), находим

$$\tilde{p}(n; s) = \frac{(1 - \tilde{P}(s))^2 \tilde{P}^{n-1}(s)}{\langle T \rangle s^2}$$

Упражнение: Проверьте, что в случае экспоненциального распределения интервала времени между последовательными событиями обратное преобразование Лапласа от выражения (28) дает распределение Пуассона.

Гауссовские процессы.

Винеровский процесс

Марковские моменты

Мартингалы

Большая часть оставшихся глав настоящего учебного пособия посвящена различным классам случайных процессов, наиболее важным с точки зрения приложений. В

предыдущих главах мы уже рассмотрели процессы с независимыми приращениями и гауссовские процессы. В этой главе речь пойдет о мартингалах. **Δ** Определение 6.1. Пусть $T \subset \mathbb{R}$. Процесс $(X_t, t \in T)$ называется мартингалом относительно фильтрации $\mathbb{F} = (\mathcal{F}_t, t \in T)$ если 1. X_t согласован с $\mathbb{F}$; 2. X_t является L^1 -процессом, т.е. для любого $t \in T$ выполнено неравенство $\mathbf{M}|X_t| < +\infty$ 3. $\mathbf{M}(X_t|\mathcal{F}_s) = X_s$ п. н. для любых $s \leq t, s, t \in T$. А Определение 6.2. Если вместо условия 3 выполнено

$$\mathbf{M}(X_t|\mathcal{F}_s) \geq X_s \text{ п. н. } \forall s \leq t, s, t \in T$$

то процесс X_t называется субмартингалом относительно $\mathbb{F}$. Если же вместо условия 3 выполнено

$$\mathbf{M}(X_t|\mathcal{F}_s) \leq X_s \text{ п. н. } \forall s \leq t, s, t \in T$$

то X_t называется супермартингалом относительно $\mathbb{F}$.

Если же вместо условия 3 выполнено

$$\mathbf{M}(X_t|\mathcal{F}_s) \leq X_s \text{ п. н. } \forall s \leq t, s, t \in T$$

то X_t называется супермартингалом относительно $\mathbb{F}$. Субмартингалы и супермартингалы - это мартингалы "наполовину". Для мартингалов условие 3 означает, что функция среднего является постоянной, $\mathbf{M}X_t = \mathbf{M}X_s \forall s, t \in T$. А Определение 6.3. Мартингалами (без указания фильтрации будем называть мартингалы относительно их естественной фильтрации).

Какие же процессы являются мартингалами? Следующая

Теорема

дает критерий мартингальности для процессов с независимыми приращениями.

независимые приращения. Тогда X_t - мартингал $\Leftrightarrow \mathbf{M}X_t = \text{const}$, т.е. функция среднего **по**— стоянна.

Несложно проверить, что если функция среднего не убывает, то процесс с независимыми приращениями является субмартингалом. Если же она не возрастает, то тогда - супермартингалом.

Теорема

6.1 сразу дает нам несколько хороших примеров мартингалов. 1. Винеровский процесс $(W_t, t \geq 0)$ - мартингал. 2. $(S_n, n \in \mathbb{Z}_+)$ - случайное блуждание, $S_n = \xi_1 + \dots + \xi_n, \xi_i$ - независимые случайные величины, $\mathbf{M}|\xi_i| < +\infty$. Тогда S_n - мартингал $\Leftrightarrow$ для любого i $\mathbf{M}\xi_i = 0$. 3. Пуассоновский процесс $(N_t, t \geq 0)$ - субмартингал. 4. Пусть $\xi_n, n \in \mathbb{N}$ - независимые неотрицательные случайные величины, $S_n = \xi_1 \cdot \dots \cdot \xi_n$. Тогда S_n - мартингал $\Leftrightarrow \mathbf{M}\xi_i = 1$ для всех i

Определение

мартингал Леви

Пусть ξ - случайная величина, $\mathbf{M}|\xi| < +\infty$ $\mathbb{F} = (\mathcal{F}_t, t \in T)$ - фильтрация. Докажите, что процесс

$$X_t = \mathbf{M}(\xi|\mathcal{F}_t)$$

Является мартингалом относительно $\mathbb{F}$.

Рассмотрим два важных свойства мартингалов. Первым является разложение Дуба-Мейера для произвольных процессов с дискретным временем. Для его формулировки нам понадобится понятие предсказуемого процесса. А Определение 6.4. Процесс $(X_n, n \in \mathbb{N})$ является предсказуемым относительно фильтрации $(\mathcal{F}_n, n \in \mathbb{Z}_+)$, если X_n измерим относительно $\mathcal{F}_{n-1}$ для всех $n \in \mathbb{N}$.

Теорема

Пусть $(X_n, n \in \mathbb{Z}_+)$ — L^1 -процесс, согласованный с фильтрацией $\mathbb{F} = (\mathcal{F}_n, n \in \mathbb{Z}_+)$. Тогда существует единственное разложение вида

$$X_n = M_n + A_n$$

где M_n — мартингал относительно $\mathbb{F}$, а

$(A_n, n \in \mathbb{N})$ — предсказуемый процесс относи- Заметим в связи с теоремой, что процесс $(X_n, n \in \mathbb{Z}_+)$ является субмартингалом относительно $\mathbb{F}$ тогда и только тогда, когда в его разложении Дуба-Мейера последовательность A_n является п. н. неубывающей.

Второе свойство связывает мартингалы с марковскими моментами и называется теоремой об остановке.

Пусть $(X_n, n \in \mathbb{Z}_+)$ — L^1 -процесс, согласованный с фильтрацией $\mathbb{F} = (\mathcal{F}_n, n \in \mathbb{Z}_+)$. Тогда X_n является мартингалом (субмартингалом) относительно $\mathbb{F} \Leftrightarrow$ для любых τ, σ — ограниченных почти наверное марковских моментов относительно $\mathbb{F}$ таких, что $\tau \leq \sigma$ п. н., выполнено

$$\mathbf{M}X_\tau = (\leq) \mathbf{M}X_\sigma$$

11.1.3 Markov Processes**Марковские процессы в двух словах (?)**

(освоюсь - напишу)

Марковские цепи с дискретным временем

Марковские цепи — это марковские процессы, принимающие не более чем счетное число значений. Пусть x — не более чем счетное множество. Для удобства всегда можно считать, что либо $X = \{1, \dots, N\}$, либо $X = \mathbb{N}$. **Определение 8.1.** Случайный процесс $(X_n, n \in \mathbb{Z}_+)$ со значениями в x называется марковской цепью, если для любого $m \in \mathbb{N}$, для любых $k_1 < \dots < k_m < k < n$ и для любых $a_1, \dots, a_m, i, j \in x$ выполнено (марковское свойство)

$$\mathbf{P}(X_n = j | X_k = i, X_{k_m} = a_m, \dots, X_{k_1} = a_1) = \mathbf{P}(X_n = j | X_k = i)$$

всегда, когда $\mathbf{P}(X_k = i, X_{k_m} = a_m, \dots, X_{k_1} = a_1) > 0$. Определение означает, что «будущее» процесса (X_n) и его «прошлое» $(X_{k_m}, \dots, X_{k_1})$ независимы при фиксированном «настоящем» (X_k) . Также из теоремы 7.1 следует, что марковская Цепь — это марковский процесс с не более чем счетным числом значений. Формальное утверждение, объясняющее обозначенную выше независимость, сформулировано ниже.

Марковские цепи с непрерывным временем

ё

Линейные преобразования в пространстве l^2 **Стационарные случайные процессы****Спектральное представление****Элементы стохастического исчисления Ито****11.1.4 On the use of Markov Processes (?)****О марковских цепях в моделях социального обслуживания (??)**

(такое тоже есть, напишу.)

11.1.5 Random Process Addendums

(всё остальное тут будет, потом буду про структуру думать уже)

11.2 methods for Setting Random Variables

(??)

Случайная величина задается производящим функционалом, который имеет вид:

$$F[\chi(t)] \equiv e^{-W[\chi(t)]} \equiv \left\langle e^{-i \int_0^t dt' \xi(t') \chi(t')} \right\rangle_{\xi}$$

Существует много частных случаев функции $W[\xi(t)]$, например, марковский:

$$W[\chi(t)] = \int_0^t dt' G(\chi(t'))$$

В марковском случае два разных интервала времени статистически независимы, то есть события на одном интервале не влияют на события на другом.

Частный случай марковского процесса является случай с гауссовой силой, когда $G(\chi) = D\chi^2$.

11.2.1 the Probability of Transition in the Presence of Random Processes

Теория

Построим простейшую модель, в которой есть случайность. Пусть у нас есть простейшее уравнение динамики системы: $\dot{x}(t) = -f(x(t))$, добавим к нему случайную величину, то, что получим называется уравнением Ланжевена:

$$\dot{x}(t) = -f(x(t)) + \xi(t).$$

Нашу добавку можно интерпретировать как случайную силу.

Простейшая задача, которую хочется решать, это задача о том, что если мы в момент времени $t = 0$ поместили частицу в точку x_i , с какой вероятностью она перейдет в окрестность точки x_f . (хз, почему, но запишем ее так):

$$P(x_f, x_i; t) = \langle \delta(x_f - x_{\xi}(t)) \rangle_{\xi}$$

Постараемся записать это в форме, которая нас сдвинет дальше. Преобразуем полный переход через сумму переходов за равные промежутки времени $\delta t = t/N$:

$$\delta(x_f - x_{\xi}(t)) = \delta(x_f - x_{n_1}) \delta(x_{n_1} - x_{\xi}(t)) = \dots = \prod_{k=1}^N \delta(x_{n_k} - x_{n_{k-1}}); \quad n_0 = i, \quad n_N = f.$$

Хочется переписать каждый переход с учетом уравнений движения. Для этого посмотрим на дискретную их версию в интегральной форме, которое приближенно

записывается в виде:

$$\begin{aligned} x(t_{k+1}) - x(t_k) &= - \int_{t_k}^{t_{k+1}} dt f(x(t)) + \int_{t_k}^{t_{k+1}} dt \xi(t) \approx \\ &\approx -\frac{1}{2} (t_{k+1} - t_k) (f(x(t_{k+1})) + f(x(t_k))) + \int_{t_k}^{t_{k+1}} dt \xi(t) \quad (11.1) \end{aligned}$$

Введем обозначение $x_k = x(t_k)$, в итоге один переход записывается как:

$$x_{k+1} - x_k = -\frac{\delta t}{2} (f(x_{k+1}) + f(x_k)) + \int_{t_k}^{t_{k+1}} dt \xi(t)$$

Вставка в дельта-функцию производится за счет ее свойства:
 $\delta(x_1 - x_2) = |f'(x_c)| \delta(f(x_c))$, поэтому

$$\begin{aligned} P(x_f, x_i; t) &= \int dx_{N-1} \dots dx_1 \prod_{k=1}^N \left(1 + \frac{\delta t}{2} f'(x_k) \right) \cdot \\ &\cdot \left\langle \delta^{N-1} \left(x_{k+1} - x_k + \frac{\delta t}{2} (f(x_{k+1}) + f(x_k)) - \int_{t_k}^{t_{k+1}} dt \xi(t) \right) \right\rangle_{\xi} \quad (11.2) \end{aligned}$$

Перепишем далее каждую дельта функцию через $\delta(x) = \int \frac{dp}{2\pi} e^{ipx}$, получим:

$$\begin{aligned} P(x_f, x_i; t) &= \frac{1}{(2\pi)^N} \int dx_{N-1} \dots dx_1 dp_N \dots dp_1 \prod_{k=1}^N \left(1 + \frac{\delta t}{2} f'(x_k) \right) \\ &\cdot e^{i \sum_{k=1}^N p_k (x_k - x_{k-1} + \frac{\delta t}{2} (f(x_k) + f(x_{k-1})))} \left\langle e^{-i \sum_{k=1}^N p_k \int_{t_{k-1}}^{t_k} dt \xi(t)} \right\rangle_{\xi} \quad (11.3) \end{aligned}$$

Далее учтем именно свойства самого случайного процесса, за счет добавления производящего функционала для марковского случая:

$$F[\chi(t)] \equiv e^{-W[\chi(t)]} \equiv \left\langle e^{-i \int_0^l dt' G(\chi(t'))} \right\rangle_{\xi}$$

Получим:

$$\begin{aligned} P(x_f, x_i; t) &= \frac{1}{(2\pi)^N} \int dx_{N-1} \dots dx_1 dp_N \dots dp_1 \prod_{k=1}^N \left(1 + \frac{\delta t}{2} f'(x_k) \right) \cdot \\ &\cdot e^{i \sum_{k=1}^N p_k (x_k - x_{k-1} + \frac{\delta t}{2} (f(x_k) + f(x_{k-1})))} \left\langle e^{-\sum_{k=1}^N \delta t G(p_k)} \right\rangle_{\xi} \quad (11.4) \end{aligned}$$

Теперь у нас есть готовое выражение для дискретного времени, которое будем решать. Его можно записать и в случае непрерывных переходов, для этого нужно перейти к непрерывной мере, записав бесконечное произведение в виде члена в экспоненте:

$$P(x_f, x_i; t) = \int_{x(0)=x_i}^{x(t)=x_f} \mathcal{D}x(\tau) \mathcal{D}p(\tau) e^{i \int_0^t d\tau p(\tau)(\dot{x}(\tau) + f(x(\tau))) - \int_0^t d\tau G(p(\tau)) + \frac{1}{2} \int_0^t d\tau f'(x(\tau))}$$

А также в случае гауссовой силы $G(\chi) = D\chi^2$ после интегрирования по импульсу получаем (???)

$$P(x_f, x_i; t) = \int_{x(0)=x_i}^{x(t)=x_f} \mathcal{D}x(\tau) e^{-\frac{1}{4D} \int_0^t d\tau (\dot{x}(\tau) + f(x(\tau)))^2 + \frac{1}{2} \int_0^t d\tau f'(x(\tau))}$$

Или в более понятном виде:

$$P(x_f, x_i; t) = e^{-\frac{1}{2D} \int_{x_i}^{x_f} dy f(y)} \int_{x(0)=x_i}^{x(t)=x_f} \mathcal{D}x(\tau) e^{-\int_0^t d\tau \left(\frac{\dot{x}^2(\tau)}{4D} + \frac{1}{4D} f^2(x(\tau)) - \frac{1}{2} f'(x(\tau)) \right)}.$$

пример: случайные блуждания

Рассмотрим разные задачи случайного блуждания и их решения.

Блуждания с нулевой силой

Рассмотрим случай $f = 0$. Вероятность перейти в конечную точку равна:

$$P(x_f, x_i; t) = \int_{x(0)=x_i}^{x(t)=x_f} \mathcal{D}x(\tau) e^{-\int_0^t d\tau \frac{\dot{x}^2(\tau)}{4D}}$$

Для вычисления данного интеграла заменим:

$$x_{\text{new}} = x + \dot{x} \tau = x_i + \frac{x_f - x_i}{t} \tau$$

После замены у нас выносятся константа:

$$P(x_f, x_i; t) = e^{-\frac{(x_f - x_i)^2}{4Dt}} \int_{x(0)=0}^{x(t)=0} \mathcal{D}x(\tau) e^{-\int_0^t d\tau \frac{\dot{x}^2(\tau)}{4D}}$$

Далее нужно посчитать интеграл по траекториям. Мы его считать не умеем, однако продвинуться нам помогут следующие соображения: во-первых заметим, что он константа; во-вторых так как частица в x_f где-нибудь на координатной оси точно

находится, просуммировав все координаты, то есть проинтегрировав, мы должны получить единицу, так что:

$$\left(\int_{x(0)=0}^{x(t)=0} \mathcal{D}x(\tau) e^{-\int_0^t d\tau \frac{\dot{x}^2(\tau)}{4D}} \right) \int_{-\infty}^{\infty} dx_f e^{-\frac{(x_f - x_i)^2}{4Dt}} = 1,$$

и поэтому

$$\int_{x(0)=0}^{x(t)=0} \mathcal{D}x(\tau) e^{-\int_0^t d\tau \frac{\dot{x}^2(\tau)}{4D}} = \frac{1}{\sqrt{4\pi Dt}}.$$

Окончательно задача решена:

$$P(x_f, x_i; t) = \frac{1}{\sqrt{4\pi Dt}} e^{-\frac{(x_f - x_i)^2}{4Dt}}$$

(место количественных расчетов и на самом деле полезных выводов)

блуждание с силой, не зависящей от времени

$$P(x_f, x_i; t) = \int_{x(0)=x_i}^{x(t)=x_f} \mathcal{D}x(\tau) e^{-\int_0^t d\tau \frac{(t(\tau)+f(\tau))^2}{4D}}$$

$$\ddot{x}_c(\tau) = -\dot{f}(\tau)$$

$$x_c(\tau) = y(\tau) + \frac{x_f - y(t)}{t} \tau \quad y(\tau) = x_i - \int_0^\tau d\tau' f(\tau')$$

$$P(x_f, x_i; t) = \frac{1}{\sqrt{4\pi Dt}} e^{-\frac{(x_f - y(t))^2}{4Dt}}$$

$$P(x_f, x_i; t) = e^{-\frac{\alpha(x^2 - x_i^2)}{4D} + \frac{\alpha t}{2}} \int_{x(0)=x_i}^{x(t)=x_f} \mathcal{D}x(\tau) e^{-\int_0^t d\tau \left(\frac{t^2(\tau)}{10} + \frac{\alpha^2 x^2(\tau)}{4D} \right)}$$

11.2.2 Fundamentals of Brownian Motion

(тут самое основное и суть)

Теория

О применениях (!!?)

(хз, пропишу тут, тоже важно.)

статистика частиц в идеальном газе

(задача с избором)

Пусть у нас есть идеальный газ который содержит N молекул. Каждая молекула может быть равномерно размазана по объёму V всего газа. Выделим в этом газе объём v и зададимся вопросом а какова вероятность того, что в этом объёме ровно m частиц.

нам нужна случайная величина, в данном случае это количество частиц в объёме назовём её ξ

11.3 Sequences of Past-dependent Tests

простейшие такие цепи - цепи маркова
дальше изучим...

11.3.1 Basics of Markov Chains

опишем цепи маркова и покажем их применения.

Теория по Горайнову (??)

Пусть некоторая система меняет свое состояние случайным образом в моменты времени $t = 1, 2, \dots, T$. Множество возможных состояний системы является конечным множеством $E = \{e_1, \dots, e_r\}$ и называется фазовым пространством.

Среди реальных систем важный класс образуют такие, у которых вероятности перехода из одного состояния в другое в данный момент времени не зависят от того, как вела себя система в предыдущие моменты времени.

Такие системы называют марковскими или цепями Маркова.

Эволюция изучаемой системы описывается траекторией $\omega = (\omega_0, \omega_1, \dots, \omega_T)$, где $\omega_t = i$, если система в момент времени t находилась в состоянии e_i . Поэтому под элементарным событием будем понимать $\omega = (\omega_0, \omega_1, \dots, \omega_T)$, т.

е.

последовательность номеров состояний, в которых находилась система в моменты времени $0, 1, \dots, T$. Можно считать, что цепь Маркова является обобщением схемы независимых испытаний, в которой условие независимости заменяется на некоторые другие естественные предположения.

Рассмотрим эти предположения и введем некоторые понятия, связанные с цепями Маркова.

Для вычисления вероятности того, что траектория $\omega = (\omega_0, \omega_1, \dots, \omega_T)$ будет реализована как $i_0, i_1, \dots, i_T$, можно воспользоваться теоремой умножения

$$\begin{aligned} P(\omega_0 = i_0, \dots, \omega_T = i_T) = \\ P(\omega_0 = i_0) P(\omega_1 = i_1 | \omega_0 = i_0) \cdots P(\omega_T = i_T | \omega_0 = i_0, \dots, \omega_{T-1} = i_{T-1}) \end{aligned}$$

Здесь мы считаем, что условная вероятность относительно события с нулевой вероятностью равна нулю.

Условие марковости (независимость от прошлого) выражается в том, что для любых $s < t$ выполняются равенства

$$P(\omega_t = j | \omega_0 = i_0, \dots, \omega_{s-1} = i_{s-1}, \omega_s = i) = P(\omega_t = j | \omega_s = i)$$

$i, j = 1, \dots, r$.

Другими словами, вероятность перехода системы, находившейся в состоянии i в момент времени s , в состояние j в момент времени t не зависит от того, как она себя вела до момента времени s .

Кроме того, мы будем рассматривать однородные марковские цепи, для которых условные вероятности перехода из состояния i в состояние j за время t не зависят от того, в какой момент времени s она находилась в состоянии i .

е.

$$P(\omega_{s+t} = j | \omega_s = i) = P(\omega_t = j | \omega_0 = i) = p_{ij}(t)$$

Вероятности $p_{ij}(t)$ обладают следующими свойствами:

$$1^\circ \cdot p_{ij}(t) \geq 0$$

$$2^\circ \cdot \sum_{j=1}^r p_{ij}(t) = 1$$

$$3^\circ \cdot p_{ij}(0) = \delta_{ij}$$

Матрицу, составленную из вероятностей $p_{ij}(t)$, будем обозначать $\Pi(t)$. Свойства 1° и 2° означают, что $\Pi(t)$ является стохастической матрицей.

Вероятности перехода системы из состояния i в состояние j за единичное время $p_{ij}(1)$ называются переходными вероятностями и обозначаются p_{ij} . Матрица $\Pi(1) = \Pi$ называется матрицей переходных вероятностей.

Однородная марковская цепь вполне определяется вектором-строкой $\vec{p}(0) = (p_1(0), \dots, p_r(0))$ начальных вероятностей ($p_i(0)$ — вероятность того, что в начальный момент времени система находилась в состоянии i) и матрицей переходных вероятностей Π .

Действительно, вероятность $p_j(t)$ того, что система будет находиться в состоянии j в момент времени t определяется равенством

$$p_j(t) = P(\omega_t = j) = \sum_{i=1}^r p_i(0) p_{ij}(t)$$

Если $\vec{p}(t)$ — вектор-строка вероятностей состояний в момент времени t , то

$$\vec{p}(t) = \vec{p}(0)\Pi(t)$$

С другой стороны, матрица $\Pi(t)$ выражается через Π .

Это следует из уравнений Колмогорова-Чепмена

$$p_{ij}(s+t) = \sum_{k=1}^r p_{ik}(s) p_{kj}(t)$$

которые выводятся непосредственно из свойств марковости и однородности

$$\begin{aligned} p_{ij}(s+t) &= P(\omega_{s+t} = j | \omega_0 = i) \\ &= \frac{P(\omega_{s+t} = j, \omega_0 = i)}{P(\omega_0 = i)} \\ &= \sum_{k=1}^r \frac{P(\omega_{s+t} = j, \omega_0 = i, \omega_s = k)}{P(\omega_0 = i)} \\ &= \sum_{k=1}^r \frac{P(\omega_s = k, \omega_0 = i) P(\omega_{s+t} = j | \omega_0 = i, \omega_s = k)}{P(\omega_0 = i)} \\ &= \sum_{k=1}^r p_{ik}(s) p_{kj}(t) \end{aligned}$$

Уравнения (11.1) можно также записать в матричном виде

$$\Pi(s+t) = \Pi(s)\Pi(t)$$

Поскольку $\Pi(1) = \Pi$, то $\Pi(t) = \Pi^t$. Следовательно

$$\vec{p}(t) = \vec{p}(0)\Pi^t$$

Важное значение в теории однородных марковских цепей имеет следующий результат. Тогда для каждого $j = 1, \dots, r$ существует предел

$$p_j = \lim_{t \rightarrow \infty} p_{ij}(t)$$

При этом $p_j > 0, j = 1, \dots, r$, не зависят от i и являются единственным решением системы

$$\sum_{k=1}^r x_k p_{kj} = x_j, \quad \sum_{k=1}^r x_k = 1$$

ДоКАЗАТЕЛЬСТВО.

Введем в рассмотрение

$$M_j(t) = \max_{1 \leq i \leq r} p_{ij}(t), \quad m_j(t) = \min_{1 \leq i \leq r} p_{ij}(t)$$

В силу уравнений (11.1)

$$p_{ij}(1+t) = \sum_{k=1}^r p_{ik} p_{kj}(t)$$

Отсюда и непосредственно из определения $m_j(t)$ и $M_j(t)$ получаем неравенства

$$m_j(t) \leq p_{ij}(t+1) \leq M_j(t)$$

Из этих неравенств, в свою очередь, следуют неравенства

$$m_j(t) \leq m_j(t+1) \leq M_j(t+1) \leq M_j(t)$$

Это означает, что $m_j(t)$ не убывает по t , а $M_j(t)$ не возрастает по t . Обозначим

$$\varepsilon = \min_{1 \leq i, j \leq r} p_{ij}(t_0)$$

По условию теоремы $0 < \varepsilon < 1$. Для $t > 0$, используя уравнения (11.1), получаем

$$\begin{aligned} p_{ij}(t_0+t) &= \sum_{k=1}^r p_{ik}(t_0) p_{kj}(t) \\ &= \sum_{k=1}^r [p_{ik}(t_0) - \varepsilon p_{jk}(t)] p_{kj}(t) + \varepsilon \sum_{k=1}^r p_{jk}(t) p_{kj}(t) \\ &= \sum_{k=1}^r [p_{ik}(t_0) - \varepsilon p_{jk}(t)] p_{kj}(t) + \varepsilon p_{jj}(2t) \end{aligned}$$

Замечая, что выражения в квадратных скобках неотрицательны, получаем неравенство

$$p_{ij}(t_0+t) \leq M_j(t) \sum_{k=1}^r [p_{ik}(t_0) - \varepsilon p_{jk}(t)] + \varepsilon p_{jj}(2t) = (1-\varepsilon)M_j(t) + \varepsilon p_{jj}(2t)$$

из которого следует, что

$$M_j(t_0 + t) \leq (1 - \varepsilon)M_j(t) + \varepsilon p_{jj}(2t)$$

Аналогично получаются неравенства

$$p_{ij}(t_0 + t) \geq (1 - \varepsilon)m_j(t) + \varepsilon p_{jj}(2t)$$

И

$$m_j(t_0 + t) \geq (1 - \varepsilon)m_j(t) + \varepsilon p_{jj}(2t)$$

Сравнивая полученные выше неравенства, приходим к следующему

$$M_j(t_0 + t) - m_j(t_0 + t) \leq (1 - \varepsilon)(M_j(t) - m_j(t))$$

Применяя последовательно k раз последнее неравенство, получаем

$$M_j(kt_0 + t) - m_j(kt_0 + t) \leq (1 - \varepsilon)^k (M_j(t) - m_j(t))$$

Отсюда сразу же следует, что

$$\lim_{k \rightarrow \infty} [M_j(kt_0 + t) - m_j(kt_0 + t)] = 0$$

а поскольку последовательность $\{M_j(t) - m_j(t)\}_{t=1}^{\infty}$ монотонно убывает, то

$$\lim_{t \rightarrow \infty} [M_j(t) - m_j(t)] = 0$$

Таким образом, доказано существование пределов

$$p_j = \lim_{t \rightarrow \infty} M_j(t) = \lim_{t \rightarrow \infty} m_j(t) = \lim_{t \rightarrow \infty} p_{ij}(t)$$

$j = 1, \dots, r$. Заметим также, что поскольку $m_j(t)$ не убывают по t , то

$$p_j \geq m_j(t_0) \geq \varepsilon > 0$$

т.

е.

все p_j строго положительны.

При этом

$$\sum_{j=1}^r p_j = \lim_{t \rightarrow \infty} \sum_{j=1}^r p_{ij}(t) = 1$$

Осуществляя в равенстве

$$p_{ij}(t+1) = \sum_{k=1}^r p_{ik}(t)p_{kj}$$

предельный переход при $t \rightarrow \infty$, получаем

$$p_j = \sum_{k=1}^r p_k p_{kj}$$

т.

е.

$p_1, \dots, p_r$ - решение системы (11.2). Допустим теперь, что $q_1, \dots, q_r$ также является решением этой системы, т.

е.

$$q_j = \sum_{k=1}^r q_k p_{kj}, \quad \sum_{k=1}^r q_k = 1$$

Заметим, что тогда $q_1, \dots, q_r$ будет решением и системы

$$q_j = \sum_{k=1}^r q_k p_{kj}(t)$$

при $t = 1, 2, \dots$. Действительно

$$\sum_{k=1}^r q_k p_{kj}(t) = \sum_{k=1}^r q_k \sum_{l=1}^r p_{kl} p_{lj}(t-1) = \sum_{l=1}^r q_l p_{lj}(t-1) = \dots = \sum_{k=1}^r q_k p_{kj} = q_j$$

Осуществляя в равенстве (11.3) предельный переход при $t \rightarrow \infty$, получаем

$$q_j = \sum_{k=1}^r q_k p_j = p_j$$

и единственность решения системы (11.2) доказана.

Доказанная Теорема относится к так называемым эргодическим теоремам.

Ее смысл состоит в том, что за длительный промежуток времени система как бы забывает, из какого начального состояния она стартовала.

Условие положительности элементов матрицы Π^{t_0} означает, что с положительной вероятностью из любого состояния в любое состояние система может перейти за время t_0 |

Марковское свойство

Простейшей моделью коррелированного случайного процесса является цепь Маркова - случайный процесс с дискретным пространством состояний и дискретным временем, обладающий марковским свойством. Рассмотрим стохастическую переменную S_n , где n играет роль дискретного времени, возможные значения которой принадлежат счетному множеству Ω . Говорят, что стохастический процесс S_n обладает марковским свойством, если условное распределение вероятностей его будущих состояний зависит только от текущего состояния, но не от последовательности более ранних состояний, то есть

$$\Pr [S_{n+1} = s_{n+1} | S_n = s_n, S_{n-1} = s_{n-1}, \dots, S_1 = s_1] = \Pr [S_{n+1} = s_{n+1} | S(n) = s_n]$$

В силу наличия влияния текущего состояния на состояние в сразу следующий за этим момент, в общем случае значения марковской переменной величины в различные моменты времени оказываются скоррелированными друг с другом. Отметим, что следствием Ур. (4) является равенство

$$\Pr [S_n = s_n | S_{n-m} = s_{n-m}, S_{n-k} = s_{n-k}] = \Pr [S_n = s_n | S_{n-m} = s_{n-m}],$$

где $k > m > 0$.

Теория по Вшэ

Управляющее уравнение по ВШЭ

Итак, наш случайный процесс совершает прыжки между двумя состояниями, причем если в момент времени процесс находится в состоянии 1, то $S_n = 1$, а если в состоянии

2, то $S_n = 0$. Обозначим вероятности обнаружить процесс на n -м шаге в состояниях 1 и 2, через $\pi_1(n)$ и $\pi_2(n)$, соответственно. Предполагая марковское свойство стохастической динамики процесса и используя закон полной вероятности, распределение вероятностей на следующем шаге однозначно можно выразить через соответствующее распределение на шаге предыдущем как

$$\begin{aligned}\pi_1(n+1) &= \Pr[S_{n+1} = 1|S_n = 1] \pi_1(n) + \Pr[S_{n+1} = 1|S_n = 2] \pi_2(n) \\ \pi_2(n+1) &= \Pr[S_{n+1} = 2|S_n = 1] \pi_1(n) + \Pr[S_{n+1} = 2|S_n = 2] \pi_2(n)\end{aligned}$$

Условные вероятности переходов между состояниями будем считать не зависящими от времени n и обозначим через $\Pr[S_{n+1} = 2|S_n = 1] = 1 - \Pr[S_{n+1} = 1|S_n = 1] = p$, $\Pr[S_{n+1} = 1|S_n = 2] = 1 - \Pr[S_{n+1} = 2|S_n = 2] = q$. Тогда

$$\begin{aligned}\pi_1(n+1) &= (1-p)\pi_1(n) + q\pi_2(n) \\ \pi_2(n+1) &= p\pi_1(n) + (1-q)\pi_2(n)\end{aligned}$$

Удобно ввести двухкомпонентный вектор вероятностей $\vec{\pi}(n) = (\pi_1(n), \pi_2(n))^T$ и перейти к матричной форме записи

$$\vec{\pi}(n+1) = \hat{P} \cdot \vec{\pi}(n),$$

где матрица $\hat{P}$ называется матрицей вероятностей перехода; ее элементы имеют вид $P_{11} = 1-p$, $P_{12} = q$, $P_{21} = p$, $P_{22} = 1-q$. Уравнение (10) называется управляющим.

Solution управляющего уравнения

Если известно распределение вероятностей $\vec{\pi}(0)$ в начальный момент времени, то решение уравнения (10) записывается как

$$\vec{\pi}(n) = \hat{P}^n \cdot \vec{\pi}(0).$$

Дальнейший анализ требует возведения матрицы $\hat{P}$ в натуральную степень n , что проще всего сделать, представив эту матрицу в виде $\hat{P} = \hat{Q}\hat{D}\hat{Q}^{-1}$, где $\hat{D}$ - диагональная матрица, на диагонали которой стоят собственные числа матрицы $\hat{P}$, $\hat{Q}$ - матрица, составленная из собственных столбцов матрицы $\hat{P}$. Решая уравнение $\det(\hat{P} - \lambda\hat{I}) = 0$ находим пару собственных чисел $\lambda_1 = 1$ и $\lambda_2 = 1-p-q$, которым соответствуют собственные столбцы $\vec{a}_1 = (1/p, 1/q)^T$, $\vec{a}_2 = (1, -1)^T$, по определению удовлетворяющие условию $\hat{P}\vec{a} = \lambda\vec{a}$. Используя эти результаты, запишем

$$\begin{aligned}\hat{P}^n &= \hat{Q}\hat{D}^n\hat{Q}^{-1} = \begin{pmatrix} \frac{1}{p} & 1 \\ \frac{1}{q} & -1 \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ 0 & 1-p-q \end{pmatrix}^n \cdot \begin{pmatrix} \frac{pq}{p+q} & \frac{pq}{p+q} \\ \frac{p}{p+q} & -\frac{q}{p+q} \end{pmatrix} \neq (12) \\ &= \begin{pmatrix} \frac{q}{p+q} + \frac{p}{p+q}(1-p-q)^n & \frac{q}{p+q} - \frac{q}{p+q}(1-p-q)^n \\ \frac{p}{p+q} - \frac{p}{p+q}(1-p-q)^n & \frac{p}{p+q} + \frac{q}{p+q}(1-p-q)^n \end{pmatrix}\end{aligned}$$

Подстановка Ур. (13) в Ур. (11) дает

$$\begin{aligned}\pi_1(n) &= \frac{q}{p+q} + \frac{p}{p+q}(1-p-q)^n\pi_1(0) - \frac{q}{p+q}(1-p-q)^n\pi_2(0) \\ \pi_2(n) &= \frac{p}{p+q} - \frac{p}{p+q}(1-p-q)^n\pi_1(0) + \frac{q}{p+q}(1-p-q)^n\pi_2(0)\end{aligned}$$

При выводе этих выражений мы использовали условие нормировки $\pi_1(0) + \pi_2(0) = 1$. Отметим также, что для решения, заданного уравнениями (32) и (33), в любой момент времени справедливо $\pi_1(n) + \pi_2(n) = 1$.

Релаксация к стационарному распределению вероятностей

Из Ур. (32) и (33) видно, что если $0 < p + q < 2$, то любое начальное условие на больших временах релаксирует к стационарному распределению вероятностей

$$\lim_{n \rightarrow +\infty} \pi_1(n) = \bar{\pi}_1^{\text{st}} = \frac{q}{p+q},$$

$$\lim_{n \rightarrow +\infty} \pi_2(n) = \bar{\pi}_2^{\text{st}} = \frac{p}{p+q},$$

которое, как легко проверить, является решением уравнения $\hat{P} \cdot \bar{\pi}^{\text{st}} = 0$. Характерное время релаксации оценивается как $\tau_{\text{relax}} \sim \frac{1}{\ln \frac{1}{|1-p-q|}}$.

Анализ автокорреляций

Продемонстрируем при помощи Ур. (32) и (33) скоррелированность случайного процесса S_n в разные моменты времени. Условная вероятность $\Pr[S_{n_2} = 1 | S_{n_1} = 1]$, где $n_2 \geq n_1$, равна решению π_1 , заданному формулой (32), в которую необходимо подставить $n \rightarrow n_2 - n_1$, $\pi_1(0) = 1$, $\pi_2(0) = 0$. Таким образом, находим

$$\Pr[S_{n_2} = 1 | S_{n_1} = 1] = \frac{q}{p+q} + \frac{p}{p+q}(1-p-q)^{n_2-n_1}$$

Аналогично можно выразить условные вероятности $\Pr[S_{n_2} = 0 | S_{n_1} = 1]$, $\Pr[S_{n_2} = 1 | S_{n_1} = 0]$ и $\Pr[S_{n_2} = 0 | S_{n_1} = 0]$. Из (18) видно, что если $p \neq 1/2$ и $q \neq 1/2$, то $\Pr[S_{n_2} = 1 | S_{n_1} = 1] \neq \Pr[S_{n_2} = 1] = \frac{q}{p+q}$, так что значения случайной функции S_n в различные моменты времени не являются статистически независимыми. Время корреляции, очевидно, дается той же оценкой, что и время релаксации, т.е. $\tau_c \sim \frac{1}{\ln \frac{1}{|1-p-q|}}$. Упражнение: Вычислите парную корреляционную функцию $C(n_1, n_2) = \langle S_{n_1} S_{n_2} \rangle$ для рассмотренного случайного процесса.

Марковские цепи с непрерывным временем

В рассмотренной модели цепи Маркова роль времени выполнял дискретный параметр n - номер наблюдения. Обсудим теперь процесс Маркова с двумя состояниями и непрерывным временем, см. Рис. 1(b). Стохастическая динамика такой цепи по определению подчиняется следующим правилам. Если в момент времени t процесс находится в состоянии 1, то в течении промежутка времени $[t, t+dt]$ с вероятностью λdt он перейдет в состояние 2, а с вероятностью $1 - \lambda dt$ останется в состоянии 1. Если же в момент t процесс находится в состоянии 2, то в течении промежутка времени $[t, t+dt]$ он с вероятностью μdt перейдет в состояние 1, и с вероятностью $1 - \mu dt$ останется в состоянии 2. Марковость процесса означает, что указанные вероятности не зависят от предшествующей истории процесса. С учетом сказанного вероятности $\pi_1(t+dt)$ и $\pi_2(t+dt)$ обнаружить процесс в состояниях 1 и 2 в момент $t+dt$ равны (см. Ур. (6) и (7))

$$\pi_1(t+dt) = (1 - \lambda dt)\pi_1(t) + \mu dt \cdot \pi_2(t)$$

$$\pi_2(t+dt) = \lambda dt \cdot \pi_1(t) + (1 - \mu dt)\pi_2(t)$$

откуда находим систему дифференциальных уравнений

$$\frac{d\pi_1(t)}{dt} = -\lambda\pi_1(t) + \mu\pi_2(t)$$

$$\frac{d\pi_2(t)}{dt} = \lambda\pi_1(t) - \mu\pi_2(t)$$

которую иногда удобно записывать в матричном виде

$$\frac{d\vec{\pi}(t)}{dt} = \hat{Q}\vec{\pi}(t),$$

где $Q_{11} = -Q_{21} = -\lambda$ и $Q_{12} = -Q_{22} = \mu$.

Решение для произвольного начального условия

Решение Ур. (21) и (22) при заданном начальном распределении вероятностей $\pi_1(0)$ и $\pi_2(0)$ может быть построено, к примеру, методом преобразования Лапласа. Умножив обе эти части уравнений на e^{-st} и проинтегрировав их по t от 0 до $+\infty$, получим,

$$\begin{aligned} s\tilde{\pi}_1(s) - \pi_1(0) &= -\lambda\tilde{\pi}_1(s) + \mu\tilde{\pi}_2(s) \\ s\tilde{\pi}_2(s) - \pi_2(0) &= \lambda\tilde{\pi}_1(s) - \mu\tilde{\pi}_2(s) \end{aligned}$$

где $\tilde{\pi}_1(s) = \int_0^{+\infty} dt e^{-st} \pi_1(t)$ и $\tilde{\pi}_2(s) = \int_0^{+\infty} dt e^{-st} \pi_2(t)$. Значит

$$\begin{aligned} \tilde{\pi}_1(s) &= \frac{s + \mu}{s(s + \lambda + \mu)} \pi_1(0) + \frac{\mu}{s(s + \lambda + \mu)} \pi_2(0) \\ \tilde{\pi}_2(s) &= \frac{s + \lambda}{s(s + \lambda + \mu)} \pi_2(0) + \frac{\lambda}{s(s + \lambda + \mu)} \pi_1(0) \end{aligned}$$

Вычисляя обратное преобразование Лапласа, получаем

$$\begin{aligned} \pi_1(t) &= \frac{1}{2\pi i} \int_{\gamma-i\infty}^{\gamma+i\infty} ds e^{st} \tilde{\pi}_1(s) = \frac{\mu}{\lambda + \mu} + \frac{\lambda}{\lambda + \mu} e^{-(\lambda+\mu)t} \pi_1(0) - \frac{\mu}{\lambda + \mu} e^{-(\lambda+\mu)t} \pi_2(0), \\ \pi_2(t) &= \frac{1}{2\pi i} \int_{\gamma-i\infty}^{\gamma+i\infty} ds e^{st} \tilde{\pi}_2(s) = \frac{\lambda}{\lambda + \mu} - \frac{\lambda}{\lambda + \mu} e^{-(\lambda+\mu)t} \pi_1(0) + \frac{\mu}{\lambda + \mu} e^{-(\lambda+\mu)t} \pi_2(0). \end{aligned}$$

Представленный путь вывода, разумеется, не единственный возможный. Заметим, что это решение можно было получить из уравнений (32) и (33), подставив $p = \lambda\Delta t$, $q = \mu\Delta t$ и перейдя к пределу $\Delta t \rightarrow 0$. Кроме того, можно было подставить соотношение $\pi_1(t) = 1 - \pi_2(t)$ в (21) и решить получившееся линейное дифференциальное уравнение первого порядка. Наконец, решение Ур. (23) может быть выписано через матричную экспоненту.

Релаксация к стационарному решению

Из формул (28) и (29) видно, что любое начальное условие экспоненциально быстро релаксирует к стационарному распределению вероятностей

$$\begin{aligned} \lim_{t \rightarrow +\infty} \pi_1(t) &= \bar{\pi}_1^{\text{st}} = \frac{\mu}{\lambda + \mu}, \\ \lim_{t \rightarrow +\infty} \pi_2(t) &= \bar{\pi}_2^{\text{st}} = \frac{\lambda}{\lambda + \mu}, \end{aligned}$$

которое, как легко проверить, является решением уравнения $\hat{Q} \cdot \vec{\pi}^{\text{st}} = 0$. Соответствующее время релаксации оценивается как $(\lambda + \mu)^{-1}$. Этот временной масштаб можно также интерпретировать как время корреляции марковской цепи, то есть время, в течении которого марковский процесс помнит свое исходное состояние.

О моделях социальных сервисов (??)

(типичные задачи на это даются, чуть что буду разбирать эту теорию.)

ВШЭ $M/M/\infty$ парковка

(мб в каталог перенесу потом)

Покажем, как происходит конструирование математической модели системы массового обслуживания на следующем простом примере. Предположим, что перед нами стоит задача обустроить автомобильную стоянку для посетителей нового торгового комплекса. На основании статистических данных мы можем оценить число посетителей, которые будут ежедневно приезжать в торговый комплекс на личных автомобилях, а также среднюю продолжительность визита типичного посетителя. Насколько большой следует сделать стоянку, чтобы быть уверенным, что большинству посетителей хватит парковочных мест?

Для ответа на этот вопрос поступим следующим образом. Оценим статистически-стационарное распределение вероятностей π_n^{st} обнаружить на стоянке ровно n автомобилей, которое имело бы место в том случае, если бы стоянка была неограниченно большой. Тогда для обеспечения большинства посетителей возможностью припарковать свой автомобиль число мест N на проектируемой нами стоянке должно быть таким, чтобы в рамках используемой вероятностной модели выполнялось условие $\Pr[n > N] \ll 1$. Вероятностное распределение π_n^{st} числа занятых парковочных мест может быть вычислено аналитически, если сделать два предположения: (1) автомобили прибывают на стоянку независимо друг от друга в случайные моменты времени со средней частотой λ , (2) каждый автовладелец оставляет свой автомобиль на случайное время, имеющее экспоненциальную функцию плотности вероятности с параметром μ . Значения λ и μ на практике должны оцениваться исходя из статистических данных. Отметим, что указанная модель в англоязычной литературе часто обозначается как $M/M/\infty$, где первая M (от англ. memoryless) обозначает пуассоновскую статистику событий поступления новых клиентов, вторая M обозначает экспоненциальность статистики времени, затрачиваемого на обслуживание одного клиента, и, наконец, знак ∞ соответствует предположению о бесконечно большом числе каналов обслуживания (т.е. парковочных мест в контексте обсуждаемой задачи).

Итак, состояние стоянки в произвольный момент времени t характеризуется числом занятых парковочных мест n . С учетом сделанных предположений, стохастические переходы между этими состояниями описываются марковской цепью, изображенной на Рис. (2). А значит эволюция распределения вероятностей $\pi_n(t)$ подчиняется следующей системе дифференциальных уравнений первого порядка

$$\begin{aligned}\frac{d\pi_0(t)}{dt} &= -\lambda\pi_0(t) + \mu\pi_1(t) \\ \frac{d\pi_n(t)}{dt} &= \lambda\pi_{n-1}(t) - (\lambda + n\mu)\pi_n(t) + (n+1)\mu\pi_{n+1}(t), \quad \text{при } n \geq 1,\end{aligned}$$

Занулив левые части этих уравнений, получим, что стационарные вероятности удовлетворяют соотношению $\pi_n^{\text{st}} = \frac{1}{n!} \left(\frac{\lambda}{\mu}\right)^n \pi_0^{\text{st}}$. Определяя π_0^{st} из условия нормировки $\sum_{n=0}^{+\infty} \pi_n^{\text{st}} = 1$, находим, что искомое стационарное решение дается распределением Пуассона

$$\pi_n^{\text{st}} = \frac{1}{n!} \left(\frac{\lambda}{\mu}\right)^n \exp\left(-\frac{\lambda}{\mu}\right)$$

Математическое ожидание и дисперсия числа автомобилей на стоянке в рамках рассматриваемой модели равны $\langle n \rangle = \frac{\lambda}{\mu}$ и $\sigma_n^2 = \frac{\lambda}{\mu}$, соответственно. Легко видеть, что $\sigma_n \ll \langle n \rangle$ при $\langle n \rangle \gg 1$. Благодаря быстрому затуханию π_n^{st} при неограниченном

возрастании n , можно ожидать, что если на практике N значительно превосходит безразмерное отношение $\frac{\lambda}{\mu}$, то ситуации, когда для очередного посетителя не найдется свободного места на стоянке, будут относительно редки.

ВШЭ M/M/1 Банкомат

Предположим, что клиенты приходят к банкомату независимо друг от друга со средней частотой λ и каждый клиент, дождавшись своей очереди, тратит на проведение банковских операций случайное время, имеющее экспоненциальную функцию плотности вероятности с параметром μ . Какова статистически-стационарная функция распределения π_n^{st} числа клиентов возле банкомата?

Как и в предыдущей задаче, будем характеризовать состояние системы количеством клиентов, в находящихся в ней данный момент. Стохастическая динамика переходов между состояниями представляет собой случайный процесс Маркова, см. Рис. (3). Соответствующая система уравнений на эволюцию распределения вероятностей нахождения системы в различных состояниях выглядит следующим образом

$$\begin{aligned}\frac{d\pi_0(t)}{dt} &= -\lambda\pi_0(t) + \mu\pi_1(t) \\ \frac{d\pi_n(t)}{dt} &= \lambda\pi_{n-1}(t) - (\lambda + \mu)\pi_n(t) + \mu\pi_{n+1}(t), \quad \text{при } n \geq 1.\end{aligned}$$

Отсюда видно, что для стационарных вероятностей имеет место соотношение $\pi_n^{\text{st}} = \left(\frac{\lambda}{\mu}\right)^n \pi_0^{\text{st}}$. С учетом условия нормировки это означает, что стационарное решение дается следующим распределением

$$\pi_n^{\text{st}} = \left(1 - \frac{\lambda}{\mu}\right) \left(\frac{\lambda}{\mu}\right)^n.$$

Отметим, что такое решение существует только, если $\mu > \lambda$. В противном случае среднее время, затрачиваемое на обслуживание одного клиента, превосходит средний интервал времени между прибытием новых клиентов, так что длина очереди неограниченно возрастает с течением времени. Последнее отражается также в том, что среднее число клиентов в системе, $\langle n \rangle = \frac{\lambda}{\mu - \lambda}$, диктуемое распределением (53), становится бесконечно большим при стремлении μ к λ сверху.

Обозначим как T время, которое клиент, пришедший к банкомату в случайный момент времени, потратит на то, чтобы отстоять в очереди и получить услугу. Очевидна, эта случайная величина может быть записана как $T = \sum_{i=0}^n \tau_i$, где n - число человек, обнаруженное возле банкомата новым клиентом, τ_i - время, которое потратит на работу с банкоматом i -й человек в очереди. Так слагаемые в указанной сумме статистически независимы и одинаково распределены с экспоненциальной плотностью вероятности $\mu e^{-\mu\tau}$, то условная плотность вероятности $\rho(T|n)$ дается распределением Эрланга следующего вида

$$\rho(T|n) = \frac{\mu^{n+1} T^n}{n!} e^{-\mu T}$$

Тогда согласно формуле полной вероятности, функция распределения случайного времени T равна

$$\begin{aligned}P(T) &= \sum_{n=0}^{+\infty} \rho(T|n) \pi_n^{\text{st}} = e^{-\mu T} \left(1 - \frac{\lambda}{\mu}\right) \sum_{n=0}^{+\infty} \frac{\mu^{n+1} T^n}{n!} \left(\frac{\lambda}{\mu}\right)^n = \\ &= e^{-\mu T} (\mu - \lambda) \sum_{n=0}^{+\infty} \frac{(\lambda T)^n}{n!} = (\mu - \lambda) e^{-(\mu - \lambda)T}.\end{aligned}$$

Среднее время, которое клиент проведет в системе:

$$\langle T \rangle = \frac{1}{\mu - \lambda}$$

Последний ответ также можно было получить посредством формулы $\langle T \rangle = (\langle n \rangle + 1) \langle \tau \rangle$.

ВШЭ $M/M/N$

Наконец, рассмотрим модель $M/M/N$, характеризуемую произвольным количеством N каналов обслуживания. Предположим, в магазине работают N касс самообслуживания. Каждый покупатель занимает кассу на случайный промежуток времени, имеющий экспоненциальное распределение с параметром μ . Покупатели подходят к кассам в пуассоновские моменты времени со средней частотой λ , выстраиваясь в случае необходимости в единую очередь. Очередной покупатель проходит на ту кассу, которая освобождается первой, либо на любую из свободных, если таковых в момент его прихода несколько. Обозначим через n полное число покупателей в зоне оплаты покупок, включая и тех кто уже стоит на кассе, и тех, кто только ожидает своей очереди. Каково стационарное вероятностное распределение π_n^{st} ? Эволюция распределения вероятностей нахождения системы в состояниях, характеризуемых различным числом покупателей n , описывается следующей системой уравнений

$$\begin{aligned} \frac{d\pi_0(t)}{dt} &= -\lambda\pi_0(t) + \mu\pi_1(t) \\ \frac{d\pi_n(t)}{dt} &= \lambda\pi_{n-1}(t) - (\lambda + n\mu)\pi_n(t) + (n+1)\mu\pi_{n+1}(t), \quad \text{при } 1 \leq n \leq N-1, \\ \frac{d\pi_n(t)}{dt} &= \lambda\pi_{n-1}(t) - (\lambda + N\mu)\pi_n(t) + N\mu\pi_{n+1}(t), \quad \text{при } n \geq N. \end{aligned}$$

Занулив левые части этих уравнений, находим стационарное решение

$$\pi_n^{\text{st}} = \begin{cases} \frac{1}{n!} \left(\frac{\lambda}{\mu}\right)^n \pi_0^{\text{st}}, & \text{при } 0 \leq n \leq N-1 \\ \frac{1}{N!N^{n-N}} \left(\frac{\lambda}{\mu}\right)^n, & \text{при } n \geq N \end{cases}$$

где $\pi_0^{\text{st}} = \left(\frac{1}{N!} \left(\frac{\lambda}{\mu}\right)^N \frac{N\mu}{N\mu - \lambda} + \sum_{n=0}^{N-1} \frac{1}{n!} \left(\frac{\lambda}{\mu}\right)^n \right)^{-1}$ в силу условия нормировки. Легко видеть, что полученное решение переходит в распределения (50) и (53) в случаях $N \rightarrow +\infty$ и $N = 1$, соответственно.

11.3.2 Branching Processes

описание модели

Описание модели.

Первоначально задача формулировалась следующим образом.

Пусть $p_0, p_1, p_2, \dots$ - вероятности того, что отец имеет соответственно 0, 1, 2, ... сыновей, каждый из которых с теми же вероятностями может иметь своих сыновей и т.

д.

Какова вероятность, что мужская линия выродится к n -ому поколению?

В общем случае можно рассматривать однотипные частицы, под которыми могут пониматься люди, животные, бактерии или нейтроны в цепных реакциях.

Вероятности $p_k, k = 0, 1, \dots$, будут связываться с событиями, что одна частица в следующем поколении превращается в k частиц.

Будем считать, что частицы размножаются независимо друг от друга и для каждой из них вероятности $p_k, k = 0, 1, \dots$, одни и те же.

Обозначим через $\xi_0, \xi_1, \xi_2, \dots$ число частиц в нулевом, первом, втором, .

поколениях соответственно.

Другими словами, ξ_n - объем популяции в n -ом поколении.

Будем считать (без особого ограничения общности), что $\xi_0 = 1$, т.

е.

начальное поколение состоит из одной частицы.

Объем последующих поколений $\xi_1, \xi_2, \dots$ будет случайным.

Заметим, что при сделанных предположениях $\{p_0, p_1, \dots\}$ будет распределением вероятностей случайной величины ξ_1 , т.

е.

$P(\xi_1 = k) = p_k, k = 0, 1, \dots$

классификация

Поскольку частицы размножаются независимо друг от друга и число потомков каждой из них имеет одно и то же распределение вероятностей, то распределение случайной величины ξ_n должно определяться из начальных условий.

Важной характеристикой ветвящегося процесса является среднее число непосредственных потомков одной частицы $E\xi_1 = \sum_{k=0}^{\infty} kp_k = m$. В зависимости от значения этой величины ветвящийся процесс называется:

докритическим, если $m < 1$; критическим, если $m = 1$;

- надкритическим, если $m > 1$ (или $m = \infty$).

Через σ^2 будем также обозначать дисперсию $D\xi_1$. Пусть A - событие, которое состоит в том, что процесс $\xi_0 = 1, \xi_1, \xi_2, \dots$ в конце концов вырождается.

Его вероятность $q = P(A)$ будем называть вероятностью вырождения процесса.

Ближайшей нашей целью будет найти зависимость q от начальных данных.

Производящая функция процесса.

Поскольку $\xi_1, \xi_2, \dots$ являются целочисленными случайными величинами, то их распределения вероятностей естественно представлять производящими функциями.

В частности, обозначим

$$f(x) = g_{\xi_1}(x) = \sum_{k=0}^{\infty} p_k x^k$$

и будем называть f производящей функцией процесса.

Для отыскания производящих функций $g_{\xi_n}, n = 2, 3, \dots$, нам потребуется следующий результат.

Теорема

12.1.

Пусть $\eta_1, \eta_2, \dots$ - независимые одинаково распределенные зависимость от них целочисленная случайная величина с производящей функцией $S_\nu = \eta_1 + \dots + \eta_\nu (S_0 = 0)$ определяется равенством

$$g_{S_\nu}(x) = g_\nu \circ g_\eta(x)$$

ДоКАЗАТЕЛЬСТВО.

Непосредственно из определения производящей функции с использованием свойств математического ожидания получаем

$$\begin{aligned} g_{S_\nu}(x) &= E x^{S_\nu} = \sum_{k=0}^{\infty} E (x^{S_\nu} 1_{\{\nu=k\}}) \\ &= \sum_{k=0}^{\infty} E (x^{\eta_1 + \dots + \eta_k} 1_{\{\nu=k\}}) \\ &= \sum_{k=0}^{\infty} (g_\eta(x))^k P(\nu = k) = g_\nu(g_\eta(x)) \end{aligned}$$

ТЕОРЕМА 12.2.

Пусть $\xi_0 = 1, \xi_1, \xi_2, \dots$ — ветвящийся процесс Гальтона-Ватсона с производящей функцией $g_{\xi_1}(x) = f(x)$. Тогда производящая функция $g_{\xi_n}(x)$ представляет собой n -ую итерацию f^n функции f , т.

е.

$$g_{\xi_n}(x) = f^n(x) = f \circ \dots \circ f(x)$$

ДоКАЗАТЕЛЬСТВО.

Представим n -ое поколение частиц в виде $\xi_n = \eta_1 + \dots + \eta_{\xi_{n-1}}$, где η_1 — число потомков от первой частицы, η_2 — от второй, ..., $\eta_{\xi_{n-1}}$ от последней из $(n-1)$ -го поколения.

Тогда по предыдущей теореме

$$g_{\xi_n}(x) = g_{\xi_{n-1}} \circ g_\eta(x) = g_{\xi_{n-1}} \circ f(x)$$

По индукции получаем требуемый результат.

Вероятность вырождения.

Пусть A_n — событие, которое заключается в том, что n -ое поколение отсутствует, т.

е.

$A_n = \{\xi_n = 0\}$. Тогда $A_n \nearrow A = \bigcup_{n=1}^{\infty} A_n$. Поэтому в силу непрерывности вероятностной меры

$$q = P(A) = \lim_{n \rightarrow \infty} P(A_n)$$

Теорема

12.3.

Вероятность q вырождения процесса Гальтона-Ватсона с производящей функцией $f(x)$ определяется равенством

$$q = \lim_{b \rightarrow \infty} f^n(0)$$

и является наименьшим неотрицательным корнем уравнения

$$f(x) = x$$

ДоКАЗАТЕЛЬСТВО.

Поскольку $P(\xi_n = 0) = f^n(0)$, то $f^n(0) \nearrow q$. Осуществляя предельный переход при $n \rightarrow \infty$ в равенстве

$$f(f^n(0)) = f^{n+1}(0)$$

получаем $f(q) = q$, т. е. q — корень уравнения $f(x) = x$. Остается показать, что q — наименьший неотрицательный корень уравнения $f(x) = x$. Допустим противное, т.е.

$f(\alpha) = \alpha$ и $0 \leq \alpha < q$. Тогда в силу монотонности функции f будем иметь

$$f(0) \leq f(\alpha) = \alpha < q = f(q)$$

Снова, используя монотонность функции f , получаем

$$f^2(0) \leq \alpha < q$$

Повторение этого процесса приводит к неравенствам

$$f^n(0) \leq \alpha < q, \quad n = 2, 3, \dots$$

Но тогда

$$\lim_{n \rightarrow \infty} f^n(0) \leq \alpha$$

что противоречит определению q .

процесс Гальтона-Ватсона критический или докритический, то $q = 1$

Теорема

12.4.

Если процесс Гальтона-Ватсона критический или докритический, то $q = 1$.

В случае надкритического процесса $0 \leq q < 1$.

ДоКАЗАТЕЛЬСТВО.

Допустим, что $m = f'(1) \leq 1$. Если процесс невырожденный, то $f'(x) < 1$ при $0 \leq x < 1$. Фиксируем произвольно $x_0 \in [0, 1)$. По теореме о среднем найдется такое $x^* \in (x_0, 1)$, что

$$1 - f(x_0) = f'(x^*)(1 - x_0)$$

Но тогда

$$1 - f(x_0) < 1 - x_0$$

и $f(x_0) > x_0$. Т.

о.

в интервале $[0, 1)$ корней уравнение $f(x) = x$ не имеет.

Если $m = f'(1) > 1$, то найдется такое $x_0 \in (0, 1)$, что $f'(x_0) > 1$. Используя теорему о среднем, получаем

$$1 - f(x_0) > 1 - x_0$$

т.

е.

$f(x_0) - x_0 < 0$. Если $f(0) = 0$, то $q = 0$. В противном случае $f(0) > 0$ и функция $\varphi(x) = f(x) - x$ принимает на концах промежутка $[0, x_0]$ значения разных знаков.

Следовательно, в этом промежутке найдется решение уравнения $f(x) = x$

11.4 Typical Applied Probabilistic Models (!!)

(тут как раз уже такие модели, которые нужны в жизни хотя бы где-то!!)

11.4.1 Particle Walking Models

(написано в разделе про случайпроцы, тут лишь готовые методы с минимумом пояснений)

11.4.2 Fokker-planck Equation

первая тут идея - что нужно перейти к пропагатору, то есть к интегралу от вероятности в прошлом. !!! здесь необходима связка с условной вероятностью, что-то я совсем там не помню никаких таких формул.

12 Mathematical Statistics: Basics

12.1 Fundamentals of Mathematical Statistics in a Nutshell

(сперва пропишу это в части про мат статистику, потом сюда суть переносить буду. пока не до неё, так что позже займусь.)

12.1.1 Some Inequalities in Mathematical Statistics

(просто помню, что их прямо много)

12.1.2 Central Limit Theorem

(хз, чего она именно в мат статистике, но ладно)

Простейший вариант ЦПТ

(из избранных пишу)

Суть ЦПТ (???)

Представим что у нас есть N независимых случайных величин $\xi_1 \dots \xi_N$ с нулевым средним и не сильно разными дисперсиями. Поищем характеристики случайной величины $\xi = \frac{1}{N} \sum_k \xi_k$.

Вычислим её производящую функцию:

$$f(p) = \left\langle e^{-ip \frac{1}{N} \sum_i \xi_i} \right\rangle = \prod_i \left\langle e^{-ip \frac{1}{N} \xi_i} \right\rangle$$

$$f(p) = \prod_i e^{-\frac{\sigma_i^2 p^2}{2N^2}} = e^{-\frac{p^2}{2N^2} \sum_i \sigma_i^2}$$

$$\rho(\xi) = \frac{\sqrt{N}}{\sigma \sqrt{2\pi}} e^{-N \frac{\xi^2}{2\sigma^2}}$$

Получили распределение Гаусса. Видим, что типичное отклонение ξ меньше чем средне-квадратичное отклонение случайных величин в $\sqrt{N}$ раз.

Проверим что наше предположение о малости - верное. При интегрировании $\sqrt{N}$ тогда условие малости верно если $\xi \propto \sqrt{N}$.

Кажется что можно использовать условие $\xi \propto \sqrt{N}$ тогда условие малости переходит в условие $\sqrt{N}$ которое выполнено. В задаче с идеальным газом ξ_i Чтобы перейти к распределению Гаусса нам потребовалось не просто условие как так-то?

Центральная предельная теорема, как следует из её названия, очень важна. Во-первых, Она позволяет понять как ведёт себя случайная величина которая зависит от многих параметров. Хорошим примером такой величины является случайная ошибка в эксперименте. ЦПТ позволяет понять что ошибка распределена Гауссовым образом и зависит всего от одного параметра. Во-вторых,

ЦПТ говорит о том как улучшается ваш результат при повторении измерений.

строгая формулировка и доказательство Цпт

(зубков)

Ранее

Теорема

Муавра-Лапласа о сходимости распределения централизованного и нормированного числа успехов в n независимых испытаниях Бернулли к нормальному распределению была доказана с помощью довольно громоздких вычислений, а результат относился к узкому классу случаев.

С помощью характеристических функций легко доказывать значительно более общие утверждения.

Теорема. (зачем она?)

Если случайные величины $\xi_1, \xi_2, \dots$ независимы, одинаково распределены, $\mathbf{M}\xi_1 = a$, $\mathbf{D}\xi_1 = \sigma^2 \in (0, \infty)$, то

$$\lim_{n \rightarrow \infty} \mathbf{P} \left\{ \frac{\xi_1 + \dots + \xi_n - na}{\sigma\sqrt{n}} \leq x \right\} = \Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-u^2/2} du$$

Доказательство.

Согласно теореме о непрерывности соответствия между функциями распределения и характеристическими функциями, достаточно показать, что характеристические функции $f_n(t) = \mathbf{M}e^{it\zeta_n}$ случайных величин

$$\zeta_n = \frac{\xi_1 + \dots + \xi_n - na}{\sigma\sqrt{n}}$$

при $n \rightarrow \infty$ для каждого $t \in \mathbb{R}$ сходятся к характеристической функции стандартного нормального распределения, равной $\exp\{-t^2/2\}$.

Очевидно,

$$\zeta_n = \frac{\xi_1 + \dots + \xi_1 - na}{\sigma\sqrt{n}} = \frac{1}{\sqrt{n}} \sum_{k=1}^n \frac{\xi_k - a}{\sigma}$$

слагаемые в сумме в правой части независимы и имеют такое же распределение, как ζ_1 .

Поэтому

$$\mathbf{M} \exp \left\{ it \sum_{k=1}^n (\xi_k - a) \right\} = (\mathbf{M} \exp \{ it (\xi_1 - a) \})^n = f_1^n(t).$$

Так как $\mathbf{M} \frac{\xi_1 - a}{\sigma} = 0$ и $\mathbf{D} \frac{\xi_1 - a}{\sigma} = \frac{1}{\sigma^2} \mathbf{D}\xi_1 = 1$, то по свойству 6 характеристических функций

$$f_1(t) = \mathbf{M} \exp \left\{ it \frac{\xi_1 - a}{\sigma} \right\} = 1 - \frac{t^2}{2} + o(t^2), \quad t \rightarrow 0$$

Следовательно, при любом фиксированном $t \in \mathbb{R}$ и $n \rightarrow \infty$

$$\begin{aligned} \mathbf{M} \exp \{ it \zeta_n \} &= \mathbf{M} \exp \left\{ \frac{it}{\sigma\sqrt{n}} \sum_{k=1}^n (\xi_k - a) \right\} = f_1^n \left(\frac{t}{\sqrt{n}} \right) = \\ &= \left(1 - \frac{1}{2} \left(\frac{t}{\sqrt{n}} \right)^2 + o \left(\frac{t^2}{n} \right) \right)^n = \left(1 - \frac{t^2}{2n} + o \left(\frac{t^2}{2n} \right) \right)^n \rightarrow e^{-t^2/2} \end{aligned}$$

Тем самым Теорема доказана.

(и что с этого? кст, суть не понял)

Теорема Ляпунова (????)

Метод характеристических функций позволяет доказывать предельные теоремы и в более общих ситуациях.

В качестве примера рассмотрим теорему Ляпунова.

Теорема Ляпунова.

Пусть при каждом $n = 1, 2, \dots$ случайные величины $\xi_{n,1}, \dots, \xi_{n,n}$ независимы

$$\mathbf{M}\xi_{n,k} = a_{n,k}, \quad \mathbf{D}\xi_{n,k} = \sigma_{n,k}^2, \quad \mathbf{M}|\xi_{n,k} - a_{n,k}|^3 = c_{n,k}^3, \quad n, k = 1, 2, \dots$$

и $S_n = \xi_{n,1} + \dots + \xi_{n,n}$. Если

$$A_n = \sum_{k=1}^n a_{n,k}, \quad B_n^2 = \sum_{k=1}^n \sigma_{n,k}^2, \quad C_n^3 = \sum_{k=1}^n c_{n,k}^3$$

и распределения слагаемых изменяются так, что $\frac{C_n}{B_n} \rightarrow 0, n \rightarrow \infty$, то

$$\lim_{n \rightarrow \infty} \mathbf{P} \left\{ \frac{S_n - A_n}{B_n} \leq x \right\} = \Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-u^2/2} du$$

Доказательство.

Достаточно доказать, что характеристические функции случайных величин $\frac{1}{B_n}(S_n - A_n)$ в каждой точке сходятся к характеристической функции стандартного нормального распределения, равной $\exp\{-t^2/2\}$.

Аналогично случаю, когда слагаемые одинаково распределены, представим характеристическую функцию в виде произведения характеристических функций центрированных и нормированных слагаемых и используем разложения характеристических функций с явной оценкой остаточного члена:

при любом $t \in \mathbb{R}$

$$\begin{aligned} \mathbf{M} \exp \left\{ it \frac{S_n - A_n}{B_n} \right\} &= \mathbf{M} \exp \left\{ it \sum_{k=1}^n \frac{\xi_{n,k} - a_{n,k}}{B_n} \right\} = \\ &= \prod_{k=1}^n \mathbf{M} \exp \left\{ it \frac{\xi_{n,k} - a_{n,k}}{B_n} \right\} = \prod_{k=1}^n \left(1 - \frac{t^2}{2B_n^2} \sigma_{n,k}^2 + \theta_{n,k} \frac{t^3}{B_n^3} c_{n,k}^3 \right) \end{aligned}$$

где $|\theta_{n,k}| \leq \frac{1}{6}$, $k = 1, \dots, n$. Согласно неравенству Ляпунова

$$\sigma_{n,k} = \sqrt{\mathbf{M}(\xi_{n,k} - a_{n,k})^2} \leq (\mathbf{M}|\xi_{n,k} - a_{n,k}|^3)^{1/3} = c_{n,k}, \quad k = 1, \dots, n$$

По условию $\frac{C_n}{B_n} \rightarrow 0$ при $n \rightarrow \infty$, значит,

$$\max_{1 \leq k \leq n} \frac{c_{n,k}^3}{B_n^3} \leq \left(\frac{C_n}{B_n} \right)^3 \rightarrow 0, \quad \text{т. е.} \quad \max_{1 \leq k \leq n} \frac{c_{n,k}}{B_n} \rightarrow 0$$

Следовательно,

$$\max_{1 \leq k \leq n} \frac{\sigma_{n,k}}{B_n} \leq \max_{1 \leq k \leq n} \frac{c_{n,k}}{B_n} \rightarrow 0$$

Таким образом, для каждого фиксированного $t \in \mathbb{R}$

$$\max_{1 \leq k \leq n} \left(\frac{t^2}{2B_n^2} \sigma_{n,k}^2 + \theta_{n,k} \frac{t^3}{B_n^3} c_{n,k}^3 \right) \rightarrow 0 \quad \text{при } n \rightarrow \infty$$

и поэтому при $n \rightarrow \infty$ и $k = 1, \dots, n$

$$\ln \left(1 - \frac{t^2}{2B_n^2} \sigma_{n,k}^2 + \theta_{n,k} \frac{t^3}{B_n^3} c_{n,k}^3 \right) = \left(-\frac{t^2}{2B_n^2} \sigma_{n,k}^2 + \theta_{n,k} \frac{t^3}{B_n^3} c_{n,k}^3 \right) (1 + o(1))$$

где в силу (47) остаточные члены стремятся к 0 при $n \rightarrow \infty$ равномерно по $k \in \{1, \dots, n\}$.

Отсюда следует, что для любого $t \in \mathbb{R}$ при некоторых $\theta = \theta(n, t) \in \left[-\frac{1}{6}, \frac{1}{6}\right]$

$$\begin{aligned} \ln \mathbf{M} \exp \left\{ it \frac{\xi_{n,1} + \dots + \xi_{n,n} - A_n}{B_n} \right\} &= \sum_{k=1}^n \ln \left(1 - \frac{t^2}{2B_n^2} \sigma_{n,k}^2 + \theta_{n,k} \frac{t^3}{B_n^3} c_{n,k}^3 \right) = \\ &= (1 + o(1)) \sum_{k=1}^n \left(-\frac{t^2}{2B_n^2} \sigma_{n,k}^2 + \theta_{n,k} \frac{t^3}{B_n^3} c_{n,k}^3 \right) = \\ &= (1 + o(1)) \left(-t^2 \frac{B_n^2}{2B_n^2} + \theta t^3 \frac{C_n^3}{B_n^3} \right) = -\frac{t^2}{2} (1 + o(1)) \end{aligned}$$

Теорема доказана.

(?? и зачем она??)

применения

12.1.3 Characteristic Functions (?!!!!!)

(чего раньше я не вынесу эту часть?!!!!! почему это мат статика? вынесу скорее всего скоро.)

(тут подробная теория. чет её прямо дохрена. чего кст раньше об этом ничего не было?)

Конструкция характеристической функции

Основные формулы характеристической функции (?)

(пока хз)

Теория

Распространим понятие математического ожидания на комплекснозначные случайные величины. Для комплексной случайной величины $\zeta = \xi + i\eta$, где $\xi = \operatorname{Re} \zeta$ и $\eta = \operatorname{Im} \zeta$ - случайные величины с действительными значениями, причем существуют и конечны $\mathbf{M} \xi$ и $\mathbf{M} \eta$. Тогда по определению полагают

$$\mathbf{M} \zeta = \mathbf{M} \xi + i \mathbf{M} \eta$$

Таким образом, $\operatorname{Re} \mathbf{M} \zeta = \mathbf{M} \operatorname{Re} \zeta$, $\operatorname{Im} \mathbf{M} \zeta = \mathbf{M} \operatorname{Im} \zeta$

Основные свойства математического ожидания (аддитивность, мультипликативность для произведения независимых случайных величин) переносятся на математические ожидания комплекснозначных случайных величин.

Лемма

Если существует математическое ожидание комплекснозначной случайной величины ζ , то

$$|\mathbf{M} \zeta| \leq \mathbf{M} |\zeta|$$

Рассмотрим функцию $m(x + iy) = |x + iy|$, $x, y \in R$, как функцию от двух действительных переменных x и y .

Графиком функции $m(x, y)$ является конус $z = (x^2 + y^2)^{1/2}$.

Для любой точки $(a, b) \in R^2$, плоскость

$$z = z(x, y) = \frac{xa + yb}{(a^2 + b^2)^{1/2}}$$

касается конуса по лучу $(ta, tb, t(a^2 + b^2)^{1/2})$, $t \geq 0$, а для остальных точек плоскости лежит строго ниже конуса, т.е.

$$z(x, y) \leq m(x + iy) \text{ для всех } x, y \in R$$

Пусть теперь ζ - комплекснозначная случайная величина и $M\zeta = a + ib$. Тогда

$$\begin{aligned} M|\zeta| &= Mm(\zeta) \geq Mz(\operatorname{Re} \zeta, \operatorname{Im} \zeta) = \frac{M(a \operatorname{Re} \zeta + ib \operatorname{Im} \zeta)}{(a^2 + b^2)^{1/2}} = \\ &= (a^2 + b^2)^{1/2} = |a + ib| = |M\zeta| \end{aligned}$$

Определение характеристической функции

(зубков)

Характеристической функцией случайной величины ξ , принимающей действительные значения, называется функция

$$f(t) = M \exp\{it\xi\}, \quad -\infty < t < \infty$$

Основные свойства характеристической функции (?)

(!! надо бы посмотреть это!!)

Перечислим несколько простых свойств характеристических функций.

1. $|f(t)| \leq 1$ при всех $t \in R$, $f(0) = 1$

Первое неравенство - простое следствие леммы, так как $|\exp\{it\xi\}| = 1$ и $|f(t)| = M \exp\{it\xi\} \leq M|\exp\{it\xi\}| = 1$. Второе равенство очевидно.

2. Если $\eta = a\xi + b$, где a, b - константы, то

$$f_\eta(t) = e^{itb} f_\xi(at)$$

Действительно,

$$f_\eta(t) = M \exp\{it\eta\} = M \exp\{it(a\xi + b)\} = e^{itb} M \exp\{ita\xi\} = e^{itb} f_\xi(at)$$

3. $f_\gamma(-t) = \overline{f_\gamma(t)}$, где черта - знак комплексного сопряжения.

Действительно: $f_\gamma(-t) = M \exp\{i(-t)\gamma\} = M \exp\{-it\gamma\} = M \exp\{it\gamma\} = \overline{M \exp\{it\gamma\}} = \overline{f_\gamma(t)}$

Следствие.

Если распределение случайной величины γ симметрично, т.е. γ и $-\gamma$ одинаково распределены, то

$$f_\gamma(t) = M \exp\{it\gamma\} = M \exp\{-it\gamma\} = f_\gamma(-t) = \overline{f_\gamma(t)}$$

т. е. характеристическая функция случайной величины с симметричным распределением принимает только действительные значения.

4. Если случайные величины $\gamma_1, \dots, \gamma_n$ независимы, то

$$f_{\gamma_1 + \dots + \gamma_n}(t) = \prod_{k=1}^n f_{\gamma_k}(t)$$

Так как $\gamma_1, \dots, \gamma_n$ независимы, то случайные величины $\exp\{it\gamma_1\}, \dots, \exp\{it\gamma_n\}$ тоже независимы, и поэтому

$$\begin{aligned} f_{\gamma_1 + \dots + \gamma_n}(t) &= \mathbf{M} \exp\{it(\gamma_1 + \dots + \gamma_n)\} = \mathbf{M} \prod_{k=1}^n \exp\{it\gamma_k\} = \\ &= \prod_{k=1}^n \mathbf{M} \exp\{it\gamma_k\} = \prod_{k=1}^n f_{\gamma_k}(t) \end{aligned}$$

свойства характеристических функций

(горяинов)

1°. $|h_{\xi}(t)| \leq 1$ и $h_{\xi}(0) = 1$

2°. $h_{\xi}(-t) = \overline{h_{\xi}(t)}$

3°. $h_{a\xi+b}(t) = e^{itb} h_{\xi}(at)$, $a, b \in \mathbb{R}$

4°. Если $\xi_1, \dots, \xi_n$ - независимые случайные величины и $S_n = \xi_1 + \dots + \xi_n$ то $h_{S_n}(t) = \prod_{k=1}^n h_{\xi_k}(t)$

5°. Функция $h_{\xi}(t)$ равномерно непрерывна на $\mathbb{R}$ и если $\mathbf{M}|\xi|^n < \infty$, то $h_{\xi}(t)$ имеет производные до n -го порядка включительно и

$$\mathbf{M}\xi^n = \frac{1}{i^n} h_{\xi}^{(n)}(0)$$

Свойства 1° – 4° следуют непосредственно из определения характеристической функции и свойств математического ожидания.

Доказательство последнего свойства опирается на теорему Лебега о предельном переходе под знаком интеграла.

Теорема

(Бохнера)

Для того, чтобы непрерывная на $\mathbb{R}$ функция $h(t)$ была характеристической функцией некоторого распределения, необходимо и достаточно, чтобы она была положительно определенной, т. е. для любых $t_1, \dots, t_n$ из $\mathbb{R}$ и $z_1, \dots, z_n$ из $\mathbb{C}$, $n = 1, 2, \dots$, выполнялось условие

$$\sum_{k,l=1}^n h(t_k - t_l) z_k \bar{z}_l \geq 0$$

Доказательство

Необходимость условия следует из простых преобразований.

$$\begin{aligned} \sum_{k,l} h_{\xi}(t_k - t_l) z_k \bar{z}_l &= \mathbf{M} \left(\sum_{k,l} z_k \bar{z}_l e^{i(t_k - t_l)\xi} \right) \\ &= \mathbf{M} \left(\sum_k z_k e^{it_k \xi} \right) \overline{\left(\sum_l z_l e^{it_l \xi} \right)} \\ &= \mathbf{M} \left| \sum_k z_k e^{it_k \xi} \right|^2 \geq 0 \end{aligned}$$

Доказательство достаточности условия теоремы выходит за рамки данного курса.

Характеристические функции для простейших распределений

а) пусть имеет симметричное двухточечное распределение:

$\mathbf{P}\{\xi = a\} = \mathbf{P}\{\xi = -a\} = \frac{1}{2}$. Тогда

$$\mathbf{M} \exp\{it\xi\} = \mathbf{P}\{\xi = a\}e^{ita} + \mathbf{P}\{\xi = -a\}e^{-ita} = \frac{(e^{ita} + e^{-ita})}{2} = \cos(at)$$

б) Пусть ξ имеет равномерное распределение на $[-a, a]$:

$$p_{\xi}(x) = \frac{1}{2a} \text{ при } x \in [-a, a] \text{ и } p_{\xi}(x) = 0 \text{ при } |x| > a$$

Тогда

$$\begin{aligned} \mathbf{M} e^{it\xi} &= \int_{-a}^a e^{itx} \frac{1}{2a} dx = \frac{1}{2a} \int_{-a}^a (\cos(tx) + i \sin(tx)) dx = \\ &= \frac{1}{2a} t (\sin(tx) - i \cos(tx)) \Big|_{-a}^a = \frac{\sin(at)}{at} \end{aligned}$$

в) Пусть ξ имеет стандартное нормальное распределение с плотностью $\varphi(x) = \frac{1}{\sqrt{2\pi}} \left(\frac{-x^2}{2} \right)$. Так как плотность $\varphi(x)$ - четная функция, то

$$\begin{aligned} f(t) &= \mathbf{M} \exp\{it\xi\} = \int_R e^{itx} \varphi(x) dx = \\ &= \int_R (\cos(tx) + i \sin(tx)) \varphi(x) dx = \int_R \varphi(x) \cos(tx) dx \end{aligned}$$

Составим уравнение для производной $f(t)$

$$\begin{aligned} f'(t) &= \frac{d}{dt} \int_R \varphi(x) \cos(tx) dx = -\frac{1}{\sqrt{2\pi}} \int_R x \sin(tx) e^{-x^2/2} dx = \\ &= -\frac{1}{\sqrt{2\pi}} \int_R \sin(tx) d(e^{x^2/2}) = \\ &= -\frac{1}{\sqrt{2\pi}} \left(\sin(tx) e^{-x^2/2} \Big|_{-\infty}^{\infty} - t \int_R \cos(tx) e^{-x^2/2} dx \right) = -tf(t) \end{aligned}$$

Дифференциальное уравнение $f'(t) = -tf(t)$ с начальным условием $f(0) = 1$ решается методом разделения переменных:

$$\frac{f'(t)}{f(t)} = (\ln f(t))' = -t \rightarrow \ln f(t) = \frac{-t^2}{2}$$

Т.е. $f(t) = \exp\left(-\frac{t^2}{2}\right), t \in R$

Если случайная величина ξ имеет стандартное нормальное распределение с математическим ожиданием 0 и дисперсией 1, то случайная величина $\eta = \sigma\xi + a$ имеет нормальное распределение с параметрами (a, σ^2) (математическим ожиданием a и дисперсией σ^2). По свойству 2) получаем:

$$\mathbf{M} \exp(it\eta) = e^{ita} f(\sigma t) = \exp\left(ita - \frac{\sigma^2 t^2}{2}\right)$$

Случайная величина η имеет функцию распределения $\mathbf{P}\{\eta \leq x\} = \Phi\left(\frac{x-a}{\sigma}\right)$ и плотность распределения $\varphi_{a, \sigma^2}(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{(x-a)^2}{2\sigma^2}\right\}$.

дальнейшие свойства

Как и функция распределения, характеристическая функция существует для любой случайной величины с действительными значениями.

Более того, характеристические функции обладают рядом аналитических свойств.

Лемма

2.

Для любых $t \in R$ и целых неотрицательных n

$$\left| e^{it} - \sum_{k=0}^{n-1} \frac{(it)^k}{k!} \right| \leq \frac{|t|^n}{n!}$$

Доказательство проведем индукцией по n . При $n = 0$ имеем

$$|e^{it} - 1| \leq |t| \text{ при всех } t \in R$$

Допустим, что для некоторого целого неотрицательного n утверждение леммы справедливо, и сделаем индуктивный переход. Заметим, что

$$\frac{d}{dt} \left(e^{it} - \sum_{k=0}^n \frac{(it)^k}{k!} \right) = ie^{it} - \sum_{k=0}^n \frac{ik(it)^{k-1}}{k!} = i \left(e^{it} - \sum_{k=0}^{n-1} \frac{(it)^k}{k!} \right)$$

и что при $t = 0$ правая часть равна 0. Поэтому

$$\left| e^{it} - \sum_{k=0}^n \frac{(it)^k}{k!} \right| = \left| i \int_0^t \left(e^{iu} - \sum_{k=0}^{n-1} \frac{(iu)^k}{k!} \right) du \right|$$

Оценивая подынтегральное выражение с помощью индуктивного предположения, находим

$$\left| e^{it} - \sum_{k=0}^n \frac{(it)^k}{k!} \right| = \left| \int_0^t \left(e^{iu} - \sum_{k=0}^{n-1} \frac{(iu)^k}{k!} \right) du \right| \leq \int_0^t \frac{|u|^n}{n!} du = \frac{|t|^{n+1}}{(n+1)!}$$

Лемма доказана.

5.

Лемма

3. Пусть случайная величина ξ имеет функцию распределения $F(x) = \mathbf{P}\{\xi \leq x\}$, $x \in R$. Положим $\mathbf{M}\xi^n = m_n$. Если $|m_n| < \infty$, то существуют производные $f^{(k)}(t)$ при $k = 1, \dots, n$ и

$$f^{(k)}(0) = i^k \mathbf{M}\xi^k = i^k m_k, \quad k = 1, \dots, n$$

более того, имеют место разложения

$$f(t) = \sum_{k=0}^n m_k \frac{(it)^k}{k!} + R_n(t), \quad R_n(t) = o(t^n), \quad t \rightarrow 0$$

$$f(t) = \sum_{k=0}^{n-1} m_k \frac{(it)^k}{k!} + r_n(t), \quad |r_n(t)| \leq \mathbf{M}|\xi|^n \cdot \frac{|t|^n}{n!}$$

Доказательство. Формальным дифференцированием равенства

$$f(t) = \mathbf{M} \exp\{it\xi\} = \int_R e^{itx} dF(x)$$

Получаем:

$$f^{(k)}(t) = \int_R (ix)^k \exp\{itx\} dF(x)$$

законность дифференцирования следует из того, что полученный интеграл сходится абсолютно при любом t :

$$\int_R |(ix)^k \exp\{itx\}| dF(x) = \int_R |x|^k dF(x) = m_k < \infty$$

Полагая $t = 0$ в (46), доказываем первое утверждение:

$$f^{(k)}(0) = \int_R (ix)^k dF(x) = i^k m_k, \quad k = 1, \dots, n$$

После того, как доказана n -кратная дифференцируемость $f(t)$, для доказательства равенства (44) достаточно использовать формулу Тейлора с остаточным членом в форме Пеано. Оценка остаточного члена в (45) следует из разложения

$$\mathbf{M} \exp\{it\xi\} = \mathbf{M} \left[\sum_{k=0}^{n-1} \frac{(it\xi)^k}{k!} + \rho_n(t\xi) \right] = \sum_{k=0}^{n-1} \frac{m_k (it)^k}{k!} + \mathbf{M}\rho_n(t\xi)$$

где $\rho_n(x) = e^{ix} - \sum_{k=0}^{n-1} \frac{(ix)^k}{k!}$ и по лемме $|\rho_n(t\xi)| \leq \frac{|t\xi|^n}{n!}$. Так как для комплекснозначной случайной величины ζ справедлива оценка $|\mathbf{M}\zeta| \leq \mathbf{M}|\zeta|$, то

$$|r_n(t)| = |\mathbf{M}\rho_n(t\xi)| \leq \mathbf{M}|\xi|^n \cdot \frac{|t|^n}{n!}$$

Важными частными случаями соотношения (44) являются

$$f(t) = f(0) + tf'(0) + o(t) = 1 + it\mathbf{M}\xi + o(t)$$

$$f(t) = f(0) + tf'(0) + t^2 \frac{f''(0)}{2!} + o(t) =$$

$$= 1 + it\mathbf{M}\xi - \frac{t^2 \mathbf{M}\xi^2}{2} + o(t^2), \quad t \rightarrow 0$$

если $\mathbf{M}\xi = 0$, то $\mathbf{M}\xi^2 = \mathbf{D}\xi = \sigma^2$ и

$$\begin{aligned} f(t) &= 1 - \frac{t^2\sigma^2}{2} + o(t^2) \quad (t \rightarrow 0) \\ f(t) &= 1 - \frac{t^2\sigma^2}{2} + \Theta t^3 \mathbf{M}|\xi|^3, \quad |\Theta| \leq \frac{1}{6} \end{aligned}$$

Следующие два свойства лежат в основе почти всех применений аппарата характеристических функций в теории вероятностей.

6. Теорема.

Различным распределениям вероятностей соответствуют различные характеристические функции.

Доказательство этой теоремы можно найти в учебниках Б.А.Севастьянова или Боровкова.

Пример. Пусть $\xi_1, \dots, \xi_n$ независимые случайные величины, имеющие нормальные распределения: ξ_k имеет нормальное распределение с параметрами (a_k, σ_k^2) , $k = 1, \dots, n$; им соответствуют характеристические функции

$$f_k(t) = \mathbf{M}e^{it\xi_k} = \exp \left\{ iat - \frac{\sigma_k^2 t^2}{2} \right\}, k = 1, \dots, n$$

Тогда сумма $\zeta = \xi_1 + \dots + \xi_n$ имеет характеристическую функцию

$$f(t) = \mathbf{M}e^{it(\xi_1 + \dots + \xi_n)} = \prod_{k=1}^n \mathbf{M}e^{it\xi_k} = \exp \left\{ i(a_1 + \dots + a_n)t - \frac{1}{2}(\sigma_1^2 + \dots + \sigma_n^2)t^2 \right\}$$

Т. е. ξ имеет нормальное распределение с параметрами $(a_1 + \dots + a_n, \sigma_1^2 + \dots + \sigma_n^2)$. Таким образом, (согласно теореме единственности) сумма независимых случайных величин, имеющих нормальные распределения, имеет нормальное распределение, параметры которого являются суммами параметров слагаемых.

Случайные величины - это функции на пространстве элементарных событий. В математическом анализе рассматриваются разные виды сходимости функций: поточечная, равномерная, почти всюду, в том или ином функциональном пространстве и т.д. В теории вероятностей тоже рассматриваются разные виды сходимости случайных величин. Пока рассмотрим только сходимость распределений (слабую сходимость).

Слабая сходимость. Определение. Последовательность $\xi_1, \xi_2, \dots$ случайных величин с действительными значениями слабо сходится к случайной величине ξ , если функции распределения $F_n(x) = \mathbf{P}\{\xi_n \leq x\}$ сходятся к функции распределения $F(x) = \mathbf{P}\{\xi \leq x\}$ в каждой точке непрерывности $F(x)$

В общем случае последовательность случайных величин $\{\xi_n\}$ со значениями в измеримом пространстве $(B, \mathcal{B})$ с метрикой ρ называют слабо сходящейся к случайной величине ξ , если для любой непрерывной ограниченной функции $g: B \rightarrow \mathbb{R}$

$$\lim_{n \rightarrow \infty} \mathbf{M}g(\xi_n) = \mathbf{M}g(\xi)$$

Пример. При слабой сходимости функции распределения могут не сходить к предельной функции в точках ее разрыва. Пусть случайные величины ξ_n - двухточечные:

$$\mathbf{P}\left\{\xi_n = -\frac{1}{n}\right\} = p_n, \quad \mathbf{P}\left\{\xi_n = \frac{1}{n}\right\} = 1 - p_n, \quad n = 1, 2, \dots$$

а $\mathbf{P}\{\xi = 0\} = 1$. Тогда при любых значениях p_n последовательность функций распределения

$$F_n(x) = \mathbf{P}\{\xi_n \leq x\} = \begin{cases} 0, & x < -\frac{1}{n} \\ p_n, & -\frac{1}{n} \leq x < \frac{1}{n} \\ 1, & x \geq \frac{1}{n} \end{cases}$$

при $n \rightarrow \infty$ сходится к функции распределения $F(x) = 0$ при $x < 0$, $F(x) = 1$ при $x \geq 1$ при каждом $x \neq 0$, но последовательность $F_n(0) = p_n$ может (в зависимости от выбора значений p_n) сходиться к любому числу из отрезка $[0, 1]$, а может не иметь предела.

7.

Теорема

непрерывности. а) Если последовательность F_n распределений случайных величин γ_n слабо сходится при $n \rightarrow \infty$ к распределению F случайной величины $f_n(t) = \mathbf{M}e^{it\gamma_n}$ сходится к характеристической функции $f(t)$ распределения F при каждом t , и эта сходимость равномерна на любом конечном интервале значений t . б) Обратно, если последовательность характеристических функций $f_n(t)$ при каждом t сходится к функции $f(t)$, непрерывной в нуле, то последовательность соответствующих распределений F_n слабо сходится к распределению F и $f(t)$ есть характеристическая функция распределения F .

Доказательство теоремы непрерывности, как и доказательство теоремы единственности, можно найти в учебниках Б.А.Севастьянова и А.А.Боровкова.

Приведем пример, показывающий, что при нарушении условия непрерывности в 0 в п.б) слабой сходимости может не быть. Пусть ξ_n имеет равномерное распределение на отрезке $[-n, n]$: $\mathbf{P}\{\xi_n \leq x\} = \frac{1}{2} + \frac{x}{2n}$, $|x| \leq n$, $n = 1, 2, \dots$; тогда (см. приведенный в начале параграфа пример б))

$$f_n(t) = \mathbf{M}e^{it\xi_n} = \frac{\sin(nt)}{nt}$$

и $f_n(0) = 1$ при всех $n = 1, 2, \dots$, но $f_n(t) \rightarrow 0$ при $n \rightarrow \infty$ для любого $t \neq 0$. Пределом функций распределения $F_n(t) = \mathbf{P}\{\xi_n \leq x\} = \frac{1}{2} + \frac{x}{2n}$ при $n \rightarrow \infty$ является функция, тождественно равная $\frac{1}{2}$, не являющаяся функцией распределения собственной случайной величины (которая с вероятностью 1 принимает конечные значения).

Следствие.

Непрерывная функция, являющаяся поточечным пределом последовательности характеристических функций, сама является характеристической.

Теорема

о непрерывности позволяет сводить доказательство слабой сходимости случайных величин или их функций распределения к доказательству поточечной сходимости характеристических функций. В ряде случаев такое доказательство оказывается проще непосредственного.

примеры

Приведем теперь характеристические функции некоторых часто встречающихся в приложениях распределений.

Врожденное распределение $\mathbf{P}(\xi = a) = 1$. Непосредственно из определения характеристической функции получаем ее вид

$$h_\xi(t) = \mathbf{M}e^{it\xi} = e^{iat}$$

Равномерное на отрезке распределение с плотностью

$$f_\xi(x) = \frac{1}{b-a} \mathbb{1}_{[a,b]}(x)$$

Ее характеристическая функция также получается непосредственно из определения. В случае $t \neq 0$ имеем

$$h_{\xi}(t) = \frac{1}{b-a} \int_a^b e^{itx} dx = \frac{e^{itx}}{it(b-a)} \Big|_{x=a}^{x=b} = \frac{e^{itb} - e^{ita}}{it(b-a)}$$

При $t = 0$ имеем $h_{\xi}(0) = 1$

Экспоненциальное распределение зависит от параметра $\lambda > 0$ и определяется плотностью

$$f_{\xi}(x) \lambda e^{-\lambda x} 1_{[0, \infty)}(x)$$

Здесь снова воспользуемся определением характеристической функции

$$h_{\xi}(t) = \lambda \int_0^{\infty} e^{itx} e^{-\lambda x} dx = \frac{\lambda e^{(it-\lambda)x}}{it-\lambda} \Big|_{x=0}^{x=\infty} = \frac{\lambda}{\lambda - it}$$

Нормальное распределение, определяемое плотностью

$$f_{\xi}(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-a)^2/2\sigma^2}$$

с двумя параметрами $a \in \mathbb{R}$ и $\sigma > 0$, имеет важное значение как в теории, так и в приложениях. Для вычисления характеристической функции нормального распределения вначале рассмотрим случай, когда $a = 0$ и $\sigma = 1$. В этом случае говорят, что случайная величина ξ имеет стандартное нормальное распределение. При этом

$$\begin{aligned} h_{\xi}(t) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{itx} e^{-x^2/2} dx \\ &= \frac{1}{\sqrt{2\pi}} \left[\int_{-\infty}^{\infty} \cos tx \cdot e^{-x^2/2} dx + i \int_{-\infty}^{\infty} \sin tx \cdot e^{-x^2/2} dx \right] \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \cos tx \cdot e^{-x^2/2} dx \end{aligned}$$

поскольку интеграл в мнимой части берется по симметричному промежутку от нечетной функции. Замечая, что выполнены условия дифференцирования несобственного интеграла, зависящего от параметра, получаем

$$\begin{aligned} h'_{\xi}(t) &= -\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} x \sin tx \cdot e^{-x^2/2} dx = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \sin tx \cdot d e^{-x^2/2} \\ &= -\frac{t}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \cos tx \cdot e^{-x^2/2} dx = -t h_{\xi}(t) \end{aligned}$$

Таким образом, характеристическая функция стандартного нормального распределения является решением дифференциального уравнения

$$h'_{\xi}(t) = -t h_{\xi}(t)$$

Это вместе с условием $h_\xi(0) = 1$ приводит к ее виду

$$h_\xi(t) = e^{-t^2/2}$$

Пусть теперь $\sigma > 0$ и $a \in \mathbb{R}$. Случайная величина $\eta = \sigma\xi + a$ будет иметь, в силу свойства 3°, характеристическую функцию

$$h_\eta(t) = e^{iat} e^{-\sigma^2 t^2/2}$$

С другой стороны,

$$F_\eta(x) = \mathbf{P}(\sigma\xi + a \leq x) = \mathbf{P}\left(\xi \leq \frac{x-a}{\sigma}\right) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{(x-a)/\sigma} e^{-u^2/2} du$$

Т. е.

$$f_\eta(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-a)^2}{2\sigma^2}}$$

и η имеет нормальное распределение с параметрами a, σ^2 . Зная вид характеристической функции, легко вычислить математическое ожидание и дисперсию случайной величины с использованием свойства 5°. В частности, если ξ имеет стандартное нормальное распределение, то из равенств

$$h'_\xi(t) = -th_\xi(t), \quad h''_\xi(t) = -h_\xi(t) - th'_\xi(t)$$

следует, что $h'_\xi(0) = 0, h''_\xi(0) = -1$. Отсюда получаем

$$\mathbf{M}\xi = 0, \quad \mathbf{D}\xi = \mathbf{M}\xi^2 = 1$$

Для произвольного нормального распределения, когда $\eta = \sigma\xi + a$, получаем

$$\mathbf{M}\eta = a, \quad \mathbf{D}\eta = \sigma^2$$

Таким образом, параметры нормального распределения имеют вероятностный смысл.

Целочисленная случайная величина с производящей функцией $g_\xi(x)$ имеет характеристическую функцию $h_\xi(t) = g_\xi(e^{it})$. Это следует непосредственно из определений характеристической и производящей функций.

Важным моментом в применении характеристических функций является то, что они вполне определяют распределение случайной величины. Приведем два результата, известные как формулы обращения. Напомним вначале, что распределение случайной величины ξ представляет собой вероятностную меру $\mathbf{P}_\xi$ определенную на σ -алгебре борелевских множеств, т. е. на измеримом пространстве $(\mathbb{R}, \mathcal{B})$. Эта мера связана с функцией распределения F_ξ равенствами

$$F_\xi(x) = \mathbf{P}_\xi((-\infty, x]), \quad \mathbf{P}_\xi((a, b]) = F_\xi(b) - F_\xi(a)$$

$$\begin{aligned} \lim_{R \rightarrow \infty} \frac{1}{2\pi} \int_{-R}^R \frac{e^{-iat} - e^{-ibt}}{it} h_\xi(t) dt &= \frac{1}{2} [\mathbf{P}_\xi([a, b)) + \mathbf{P}_\xi((a, b])] \\ &= \frac{1}{2} [F_\xi(b-0) - F_\xi(a-0) + F_\xi(b) - F_\xi(a)] \end{aligned}$$

Теорема

Пусть характеристическая функция h_ξ абсолютно интегрируема, т. е.

$$\int_{-\infty}^{\infty} |h_\xi(t)| dt < \infty$$

Тогда ξ имеет абсолютно непрерывное распределение, а ее плотность определяется по формуле

$$f_\xi(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-ixt} h_\xi(t) dt$$

Из формулы обращения следует, что между множеством функций распределения F и множеством характеристических функций h имеет место взаимно-однозначное соответствие. Это соответствие является даже непрерывным относительно соответствующих сходимостей.

(!!) Определение слабой сходимости

Будем говорить, что последовательность функций распределения $F_n, n = 1, 2, \dots$, слабо сходится к функции распределения F и писать $F_n \rightarrow F$, если $F_n(x) \rightarrow F(x)$ при $n \rightarrow \infty$ в каждой точке непрерывности функции F .

Теорема

9.4. Пусть $F_1, F_2, \dots$ - функции распределения, а $h_1, h_2, \dots$ - соответствующие им характеристические функции. Тогда имеют место следующие утверждения:

(i) Если $F_n \rightarrow F$ и h - характеристическая функция, соответствующая функции распределения F , то $h_n(t) \rightarrow h(t)$ при $n \rightarrow \infty$ для всех $t \in \mathbb{R}$ (ii) Если $h_n(t) \rightarrow h(t)$ при $n \rightarrow \infty$ для всех $t \in \mathbb{R}$ и функция h непрерывна в точке $t = 0$, то h является характеристической функцией некоторой функции распределения F и $F_n \rightarrow F$.

Важную роль в приложениях играют результаты, утверждающие нормальность распределения сумм независимых случайных величин. Обычно их объединяют в одну группу с названием центральной предельной теоремы. Мы рассмотрим один из вариантов центральной предельной теоремы.

Теорема

9.5. Пусть $\xi_1, \xi_2, \dots$ - последовательность независимых одинаково распределенных случайных величин с конечными $\mathbf{M}\xi = a$ и $\mathbf{D}\xi = \sigma^2$. Тогда для $S_n = \xi_1 + \dots + \xi_n$ выполняется следующее предельное соотношение

$$\lim_{n \rightarrow \infty} \mathbf{P} \left(\frac{S_n - na}{\sigma\sqrt{n}} \leq x \right) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-u^2/2} du$$

для всех $x \in \mathbb{R}$ ДОКАЗАТЕЛЬСТВО. Обозначим $\zeta_n = (S_n - na) / \sigma\sqrt{n}, n = 1, 2, \dots$ Тогда утверждение теоремы эквивалентно тому, что $F_{\zeta_n} \rightarrow \Phi$, где Φ - функция Лапласа, т. е. функция распределения случайной величины η , имеющей стандартное нормальное распределение. В силу теоремы сходимости это эквивалентно условию $h_{\zeta_n}(t) \rightarrow h_\eta(t)$

при $n \rightarrow \infty$ для всех $t \in \mathbb{R}$. Рассмотрим случайные величины $\hat{\xi}_n = (\xi_n - a)/\sigma, n = 1, 2, \dots$. Очевидно, что $\mathbf{M}\hat{\xi}_n = 0$ и $\mathbf{D}\hat{\xi}_n = \mathbf{M}\hat{\xi}_n^2 = 1$. При этом

$$\zeta_n = \frac{\hat{\xi}_1}{\sqrt{n}} + \dots + \frac{\hat{\xi}_n}{\sqrt{n}}$$

Пусть

$$h(t) = h_{\hat{\xi}_n}(t) = \mathbf{M}e^{it\hat{\xi}_n}$$

Поскольку $h(0) = 1, h'(0) = i\mathbf{M}\hat{\xi}_n = 0, h''(0) = -\mathbf{M}\hat{\xi}_n^2 = -1$, То

$$h(t) = 1 - \frac{1}{2}t^2 + o(t^2)$$

Но тогда

$$h_{\zeta_n}(t) = \left(h_{\hat{\xi}_n/\sqrt{n}}(t)\right)^n = \left(h\left(\frac{t}{\sqrt{n}}\right)\right)^n = \left(1 - \frac{1}{2}\frac{t^2}{n} + o\left(\frac{t^2}{n}\right)\right)^n$$

При фиксированном $t \in \mathbb{R}$ имеем $t^2/n \rightarrow 0$ при $n \rightarrow \infty$ и

$$\lim_{n \rightarrow \infty} h_{\zeta_n}(t) = \lim_{n \rightarrow \infty} e^{n \ln\left(1 - \frac{t^2}{2n} + o\left(\frac{t^2}{n}\right)\right)} = e^{-t^2/2}$$

Замечая, что $h_{\eta}(t) = e^{-t^2/2}$, приходим к утверждению теоремы. хов в схеме Бернулли из n независимых испытаний с вероятностью $p, 0 < p < 1$, успеха в отдельном испытании. Тогда для $-\infty < a < b < \infty$ выполняется предельное соотношение

$$\lim_{n \rightarrow \infty} \mathbf{P}\left(a \leq \frac{S_n - np}{\sqrt{npq}} \leq b\right) = \Phi(b) - \Phi(a)$$

где $q = 1 - p$

Доказательство. Заметим, что S_n можно представить как сумму независимых случайных величин ξ_k , которые являются индикаторами успеха в k -ом испытании: $\mathbf{P}(\xi_k = 1) = p, \mathbf{P}(\xi_k = 0) = q, k = 1, 2, \dots$. Тогда $\mathbf{M}\xi_k = p, \mathbf{D}\xi_k = pq$. Применение центральной предельной теоремы дает требуемый результат.

Характеристические функции распределений случайных векторов

(зубков, быстро разберу)

Характеристические функции распределений случайных векторов вводятся аналогично скалярному случаю с заменой произведения аргумента характеристической функции и случайной величины скалярным произведением векторного аргумента характеристической функции и случайного вектора.

Характеристическая функция распределения случайного вектора

Определение.

Характеристической функцией распределения случайного вектора $\bar{\gamma} = (\gamma_1, \dots, \gamma_d)$, принимающего значения в $\mathbb{R}^d$, называется

$$f(t) = \mathbf{M}e^{i(\bar{t}, \bar{\gamma})}, \quad \bar{t} \in \mathbb{R}^d$$

О случаях применения (???)

(?? потом разберу!!)

Свойства

Свойства характеристических функций и их доказательство аналогичны одномерным случайным величинам:

1) $|f(\bar{t})| \leq 1$ при всех $\bar{t} \in R^d$, $f(\bar{0}) = 1$.

Если $d = 1$ и $f(t_0) = \mathbf{M}e^{it_0\gamma} = e^{i\lambda}$ при некоторых $t_0 \neq 0, \lambda \in R$, то в силу выпуклости окружности $\{e^{i\beta}, \beta \in [0, 2\pi)\}$

$$\mathbf{P}\{t_0\gamma \in \{\lambda + 2k\pi, k \in \mathbb{Z}\}\} = 1$$

т. е. распределение γ сосредоточено на арифметической прогрессии $\{t_0^{-1}\lambda + 2kt_0^{-1}\pi, k \in \mathbb{Z}\}$

Аналогично, если $d > 1$ и $f(\bar{t}_0) = \mathbf{M}e^{i(\bar{t}_0, \bar{\gamma})} = e^{i\lambda}$ при некотором $\bar{t}_0 \neq \bar{0}$, то

$$\mathbf{P}\{(\bar{t}_0, \bar{\gamma}) \in \{\lambda + 2\pi k, k \in \mathbb{Z}\}\} = 1$$

т. е. распределение случайного вектора $\bar{\gamma}$ сосредоточено на семействе параллельных гиперплоскостей вида $\{\bar{x} \in \mathbb{R}^d : (\bar{t}_0, \bar{x}) = \lambda + 2\pi k\}, k \in \mathbb{Z}$.

2) Если $\bar{\gamma}$ - случайный вектор (столбец) и $\bar{\kappa} = B\bar{\gamma} + \bar{a}$, где $B = \|b_{ij}\|_{i,j=1}^d, \bar{a} \in R^d$ то

$$f_{\bar{\kappa}}(\bar{t}) = \mathbf{M}e^{i(\bar{t}, \bar{\kappa})} = e^{i(\bar{t}, \bar{a})} f_{\bar{\gamma}}(B^T \bar{t}),$$

где ^{T} - знак транспонирования.

Действительно, так как

$$(\bar{t}, B\bar{\gamma}) = \sum_{i=1}^d t_i \sum_{j=1}^d b_{ij} \gamma_j = \sum_{j=1}^d \left(\sum_{i=1}^d b_{ij} t_i \right) \gamma_j = (B^T \bar{t}, \bar{\gamma})$$

то

$$f_{\bar{\kappa}}(\bar{t}) = \mathbf{M}e^{i(\bar{t}, \bar{\kappa})} = \mathbf{M}e^{i(\bar{t}, B\bar{\gamma} + \bar{a})} = e^{i(\bar{t}, \bar{a})} \mathbf{M}e^{i(\bar{t}, B\bar{\gamma})} = e^{i(\bar{t}, \bar{a})} \mathbf{M}e^{i(B^T \bar{t}, \bar{\gamma})} = e^{i(\bar{t}, \bar{a})} f_{\bar{\gamma}}(B^T \bar{t})$$

Пример

Случайная величина ζ имеет стандартное нормальное распределение с параметрами $(0,1)$, если ее плотность распределения равна $\varphi(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}, -\infty < x < \infty$; функция стандартного нормального распределения есть

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-u^2/2} du$$

а его характеристическая функция

$$\psi(t) = \mathbf{M}e^{it\zeta} = e^{-t^2/2}, \quad -\infty < t < \infty$$

Среди вероятностных распределений в $\mathbb{R}^d$ аналогом стандартного нормального распределения является распределение случайного вектора $\bar{\zeta} = (\zeta_1, \dots, \zeta_d)$, компоненты которого - независимые случайные величины, имеющие стандартное нормальное распределение. Такой случайный вектор имеет нулевой вектор математических ожиданий и единичную матрицу ковариаций:

$$\mathbf{M}\bar{\zeta} = \bar{0}, \quad \text{Cov}(\bar{\zeta}) = \|\text{cov}(\zeta_i, \zeta_j)\| = \begin{vmatrix} 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & 1 \end{vmatrix}$$

Плотность и характеристическая функция распределения вектора $\bar{\zeta} = (\zeta_1, \dots, \zeta_d)$ имеют вид

$$\varphi(x_1, \dots, x_d) = \prod_{k=1}^d \varphi(x_k) = \frac{1}{(2\pi)^{d/2}} e^{-\frac{x_1^2 + \dots + x_d^2}{2}}$$

$$\psi(t_1, \dots, t_d) = \mathbf{M} e^{i(\bar{t}, \bar{\zeta})} = \mathbf{M} e^{i(t_1 \zeta_1 + \dots + t_d \zeta_d)} = \mathbf{M} \prod_{k=1}^d e^{it_k \zeta_k} = e^{-\frac{t_1^2 + \dots + t_d^2}{2}}$$

Плотность $\varphi(\bar{x})$ распределения $\bar{\zeta}$ зависит только от расстояния точки $\bar{x} = (x_1, \dots, x_d)$ от начала координат, т.е. инвариантна относительно ортогональных преобразований $\mathbb{R}^d$, оставляющих точку $\bar{0}$ на месте. Такие распределения называются сферически симметричными. Если случайный вектор имеет сферически симметричное распределение, то ортогональное преобразование с неподвижной точкой $\bar{0}$ переводит этот случайный вектор в другой случайный вектор, имеющий то же самое распределение.

Согласно определению случайный вектор $\bar{\zeta}$ с $\mathbf{M}\bar{\zeta} = \bar{0}$, $\text{Cov}(\bar{\zeta}) = I$ можно представить в виде $\bar{\zeta} = \zeta_1 \bar{e}_1 + \dots + \zeta_d \bar{e}_d$, где $\zeta_1, \dots, \zeta_d$ - независимые случайные величины, имеющие нормальное распределение с параметрами $(0, 1)$, а $\bar{e}_1, \dots, \bar{e}_d$ - единичные векторы координатных осей. Из сферической симметричности распределения $\bar{\zeta}$ следует, что для любого другого ортонормированного базиса $\bar{e}'_1, \dots, \bar{e}'_d$ в соответствующем разложении $\bar{\zeta} = \zeta'_1 \bar{e}'_1 + \dots + \zeta'_d \bar{e}'_d$ случайные величины $\zeta'_1, \dots, \zeta'_d$ тоже независимы и имеют то же самое нормальное распределение с параметрами $(0, 1)$.

Распределение любого случайного вектора $c\bar{\zeta}$ с математическим ожиданием $\bar{0}$ и матрицей ковариаций $c^2 I$ тоже является нормальным и сферически симметричным. Такое распределение имеет плотность

$$\varphi_{\bar{0}, c^2 I}(x_1, \dots, x_d) = \frac{1}{(2\pi)^{d/2} c^d} e^{-\frac{x_1^2 + \dots + x_d^2}{2c^2}}$$

и характеристическую функцию

$$\psi_{\bar{0}, c^2 I}(t_1, \dots, t_d) = e^{-\frac{(t_1^2 + \dots + t_d^2)c^2}{2}}$$

Определение.

Распределение случайного d -мерного вектора $\bar{\kappa}$ называется многомерным нормальным, если его можно представить в виде линейного преобразования случайного вектора $\bar{\zeta}_{\bar{0}, I}$ с нулевым математическим ожиданием и единичной матрицей ковариаций: $\bar{\kappa} = B\bar{\zeta}_{\bar{0}, I} + \bar{a}$, где $\bar{a} \in \mathbb{R}^d$ - неслучайный вектор и B - $(d \times d)$ -матрица.

Найдем $\mathbf{M}\bar{\kappa}$ и ковариационную матрицу $\text{Cov}(\bar{\kappa})$. Переход от вектора $\bar{\zeta}_{\bar{0}, I}$ к вектору $\bar{\kappa}$ осуществляется преобразованием $g(\bar{x}) = B\bar{x} + \bar{a}$. Значит,

$$\mathbf{M}\bar{\kappa} = \mathbf{M}g(\bar{\zeta}_{\bar{0}, I}) = \mathbf{M}(B\bar{\zeta}_{\bar{0}, I} + \bar{a}) = \bar{a}$$

В силу утверждения 2 из §18

$$\text{Cov}(\bar{\kappa}) = B \text{Cov}(\bar{\zeta}_{\bar{0}, I}) B^T = BB^T$$

В частности, отсюда следует, что $\det C(\bar{\kappa}) = (\det B)^2 \geq 0$.

Согласно утверждению 1 из §24 при невырожденной $d \times d$ -матрице B плотность распределения вектора $\bar{\kappa}$ имеет вид

$$\begin{aligned} \varphi_{\bar{\kappa}}(\bar{x}) &= \frac{\varphi_{\bar{\zeta}_{\bar{0}, I}}(B^{-1}(\bar{x} - \bar{a}))}{|\det B|} = \\ &= \frac{1}{(2\pi)^{d/2} |\det B|} \exp \left\{ -\frac{1}{2} (B^{-1}(\bar{x} - \bar{a}), B^{-1}(\bar{x} - \bar{a})) \right\} = \\ &= \frac{1}{\sqrt{(2\pi)^d \det C(\bar{\kappa})}} \exp \left\{ -\frac{1}{2} (\bar{x} - \bar{a}) C^{-1}(\bar{\kappa}) (\bar{x} - \bar{a})^T \right\} \end{aligned}$$

По теореме из линейной алгебры для любой неотрицательно определенной симметричной матрицы C существует такая матрица B , что $C = BB^\top$. Поэтому для любой симметричной неотрицательно определенной $d \times d$ -матрицы C и любого вектора $\bar{a} \in \mathbb{R}^d$ существует d -мерное нормальное распределение с математическим ожиданием $\bar{a}$ и ковариационной матрицей C .

Так как плотность $\varphi_{\bar{k}}(\bar{x})$ выражается через $\bar{a} = M\zeta$ и $\text{Cov}(\bar{\zeta})$, то многомерное нормальное распределение однозначно определяется вектором математических ожиданий и матрицей ковариаций. В отличие от плотности сферически симметричного нормального распределения, у которой поверхностями уровня являлись сферы с центром в начале координат, для плотности многомерного нормального распределения (48) поверхностями уровня являются эллипсоиды.

Пользуясь свойством 2) характеристических функций, найдем характеристическую функцию многомерного нормального распределения с плотностью (48): так как $\bar{k} = B\bar{\zeta}_{0,I}^\top + \bar{a}$, то

$$\begin{aligned}\psi_{\bar{k}}(\bar{t}) &= M e^{i(\bar{t}, \bar{k})} = e^{i(\bar{t}, \bar{a})} \psi_{\bar{\zeta}}(B^\top \bar{t}) = e^{i(\bar{t}, \bar{a})} e^{-|B^\top \bar{t}|^2/2} = \\ &= e^{i(\bar{t}, \bar{a})} e^{-(B^\top \bar{t}, B^\top \bar{t})/2} = e^{i(\bar{t}, \bar{a})} e^{-(BB^\top \bar{t}, \bar{t})/2} = e^{i(\bar{t}, \bar{a})} e^{-t \text{Cov}(\bar{k}) \bar{t}^\top/2}, \quad \bar{t} \in \mathbb{R}^d\end{aligned}$$

Равенство (49) можно использовать как

определение

многомерного нормального распределения с математическим ожиданием $\bar{a}$ и ковариационной матрицей $C(\bar{k})$. Нужно иметь в виду, что плотность (48) существует только для многомерных нормальных распределений с невырожденными матрицами ковариаций.

Далее для случайного вектора с математическим ожиданием a и матрицей ковариаций C мы будем использовать обозначение $\zeta_{a,C}$. 3) При любом $\bar{t} \in \mathbb{R}^d$ значения $f(-\bar{t}) = M e^{i(-\bar{t}, \gamma)} = M e^{-i(\bar{t}, \gamma)}$ и $f(\bar{t}) = M e^{i(\bar{t}, \gamma)}$ комплексно сопряжены. 4) Если случайные векторы $\bar{\gamma}^{(1)}, \dots, \bar{\gamma}^{(n)}$ независимы, то

$$f_{\bar{\gamma}^{(1)} + \dots + \bar{\gamma}^{(n)}}(\bar{t}) = M e^{i(\bar{t}, \bar{\gamma}^{(1)} + \dots + \bar{\gamma}^{(n)})} = M \prod_{k=1}^n e^{i(\bar{t}, \bar{\gamma}^{(k)})} = \prod_{k=1}^n f_{\bar{\gamma}^{(k)}}(\bar{t})$$

Замечание. Из свойств 3) и 4) следует, что если $f(\bar{t}) = M e^{i(\bar{t}, \bar{\gamma})}$, то $|f(\bar{t})|^2$ - характеристическая функция случайного вектора $\bar{\gamma}_1 - \bar{\gamma}_2$, где $\bar{\gamma}_1$ и $\bar{\gamma}_2$ независимы и имеют такое же распределение, как $\bar{\gamma}$. Распределение разности $\bar{\gamma}_1 - \bar{\gamma}_2$ симметрично:

$-(\bar{\gamma}_1 - \bar{\gamma}_2) = \bar{\gamma}_2 - \bar{\gamma}_1$, и характеристическая функция любого симметричного распределения принимает только действительные неотрицательные значения. 5) Теорема. Различным распределениям вероятностей соответствуют различные характеристические функции. 6)

Теорема

непрерывности. а) Если последовательность F_n распределений случайных векторов $\bar{\gamma}_n$ со значениями в $\mathbb{R}^d$ слабо сходится при $n \rightarrow \infty$ к распределению F при каждом $\bar{t}$. б) Обратно, если последовательность характеристических функций $f_n(\bar{t})$ при каждом $\bar{t}$ сходится к функции $f(\bar{t})$, непрерывной в нуле, то последовательность соответствующих распределений F_n слабо сходится к распределению F и $f(\bar{t})$ есть

характеристическая функция распределения F . Как и в скалярном случае, теоремы единственности и непрерывности в этом курсе приводятся без доказательств.

Следствие. Непрерывная функция, являющаяся поточечным пределом последовательности характеристических функций, сама является характеристической.

7) Лемма 1. Если случайный вектор $\bar{\gamma} = (\gamma_1, \dots, \gamma_d)$ имеет конечные смешанные моменты

$$m_{k_1, \dots, k_d} = \mathbf{M}\gamma_1^{k_1} \dots \gamma_d^{k_d}, \quad k_1, \dots, k_d \in \{0, 1, \dots\}, \quad \sum_{j=1}^d k_j \leq r$$

и характеристическую функцию $f(\bar{t}) = \mathbf{M}e^{i(\bar{t}, \bar{\gamma})}$, $\bar{t} = (t_1, \dots, t_d) \in \mathbb{R}^d$, то существуют частные производные характеристической функции вплоть до r -го порядка и

$$\frac{\partial^{k_1 + \dots + k_d}}{\partial t_1^{k_1} \dots \partial t_d^{k_d}} f(t_1, \dots, t_d) = i^{k_1 + \dots + k_d} \mathbf{M}e^{i(\bar{t}, \bar{\gamma})} \gamma_1^{k_1} \dots \gamma_d^{k_d}, \quad \sum_{j=1}^d k_j \leq r$$

$$m_{k_1, \dots, k_d} = i^{-k} \left. \frac{\partial^k f(\bar{t})}{\partial t_1^{k_1} \dots \partial t_d^{k_d}} \right|_{\bar{t}=\bar{0}}, \quad k = \sum_{j=1}^d k_j \leq r$$

и

$$f(\bar{t}) = 1 + \sum_{k=1}^r i^k \sum_{k_1 + \dots + k_d = k} m_{k_1, \dots, k_d} \frac{t_1^{k_1} \dots t_d^{k_d}}{k_1! \dots k_d!} + R_k(\bar{t})$$

где $R_k(\bar{t}) = o(|\bar{t}|^k)$ при $|\bar{t}| = |t_1| + \dots + |t_d| \rightarrow 0$. Доказательство леммы. Доказательство проводится индукцией. Формула для частной производной первого порядка, например, по t_1 следует из равенства

$$\frac{f(t_1 + h, t_2, \dots, t_d) - f(t_1, t_2, \dots, t_d)}{h} = \mathbf{M}e^{i(\bar{t}, \bar{\gamma})} \frac{e^{ih\gamma_1} - 1}{h}$$

Согласно лемме 2 из §29 дробь под знаком математического ожидания не превосходит $|h|$ по абсолютной величине и стремится к $i\gamma_1$ при $h \rightarrow 0$, поэтому под знаком математического ожидания можно перейти к пределу:

$$\frac{\partial}{\partial t_1} f(t_1, \dots, t_d) = \mathbf{M}e^{i(\bar{t}, \bar{\gamma})} \gamma_1$$

Получив такие формулы для всех частных производных первого порядка, аналогично можно доказывать формулы для производных второго порядка, и по индукции для частных производных всех таких порядков, при которых конечны соответствующие смешанные моменты. Лемма доказана.

Формулы для моментов получаются из формул леммы 1 подстановкой $\bar{t} = \bar{0}$. После того, как доказаны формулы для частных производных, формула (51) становится формулой Тейлора для функции d переменных.

Следствие. Пусть случайный вектор $\bar{\gamma} = (\gamma_1, \dots, \gamma_d)$ имеет конечные математическое ожидание $\bar{a} = (a_1, \dots, a_d) = (\mathbf{M}\gamma_1, \dots, \mathbf{M}\gamma_d)$ и матрицу ковариаций $C(\bar{\gamma})$. Тогда при $|\bar{t}| = |t_1| + \dots + |t_d| \rightarrow 0$

$$f(\bar{t}) = e^{i(\bar{t}, \bar{a})} \left(1 - \frac{1}{2} \bar{t} C(\bar{\gamma}) \bar{t}^T + o(|\bar{t}|^2) \right)$$

Доказательство. Представим вектор $\bar{\gamma}$ в виде $\bar{\gamma} = \bar{a} + (\bar{\gamma} - \bar{a})$. По свойству 2)

$$\mathbf{M}e^{i(\bar{t}, \bar{\gamma})} = e^{i(\bar{t}, \bar{a})} \mathbf{M}e^{i(\bar{t}, \bar{\gamma} - \bar{a})}$$

Далее, $\mathbf{M}(\bar{\gamma} - \bar{a}) = (m_1, \dots, m_d) = 0$ и

$$C(\bar{\gamma}) = \|c_{ij}\| = \|\text{Cov}(\gamma_i, \gamma_j)\| = \|\mathbf{M}(\gamma_i - a_i)(\gamma_j - a_j)\|$$

т.е. ковариации компонент вектора $\bar{\gamma}$ совпадают со смешанными моментами второго порядка вектора $\bar{\gamma} - \bar{a}$; поэтому по свойству 6)

$$\begin{aligned} \text{Me}^{i(\bar{t}, \bar{\gamma} - \bar{a})} &= 1 - \sum_{k_1 + \dots + k_d = 2} m_{k_1, \dots, k_d} \frac{t_1^{k_1} \dots t_d^{k_d}}{k_1! \dots k_d!} + o(|\bar{t}|^2) = \\ &= 1 - \frac{1}{2} \sum_{i,j=1}^d c_{ij} t_i t_j + o(|\bar{t}|^2) \end{aligned}$$

Тем самым следствие доказано. Теперь докажем центральную предельную теорему для сумм независимых одинаково распределенных случайных векторов.

12.1.4 law of Large Numbers

простейший закон больших чисел

(зубков 19)

просто сперва сформулируем и докажем ее.

Закон больших чисел

Пусть случайные величины $\xi_1, \xi_2, \dots$ попарно независимы или попарно некоррелированы и имеют одинаковые математические ожидания и дисперсии: $\mathbf{M}\xi_t = a, \mathbf{D}\xi_t = \sigma^2 < \infty$. Тогда для каждого $\varepsilon > 0$

$$\mathbf{P} \left\{ \left| \frac{1}{n} (\xi_1 + \dots + \xi_n) - a \right| < \varepsilon \right\} \rightarrow 1, \quad n \rightarrow \infty$$

Замечание.

Для случайных величин ξ_t , множество значений которых конечно, условие $\mathbf{D}\xi_t < \infty$ выполняется автоматически.

Закон больших чисел в указанной форме является следствием более общего утверждения.

Теорема.

Пусть случайные величины $\xi_1, \xi_2, \dots$ попарно независимы или попарно некоррелированы и $\sup_t \mathbf{D}\xi_t \leq \sigma^2 < \infty$. Тогда для любого $\varepsilon > 0$

$$\lim_{n \rightarrow \infty} \mathbf{P} \left\{ \left| \frac{1}{n} (\xi_1 + \dots + \xi_n) - \frac{1}{n} (\mathbf{M}\xi_1 + \dots + \mathbf{M}\xi_n) \right| < \varepsilon \right\} = 1$$

Доказательство.

Положим $\zeta_n = \xi_1 + \dots + \xi_n$. Тогда

$$\mathbf{M}\zeta_n = \mathbf{M}\xi_1 + \dots + \mathbf{M}\xi_n, \quad \mathbf{D}\zeta_n = \mathbf{D}\xi_1 + \dots + \mathbf{D}\xi_n \leq n\sigma^2$$

Следовательно,

$$\begin{aligned} \mathbf{M} \left(\frac{1}{n} \zeta_n \right) &= \frac{1}{n} (\mathbf{M}\xi_1 + \dots + \mathbf{M}\xi_n) \\ \mathbf{D} \left(\frac{1}{n} \zeta_n \right) &\leq \frac{1}{n^2} (\mathbf{D}\xi_1 + \dots + \mathbf{D}\xi_n) \leq \left(\frac{1}{n^2} \right) n\sigma^2 = \frac{\sigma^2}{n} \end{aligned}$$

Применим к случайной величине ζ_n/n неравенство Чебышева:

$$\mathbf{P} \left\{ \left| \frac{1}{n} \zeta_n - \mathbf{M} \left(\frac{1}{n} \zeta_n \right) \right| \geq \varepsilon \right\} \leq \frac{\mathbf{D} \left(\frac{1}{n} \zeta_n \right)}{\varepsilon^2} \leq \frac{\sigma^2}{n\varepsilon^2}$$

Правая часть этого неравенства стремится к 0 при $n \rightarrow \infty$. Следовательно, при $n \rightarrow \infty$

$$\mathbf{P} \left\{ \left| \frac{1}{n} \zeta_n - \mathbf{M} \left(\frac{1}{n} \zeta_n \right) \right| < \varepsilon \right\} = 1 - \mathbf{P} \left\{ \left| \frac{1}{n} \zeta_n - \mathbf{M} \left(\frac{1}{n} \zeta_n \right) \right| \geq \varepsilon \right\} \rightarrow 1$$

что завершает доказательство.

Следствие (

Теорема

Бернулли).

(теперь мы можем доказать то, что выглядит максимально очевидно. а когда оно могло бы нарушаться?)

Пусть проводятся испытания по схеме Бернулли с вероятностью успеха $p \in (0, 1)$ и μ_n - число успехов в первых n испытаниях. Тогда при любом $\varepsilon > 0$

$$\mathbf{P} \left\{ \left| \frac{1}{n} \mu_n - p \right| < \varepsilon \right\} \rightarrow 1, \quad n \rightarrow \infty$$

Доказательство.

Свяжем с k -м испытанием Бернулли индикатор $\mathbb{I}_k$, который равен 1, если испытание оканчивается успехом, и равен 0 в противном случае.

Тогда $\mu_n = \mathbb{I}_1 + \dots + \mathbb{I}_n$, случайные величины $\mathbb{I}_1, \mathbb{I}_2, \dots$ независимы в совокупности

$$\mathbf{P} \{ \mathbb{I}_k = 1 \} = p_k, \quad \mathbf{P} \{ \mathbb{I}_k = 0 \} = 1 - p_k, \quad k = 1, 2, \dots$$

и поэтому $\mathbf{M} \mathbb{I}_k = p_k, \mathbf{D} \mathbb{I}_k = \mathbf{M} \mathbb{I}_k^2 - (\mathbf{M} \mathbb{I}_k)^2 = p - p^2 = p(1 - p) \leq \frac{1}{4}$. Теперь для доказательства следствия остается применить закон больших чисел.

усиленный закон больших чисел

(зубков)

Пусть случайные величины $\xi_1, \xi_2, \dots$ независимы, $\mathbf{M} \xi_n \equiv 0, \mathbf{D} \xi_n = \sigma_n^2, S_n = \xi_1 + \dots + \xi_n$ и $\sum_{n \geq 1} \frac{\sigma_n^2}{n^2} < \infty$. тогда

$$\frac{1}{n} S_n \xrightarrow{\text{п.н.}} 0$$

Доказательство.

Согласно утверждению 1 для доказательства того, что последовательность $\frac{1}{n} S_n$ сходится с вероятностью 1, достаточно показать, что при любом $\varepsilon > 0$

$$\lim_{n \rightarrow \infty} \mathbf{P} \left\{ \omega : \sup_{k \geq n} \left| \frac{1}{k} S_k \right| > \varepsilon \right\} = 0$$

Введем события

$$A_m = \left\{ \omega : \max_{2^{m-1} \leq k < 2^m} \left| \frac{1}{k} S_k \right| > \varepsilon \right\}, \quad m = 1, 2, \dots$$

Тогда (54) равносильно тому, что

$$\mathbf{P} \left\{ \bigcup_{m \geq n} A_m \right\} \rightarrow 0, \quad n \rightarrow \infty$$

Вероятность объединения событий не больше суммы вероятностей этих событий:

$$\mathbf{P} \left\{ \bigcup_{m \geq n} A_m \right\} \leq \sum_{m \geq n} \mathbf{P} \{A_m\}$$

поэтому (55) будет доказано, если показать, что ряд $\sum_{m \geq 1} \mathbf{P} \{A_m\}$ сходится.

Оценим члены этого ряда с помощью неравенства Колмогорова:

$$\begin{aligned} \mathbf{P} \{A_m\} &= \mathbf{P} \left\{ \max_{2^{m-1} \leq k < 2^m} \left| \frac{1}{k} S_k \right| > \varepsilon \right\} \leq \mathbf{P} \left\{ \max_{2^{m-1} \leq k < 2^m} |S_k| > 2^{m-1} \varepsilon \right\} \leq \\ &\leq \mathbf{P} \left\{ \max_{1 \leq k < 2^m} |S_k| > 2^{m-1} \varepsilon \right\} \leq \frac{\mathbf{D} S_{2^m}}{(\varepsilon 2^{m-1})^2} = \frac{1}{\varepsilon^2 2^{2(m-1)}} \sum_{k=1}^{2^m} \sigma_k^2 \end{aligned}$$

Из полученных оценок и условия теоремы следует, что

$$\begin{aligned} \sum_{m \geq 1} \mathbf{P} \{A_m\} &\leq \frac{1}{\varepsilon^2} \sum_{m \geq 1} \sum_{k=1}^{2^m} \frac{\sigma_k^2}{2^{2(m-1)}} = \frac{1}{\varepsilon^2} \sum_{k \geq 1} \sigma_k^2 \sum_{m: 2^m \geq k} \frac{1}{2^{2(m-1)}} < \\ &< \frac{1}{\varepsilon^2} \sum_{k \geq 1} \frac{4}{1 - \frac{1}{4}} \frac{\sigma_k^2}{k^2} < \frac{16}{3\varepsilon^2} \sum_{k \geq 1} \frac{\sigma_k^2}{k^2} < \infty \end{aligned}$$

Теорема

доказана.

Следствие

Для числа ν_n успехов в схеме Бернулли с вероятностью успеха p справедливы как слабый закон больших чисел

$$\frac{\nu_n}{n} \xrightarrow{P} p, \quad n \rightarrow \infty$$

так и усиленный закон больших чисел

$$\frac{\nu_n}{n} \xrightarrow{\text{п. н.}} p, \quad n \rightarrow \infty$$

Проведенные рассуждения можно дополнить и доказать, что усиленный закон больших чисел для одинаково распределенных случайных величин верен при тех же условиях, при которых верен слабый закон больших чисел.

Усиленный закон больших чисел Колмогорова.

Пусть случайные величины $\xi_1, \xi_2, \dots$ независимы и одинаково распределены.

Для справедливости усиленного закона больших чисел $\frac{1}{n} S_n \xrightarrow{\text{п. н.}}$ а необходимо и достаточно, чтобы существовало конечное математическое ожидание $\mathbf{M} \xi_k = a$.

12.1.5 types of Convergence of Sequences of Random Variables

(горяйнов. потом укажу применения, пока хз!!! мб в дополнения потом перекину, хз)

В этом параграфе рассматриваются основные виды сходимости последовательностей случайных величин и устанавливается взаимосвязь между ними.

Определение сходимости по вероятности

Будем говорить, что последовательность случайных величин $\xi_1, \xi_2, \dots$ сходится по вероятности к случайной величине ξ , и писать $\xi_n \xrightarrow{P} \xi$, если для всякого $\varepsilon > 0$ выполняется предельное соотношение

$$\lim_{n \rightarrow \infty} P(|\xi_n - \xi| \geq \varepsilon) = 0$$

Именно этот вид сходимости имеет место в законе больших чисел Чебышёва.

Заметим, что из сходимости по вероятности последовательности случайных величин мы ничего не можем сказать о том, как ведет себя последовательность в отдельных экспериментах.

В этом отношении более содержательную информацию дает следующий вид сходимости.

определение сходимости почти наверное

(ну и название...)

Будем говорить, что последовательность случайных величин $\xi_1, \xi_2, \dots$ сходится к случайной величине ξ с вероятностью 1 или почти наверное, и писать $\xi_n \xrightarrow{п.н.} \xi$, если

$$P\left(\left\{\omega : \lim_{n \rightarrow \infty} \xi_n(\omega) = \xi(\omega)\right\}\right) = 1$$

Следующий вид сходимости требует от случайных величин существования средних порядка p .

определение

10.3.

Будем говорить, что последовательность случайных величин $\xi_1, \xi_2, \dots$ сходится к случайной величине ξ в среднем порядка $p, p > 0$ и писать $\xi_n \xrightarrow{L_p} \xi$, если $M|\xi_n - \xi|^p \rightarrow 0$ при $n \rightarrow \infty$. При $p = 2$ эта сходимость называется также в среднем квадратичном.

определение

10.4.

Будем говорить, что последовательность случайных величин $\xi_1, \xi_2, \dots$ сходится к случайной величине ξ по распределению, и писать $\xi_n \xrightarrow{d} \xi$, если $F_{\xi_n} \rightarrow F_\xi$

Теорема

Пусть $\xi_1, \xi_2, \dots$ — последовательность случайных величин. Тогда имеют место следующие утверждения:

- (i) Если $\xi_n \xrightarrow{п.н.} \xi$, то $\xi_n \xrightarrow{P} \xi$
- (ii) Если $\xi_n \xrightarrow{L_p} \xi$, то $\xi_n \xrightarrow{P} \xi$
- (iii) Если $\xi_n \xrightarrow{P} \xi$, то $\xi_n \xrightarrow{d} \xi$.

Доказательство.

- (i) Для $\varepsilon > 0$ введем в рассмотрение последовательность событий

$$A_n^\varepsilon = \{\omega : |\xi_n(\omega) - \xi(\omega)| \geq \varepsilon\}$$

$n = 1, 2, \dots$ Заметим, что условие $\xi_n \xrightarrow{\mathbf{P}} \xi$ эквивалентно тому, что для любого $\varepsilon > 0$ должно выполняться предельное соотношение $\mathbf{P}(A_n^\varepsilon) \rightarrow 0$ при $n \rightarrow \infty$. Также для каждого $\varepsilon > 0$ и $n = 1, 2, \dots$ рассмотрим

$$B_n^\varepsilon = \left\{ \omega : \sup_{k \geq n} |\xi_k(\omega) - \xi(\omega)| \geq \varepsilon \right\}$$

Очевидно, что $A_n^\varepsilon \subset B_n^\varepsilon$ и $\{B_n^\varepsilon\}_{n=1}^\infty$ является монотонно убывающей последовательностью.

Пусть $B^\varepsilon = \bigcap_{n=1}^\infty B_n^\varepsilon$ - предел этой последовательности.

В силу свойства непрерывности вероятностной меры $\mathbf{P}(B_n^\varepsilon) \rightarrow \mathbf{P}(B^\varepsilon)$ при $n \rightarrow \infty$.

С другой стороны, если $\omega \in B^\varepsilon$, то $\xi_n(\omega) \not\xrightarrow{\mathbf{P}} \xi(\omega)$. Это означает, что

$$B^\varepsilon \subset \{\omega : \xi_n(\omega) \not\xrightarrow{\mathbf{P}} \xi(\omega)\}$$

Однако, из условия $\xi_n \xrightarrow{\text{п.н.}} \xi$ следует

$$\mathbf{P}(\{\omega : \xi_n(\omega) \not\xrightarrow{\mathbf{P}} \xi(\omega)\}) = 0$$

Но тогда

$$\lim_{n \rightarrow \infty} \mathbf{P}(B_n^\varepsilon) = \mathbf{P}(B^\varepsilon) = 0$$

Из неравенств

$$0 \leq \mathbf{P}(A_n^\varepsilon) \leq \mathbf{P}(B_n^\varepsilon)$$

следует, что и $\mathbf{P}(A_n^\varepsilon) \rightarrow 0$ при $n \rightarrow \infty$, т.

е.

$$\xi_n \xrightarrow{\mathbf{P}} \xi.$$

(ii) Допустим теперь, что $\xi_n \xrightarrow{L^{\mathcal{R}}} \xi$, т.

е.

$\mathbf{M}|\xi_n - \xi|^p \rightarrow 0$ при $n \rightarrow \infty$. Фиксируем произвольно $\varepsilon > 0$. Используя неравенство (8.1) Маркова, получаем

$$\mathbf{P}(|\xi_n - \xi| \geq \varepsilon) = \mathbf{P}(|\xi_n - \xi|^p \geq \varepsilon^p) \leq \frac{\mathbf{M}|\xi_n - \xi|^p}{\varepsilon^p} \rightarrow 0$$

при $n \rightarrow \infty$, т.

е.

$$\xi_n \xrightarrow{\mathbf{P}} \xi. \quad (\text{iii}) \text{ Пусть } \xi_n \xrightarrow{\mathbf{P}} \xi, \text{ т.}$$

е.

$\mathbf{P}(A_n^\varepsilon) \rightarrow 0$ при $n \rightarrow \infty$ для всех $\varepsilon > 0$. Обозначим $F_n = F_{\xi_n}$ и $F = F_\xi$.

Доказательство $F_n \rightarrow F$ будет следовать из того, что для всех $\varepsilon > 0$ выполняются неравенства

$$F(x - \varepsilon) \leq \liminf_{n \rightarrow \infty} F_n(x) \leq \overline{\lim}_{n \rightarrow \infty} F_n(x) \leq F(x + \varepsilon), \quad x \in \mathbb{R}$$

Эти неравенства, в свою очередь, сразу же следуют из включений

$$\{\xi \leq x - \varepsilon\} \subset \{\xi_n \leq x\} \cup A_n^\varepsilon, \quad \{\xi_n \leq x\} \subset \{\xi \leq x + \varepsilon\} \cup A_n^\varepsilon$$

и того, что $\mathbf{P}(A_n^\varepsilon) \rightarrow 0$ при $n \rightarrow \infty$. В точках непрерывности функции F предельный переход в неравенствах (10.1) при $\varepsilon \rightarrow 0$ дает

$$F(x) \leq \liminf_{n \rightarrow \infty} F_n(x) \leq \overline{\lim}_{n \rightarrow \infty} F_n(x) \leq F(x)$$

что означает $F_n(x) \rightarrow F(x)$ при $n \rightarrow \infty$ и

Теорема доказана.

Пример-контрпример 2 (?)

Следующие два примера показывают, что из сходимости последовательности случайных величин почти наверное не следует сходимость в среднем, а из сходимости в среднем не следует сходимость почти наверное.

Пример 1. Пусть $\Omega = [0, 1]$, $\mathcal{A} = \mathcal{B}([0, 1])$, $\mathbf{P}$ - мера Лебега.

Рассмотрим последовательность

$$\xi_n(\omega) = n^{1/p} \cdot 1_{[0, 1/n]}(\omega)$$

$$n = 1, 2, \dots \text{ Тогда } \xi_n \xrightarrow{\text{п.н.}} \xi, \xi(\omega) \equiv 0, \text{ но } \xi_n \rightarrow \xi \text{ в } L_p, p > 0$$

Пример 2

Пусть $\Omega = \mathbb{T}$ - единичная окружность, $\mathcal{A} = \mathcal{B}(\mathbb{T})$, $\mathbf{P} = \mu/2\pi$ - нормированная мера Лебега.

Для $n = 1, 2, \dots$ рассмотрим

$$E_n = \left\{ \omega = e^{i\theta} : \sum_{k=1}^{n-1} \frac{1}{k} \leq \theta < \sum_{k=1}^n \frac{1}{k} \right\} \subset \mathbb{T}$$

Заметим, что $\mathbf{P}(E_n) = 1/(2\pi n) \rightarrow 0$ при $n \rightarrow \infty$.

Следовательно, для любого $p > 0$ последовательность $\xi_n = 1_{E_n}$ сходится в среднем порядка p и $\xi(\omega) \equiv 0$. С другой стороны, поскольку $\sum_{k=1}^{\infty} 1/k = \infty$, то каждое элементарное событие ω принадлежит бесконечному числу событий из последовательности $\{E_n\}$, т.е.

$\overline{\lim}_{n \rightarrow \infty} E_n = \Omega$. Это означает, что $\xi_n(\omega) \not\xrightarrow{\text{п.н.}} \xi(\omega)$ ни при каком $\omega \in \Omega$, т.е. последовательность ξ_n не сходится почти наверное.

Закон больших чисел Чебышёва касается последовательности независимых случайных величин с равномерно ограниченными дисперсиями.

Метод характеристических функций позволяет снять условие конечности вторых моментов для одинаково распределенных случайных величин.

Теорема Хинчина

(????)

Теорема

(Хинчина)

Пусть $\xi_1, \xi_2, \dots$ - последовательность независимых одинаково распределенных случайных величин с конечным $\xi = a$.

Тогда для $S_n = \xi_1 + \dots + \xi_n$ имеет место предельное соотношение

$$\frac{S_n}{n} \xrightarrow{\mathbf{P}} a \quad \text{при} \quad n \rightarrow \infty$$

Доказательство

ДоКАЗАТЕльСТВО.

Поскольку все $\xi_k, k = 1, 2, \dots$, одинаково распределены, то они имеют одну и ту же характеристическую функцию

$$h(t) = h_{\xi_k}(t) = \mathbf{M}e^{it\xi_k}$$

В силу того, что ξ_k имеет конечное математическое ожидание, характеристическая функция h дифференцируема и

$$h(t) = 1 + iat + o(t) \quad \text{при} \quad t \rightarrow 0$$

По условию $\xi_1, \dots, \xi_n$ независимы.

Поэтому

$$h_{S_n}(t) = (h(t))^n$$

и

$$h_{\frac{S_n}{n}}(t) = \left(h\left(\frac{t}{n}\right) \right)^n = \left(1 + ia\frac{t}{n} + o\left(\frac{t}{n}\right) \right)^n$$

Используя асимптотическую формулу $\ln(1+z) = z + o(|z|)$, при $z \rightarrow 0$, получаем

$$\lim_{n \rightarrow \infty} h_{\frac{S_n}{n}}(t) = \lim_{n \rightarrow \infty} e^{n \ln\left(1 + \frac{iat}{n} + o\left(\frac{t}{n}\right)\right)} = e^{iat}$$

Предельная функция e^{iat} непрерывна и по теореме сходимости она является характеристической функцией некоторой случайной величины η , причем $S_n/n \xrightarrow{d} \eta$. Заметим также, что $h_{\eta}(t) = e^{iat}$ является характеристической функцией вырожденного распределения, сосредоточенного в точке a , т.е.

$$F_{\eta}(x) = \mathbb{1}_{[a, \infty)}(x)$$

Остается показать, что $S_n/n \xrightarrow{\mathbf{P}} \eta$. Для $\varepsilon > 0$ имеем

$$\begin{aligned} \mathbf{P}\left(\left|\frac{S_n}{n} - a\right| \geq \varepsilon\right) &= \mathbf{P}\left(\frac{S_n}{n} \leq a - \varepsilon\right) + \mathbf{P}\left(\frac{S_n}{n} \geq a + \varepsilon\right) \\ &\leq F_n(a - \varepsilon) + \left(1 - F_n\left(a + \frac{\varepsilon}{2}\right)\right) \end{aligned}$$

где F_n - функция распределения случайной величины S_n/n . Поскольку точки $a - \varepsilon$ и $a + \varepsilon/2$ являются точками непрерывности функции F_{η} , то

$$F_n(a - \varepsilon) \rightarrow 0, \quad F_n(a + \varepsilon/2) \rightarrow 1$$

при $n \rightarrow \infty$. Отсюда следует, что

$$\lim_{n \rightarrow \infty} \mathbf{P}\left(\left|\frac{S_n}{n} - a\right| \geq \varepsilon\right) = 0$$

и $S_n/n \xrightarrow{\mathbf{P}} a$.

О применениях теоремы Хинчина

(хз вообще)
и что с этого?

12.1.6 Some Problems About Mathematical Statistics

(потом буду разбирать, пока: раз встретилось, то просто сюда вставлю и всё.)

Равенство для порядковых статистик (???)

(в матстатистике закину и забуду, ибо не до этого)

Пусть $X_1, \dots, X_n$ - независимые одинаково распределенные случайные величины с плотностью $f(x)$. Для каждого элементарного события $\omega \in \Omega$ вектор $(X_1(\omega), \dots, X_n(\omega))$ преобразуем в упорядоченный $(X_{(1)}(\omega), \dots, X_{(n)}(\omega))$, где $X_{(1)}(\omega) \leq \dots \leq X_{(n)}(\omega)$. Упорядоченный вектор $(X_{(1)}, \dots, X_{(n)})$ в математической статистике называют вариационным рядом, а случайные величины $X_{(k)}, k = 1, \dots, n$ - порядковыми статистиками. Плотность совместного распределения порядковых статистик определяется равенством.

$$f_{X_{(1)}, \dots, X_{(n)}}(x_1, \dots, x_n) = n! f(x_1) \cdot \dots \cdot f(x_n) \mathbb{1}_{\{x_1 \leq \dots \leq x_n\}}(x_1, \dots, x_n)$$

Действительно, множество, на котором совпали X_i, X_j $i \neq j$ имеет меру ноль. Для остальных каждому корректному вариационному ряду соответствует $n!$ неупорядоченных векторов. Из независимости $X_1, \dots, X_n$

$$f_{X_1, X_2, \dots, X_n}(x_1, x_2, \dots, x_n) = \prod_{k=1}^n f(x_k)$$

Поэтому $f_{X_{(1)}, \dots, X_{(n)}}(x_1, \dots, x_n) = n! f(x_1) \cdot \dots \cdot f(x_n) \mathbb{1}_{\{x_1 \leq \dots \leq x_n\}}(x_1, \dots, x_n)$.
(также тут укажу приложения)

12.2 Fundamentals of Mathematical Statistics

Разберем основы математической статистики, укажем примеры и приложения.

(пока что еще не дошел, но мб этим буду очень много заниматься.)

(структура по Жуковскому мфти, потом буду развивать нормально)

12.2.1 Confidence Interval

8.2.

12.2.2 Asymptotic Confidence Intervals

12.2.3 Linear Regression

12.2.4 Statistical Hypothesis Testing

12.2.5 Equally Most Powerful Criteria

12.2.6 Testing Simple Hypotheses

12.2.7 Testing Complex Hypotheses

12.2.8 Verification of Linear Hypotheses in the Gaussian Regression Model

12.2.9 Consent Criteria

12.2.10 Kolmogorov's Goodness-of-fit Criterion.

12.2.11 χ^2 Test.

(жуковский)

Пусть $X_1, \dots, X_n$ - выборка из схемы Бернулли с $m \geq 2$ исходами, т.е. $a_1, \dots, a_m$ - исходы и $P(X_i = a_j) = p_j$ для $j = 1, \dots, m$. Построим критерий для проверки гипотезы

$$H_0 : p_j = p_j^0, \quad j \in \{1, \dots, m\}$$

определение

12.3. Статистикой хи-квадрат называется

$$\hat{\chi} = \sum_{j=1}^m \frac{(\mu_j - np_j^0)^2}{np_j^0}$$

где

$$\mu_j = \sum_{i=1}^n I\{X_i = a_j\}$$

Теорема

В предположении верности гипотезы H_0 имеет место сходимость по распределению

$$\hat{\chi} \xrightarrow[n \rightarrow \infty]{d} \chi_{m-1}^2$$

А Критерий хи-квадрат Пирсона записывается следующим образом: $\{\hat{\chi} > u_{1-\alpha}\}$, где $u_{1-\alpha} = (1 - \alpha)$ -Квантиль распределения χ_{m-1}^2 . Из теоремы 12.2 следует, что критерий хи-квадрат является состоятельным. Напоследок заметим, что критерий хи-квадрат принято применять, если $n \geq 50$ и $\mu_j \geq 5$ для всех $j \in \{1, \dots, m\}$.

12.2.12 Other Mathematical Statistics (???)

(там много всего есть, так что тут укажу это, подробнее - в записи по мат статистике)

Коэффициенты корреляции

Коэффициент корреляции Пирсона

Коэффициент корреляции Спирмена

Коэффициент корреляции Кэндалла

Байесовские оценки

12.3 Typical Distributions

Приведем часто встречающиеся распределения.

(в порядке пользы отсортирую потом. укажу, зачем они нужны. графики добавлю потом! удаляю разделы про распределения, которые не так нужны.)

12.3.1 Normal Distribution

Суть нормального распределения

Нормальное (или гауссовское) с плотностью

$$f_{\xi}(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-a)^2/2\sigma^2}, \quad a \in \mathbb{R}, \quad \sigma > 0$$

Случайная величина ζ имеет стандартное нормальное распределение с параметрами $(0,1)$, если ее плотность распределения равна $\varphi(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$, $-\infty < x < \infty$; функция стандартного нормального распределения есть

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-u^2/2} du$$

а его характеристическая функция

$$\psi(t) = \mathbf{M}e^{it\zeta} = e^{-t^2/2}, \quad -\infty < t < \infty$$

Свойства (!!!)

(очень часто могут быть задачи про это. мб в разделе выше это напишу, посмотрим.)

Случаи применения

Вывод

12.3.2 Homogeneous Distribution

Конструкция и свойства равномерного распределения

Равномерное на $[a, b]$

$$f_{\xi}(x) = \frac{1}{b-a} \mathbb{1}_{[a,b]}(x)$$

Равномерное распределение на конечном множестве $\{1, \dots, N\}$ определяется равенствами

$$\mathbf{P}\{\xi = m\} = \frac{1}{N}, \quad m = 1, 2, \dots, N$$

Функция распределения $F_{\xi}(x) = \mathbf{P}\{\xi \leq x\}$ такой случайной величины равна 0 при $x < a$, равна 1 при $x > b$ и равна $\frac{x-a}{b-a}$ при $a \leq x \leq b$. Плотность распределения этой случайной величины равна $\frac{1}{b-a}$ при $a \leq x \leq b$ и равна 0 вне отрезка $[a, b]$.

Равномерное распределение не является устойчивым относительно суммирования.

Найдём функцию и плотность распределения суммы двух независимых случайных величин с одинаковым равномерным на отрезке $[0, 1]$ распределением, но не по формуле свёртки, а используя геометрическую вероятность.

Пусть $\xi_1, \xi_2 \in U_{0,1}$ - независимые случайные величины. Пару (ξ_1, ξ_2) можно считать координатой точки, брошенной наудачу в единичный квадрат. Тогда $F_{\xi_1 + \xi_2}(x) = P(\xi_1 + \xi_2 < x)$ равна площади области внутри квадрата под прямой $x_2 = x - x_1$. Эта область - заштрихованный треугольник при $0 < x \leq 1$, и пятиугольник при $1 < x \leq 2$. Окончательно получаем:

$$F_{\xi_1 + \xi_2}(x) = \begin{cases} 0, & x \leq 0 \\ x^2/2, & 0 < x \leq 1 \\ 2x - x^2/2 - 1, & 1 < x \leq 2 \\ 1, & x > 2. \end{cases}$$

Плотность распределения суммы равна

$$f_{\xi_1 + \xi_2}(x) = \begin{cases} 0, & x \notin (0, 2) \\ x, & 0 < x \leq 1 \\ 2 - x, & 1 < x \leq 2 \end{cases}$$

Случаи применения

12.3.3 Demonstrative

Конструкция

Показательное

$$f_{\xi}(x) = \lambda e^{-\lambda x} \mathbb{1}_{[0, \infty)}(x), \quad \lambda > 0$$

Случаи применения

Вывод

12.3.4 Cauchy Distribution

Конструкция

$$f_{\xi}(x) = \frac{1}{\pi} \frac{1}{x^2 + 1}$$

Вырожденное распределение — это распределение случайной величины, которая с вероятностью 1 принимает какое-то одно значение (т. е. как функция на пространстве элементарных событий Ω она равна константе на множестве меры 1)

Случаи применения**Вывод****12.3.5 Binomial (???)****Конструкция**

Пусть $\eta_1, \eta_2, \dots$ - индикаторы успехов в последовательности независимых испытаний Бернулли с вероятностью успеха p :

$$\mathbf{P}\{\eta_k = 1\} = p, \quad \mathbf{P}\{\eta_k = 0\} = 1 - p, k = 1, 2, \dots$$

Биномиальный закон с параметрами (n, p) соответствует распределению числа $\mu_n = \eta_1 + \dots + \eta_n$ успехов в n независимых испытаниях по схеме Бернулли, если в каждом испытании вероятность успеха равна p :

$$\mathbf{P}\{\mu_n = m\} = C_n^m p^m (1 - p)^{n-m}, \quad m = 0, 1, \dots, n$$

Случаи применения**12.3.6 Geometric Distribution****Конструкция**

Геометрическое распределение с параметром p определяется равенствами

$$\mathbf{P}\{\nu = m\} = p(1 - p)^m, \quad m = 0, 1, 2, \dots$$

Одна из интерпретаций случайной величины с геометрическим распределением Число неудач (нулей) до первого успеха (единицы) в последовательности испытаний Бернулли: $\nu = \min\{n \geq 1 : \eta_n = 1\} - 1$, где $\eta_1, \eta_2, \dots$ - те же, что в предыдущем примере.

Случаи применения**Вывод****12.3.7 Hypergeometric Distribution**

Гипергеометрическое распределение имеет вид:

$$\mathbf{P}\{\xi = m\} = \frac{C_M^m C_{N-M}^{n-m}}{C_N^n}, \quad \max\{0, n - N + M\} \leq m \leq \min\{n, M\}$$

(???)

$$\mathbf{P}\{\xi = m\} = \left(\frac{\lambda^m}{m!}\right) e^{-\lambda}, \quad m = 0, 1, 2, \dots$$

Очевидно,

$$\sum_{m \geq 0} \mathbf{P}\{\xi = m\} = \sum_{m \geq 0} \frac{\lambda^m}{m!} e^{-\lambda} = e^{\lambda} e^{-\lambda} = 1$$

Случаи применения

Гипергеометрическое распределение с параметрами n, M и N является распределением числа ξ белых шаров среди n шаров, извлеченных по схеме равновероятного выбора без возвращения урны, содержавшей вначале M белых и $N - M$ черных Шаров:

Вывод**12.3.8 Gamma Distribution****Суть**

Пусть распределение случайной величины X задаётся плотностью вероятности, имеющей вид $f_X(x) = \begin{cases} x^{k-1} \frac{e^{-x/\theta}}{\theta^k \Gamma(k)}, & x \geq 0 \\ 0, & x < 0 \end{cases}$, где $\Gamma(k)$ - гамма-функция Эйлера. Тогда говорят, что случайная величина X имеет гамма-распределение с положительными параметрами θ и k . Пишут $X \sim \Gamma(k, \theta)$.

Замечание. Иногда используют другую параметризацию семейства гаммараспределений. Или вводят третий параметр - сдвиг.

О случаях применения (!?!?!?)**Свойства**

(вики)

Математическое ожидание и дисперсия случайной величины X , имеющей гаммараспределение, имеют вид

$$\begin{aligned}\mathbb{E}[X] &= k\theta, \\ \mathbb{D}[X] &= k\theta^2.\end{aligned}$$

- Если $X_1, \dots, X_n$ - независимые случайные величины, такие что $X_i \sim \Gamma(k_i, \theta)$, $i = 1, \dots, n$, то

$$Y = \sum_{i=1}^n X_i \sim \Gamma\left(\sum_{i=1}^n k_i, \theta\right).$$

- Если $X \sim \Gamma(k, \theta)$, и $a > 0$ - произвольная константа, то $aX \sim \Gamma(k, a\theta)$ - Гамма-распределение бесконечно делимо.

12.3.9 Other Distributions: Overview

(о тех, про которых что-то слышал, но пока хз)

О распределении Ландау

(понадобится - пропишу, пока мало касался его)

О распределении Бозе-Эйнштейна

(подготовлю его свойства!!!)

$$\bar{n}_i = \frac{1}{e^{(\epsilon_i - \mu)/kT} - 1}$$

М ξ и D ξ для статистики Бозе-Эйнштейна (?!!!!)

(пока хз, но посмотрю, это важно. потом добавлю в пар-ф выше!)

В N ячеек случайно в соответствии со статистикой Бозе-Эйнштейна (частицы неразличимы и размещение без ограничений) размещаются n частиц. Пусть ξ - число пустых ячеек. Найти $M\xi$ и $D\xi$.

$$E\xi = \sum_i E\mathbb{1}_{A_i} = \sum_i P(A_i) = N \frac{|A|}{|\Omega|} = N \left(1 - \frac{1}{N}\right)^n$$

$$E\xi^2 = E\left(\sum_i \mathbb{1}_{A_i} + \sum_{i \neq j} \mathbb{1}_{A_i A_j}\right) = \sum_{i=k} E\mathbb{1}_{A_i}^2 + \sum_{i \neq k} E\mathbb{1}_{A_i} \mathbb{1}_{A_k} = (N^2 - N) \left(1 - \frac{2}{N}\right)^n$$

$$D\xi = E\xi^2 - (E\xi)^2 = (N^2 - N) \left(1 - \frac{2}{N}\right)^n - N^2 \left(1 - \frac{1}{N}\right)^{2n}$$

(хочется цифру в конце получить!!!)

О распределении Ферми-Дирака

(подготовлю его свойства!!!)

Мξ и Dξ для статистики Ферми-Дирака (?!!!!)

(пока хз, даже по аналогии хз, как это сделать)

12.4 Statistical Methods of Hypothesis Testing

Посмотрим, как можно статистически проверить, верна гипотеза или нет?

(пропишу сюда структуру и ответы как только разберусь с выгрузками из курсов мфти)

12.4.1 About Hypothesis Testing in Life (???)

как это делать?

в жизни где применять?

интересно.

тут как раз примеры

(раздел мелкий ибо по факту так никто не делает, а значит я в последнюю очередь буду про это думать, ибо тут отдельное скорее всего развитие нужно.)

напоминание про адекватность

Хоть методы проверки гипотез и будут хорошие, не нужно их применять к любой мелочи. Они нужны для на самом деле важных вопросов. К сожалению находятся отбитые от жизни люди, которые проверяют всякие мелочи, которые ни на что не влияют. Чтобы этого избежать - следует развиваться в психологии и духовности.

12.4.2 hypothesis Testing Theory

13 Applications in Physics (!?!?!?)

(отдельный большой раздел про это)

14 Analytics and Data Analysis with Probability Theory

(main methods that are used by companies)

15 Other Applications in Programming (???)

(тоже важное указание, потому что только для этого многие её и изучают.)

15.1 Typical Applications in Programming

(пока не до этого, чуть что - сюда писать буду)

15.1.1 Visualization of Probabilistic Model Results

Построение распределений на Питоне (??)

(идем по [?] - там отличное введение в приложения.)

как и обещанно в введении, будет много кода и примеров, позже раскидаю многое по другим главам.

выглядит так:

```
from IPython.core.pylabtools import figsize
import numpy as np
from matplotlib import pyplot as plt
figsize(11, 9)

import scipy.stats as stats

dist = stats.beta
n_trials = [0, 1, 2, 3, 4, 5, 8, 15, 50, 500]
data = stats.bernoulli.rvs(0.5, size=n_trials[-1])
x = np.linspace(0, 1, 100)

# For the already prepared, I'm using Binomial's conj. prior.
for k, N in enumerate(n_trials):
    sx = plt.subplot(len(n_trials) / 2, 2, k + 1)
    plt.xlabel("$p$, probability of heads") \
    if k in [0, len(n_trials) - 1] else None
    plt.setp(sx.get_yticklabels(), visible=False)
    heads = data[:N].sum()
    y = dist.pdf(x, 1 + heads, 1 + N - heads)
    plt.plot(x, y, label="observe %d tosses, \n %d heads" % (N, heads))
    plt.fill_between(x, 0, y, color="#348ABD", alpha=0.4)
    plt.vlines(0.5, 0, 4, color="k", linestyle="--", lw=1)

leg = plt.legend()
leg.get_frame().set_alpha(0.4)
plt.autoscale(tight=True)

plt.suptitle("Bayesian updating of posterior probabilities",
y=1.02,
fontsize=14)

plt.tight_layout()
```

и это уже круто. вот - прикладной теорвер!!!

15.2 Machine Learning

часто этим мы и занимаемся, ведь МЛ это круто и полезно.

15.3 Bayesian Machine Learning

обзор

есть курс на ютубе от ВШЭ про это, так что можно посмотреть.

статья.

вообще так или иначе - байесовские темы достойны отдельной подглавы.

степень нашей уверенности в событии A обозначим цифрой $P(A)$.

а далее мы получаем какую-то дополнительную информацию, и с ней наше видение вероятности уже становится другим. обозначим новое видение $P(A|X)$, где X - новое событие. оно называется *условная вероятность (posterior probability)*.

(какие-то рассуждения про величину количества испытаний)

вообще, следует заметить, что тут стиль более программирования и мало физики, мало сути вероятности. для сути вероятности нужно читать другую главу в разделе про квантмех. тут какой-то поверхностный такой метод.

но в этом и его польза.

16 Special constructions

(по идее хватит такой структуры на них, все равно этим почти не занимался.)

16.0.1 Frobenius-perron Theorem

Материал из Википедии — свободной энциклопедии [Перейти к навигации](#)[Перейти к поиску](#) Теорема Фробениуса — Перрона — теорема о наибольшем собственном значении вещественной квадратной матрицы с положительными компонентами. Эта теорема имеет многочисленные приложения в теории вероятностей (эргодичность цепей Маркова); в теории динамических систем; в экономике; в демографии; в социальных сетях; в поисковых системах.

Доказана Оскаром Перроном (1907) и независимо Георгом Фробениусом (1912). Идея использования этой теоремы для определения порядка игроков в турнирах принадлежит Эдмунду Ландау.

Формулировка

Пусть A - квадратная матрица, со строго положительными вещественными элементами, тогда справедливы утверждения: - наибольшее по модулю собственное значение r является вещественным и строго положительным; - это собственное значение является простым корнем характеристического многочлена; - соответствующий r собственной вектор имеет (точнее говоря, может быть выбран таким образом, чтобы иметь) строго положительные координаты, все остальные собственные векторы таким свойством не обладают; - собственное значение r удовлетворяет неравенствам $\min_i \sum_j a_{ij} \leq r \leq \max_i \sum_j a_{ij}$.

Доказательство (??)

16.1 Random Matrices

(почитаю избры, потом впишу)

и да, так как тема сложная, будем излагать просто, без четких определений как в прошлых разделах.

и тут важный взгляд такой: у нас была случайная величина, мы там интегрировали, что-то с ней делали, так же и с матрицами хочется делать, поэтому

Definition 16.1. дифференциальная вероятность от случайной матрицы

Definition 16.2. *дифференциал случайной матрицы*

(пока просто понятно что как-то определяется, потом пойму)

но в сути мы берем тип матриц, закидываем в элементы случайные величины и смотрим характеристики.

Definition 16.3. *гауссовый унитарный ансамбль - GUE - набор эрмитовых матрицы, с такой-то функцией распределения*

функция экспонента на что-то на след квадрата этой матрицы. ну типа напоминает функцию гауссова распределения

Definition 16.4. *среднее от функции от случайной матрицы*

ок, на этом этапе мы поняли, как ставится задача, ввели определения.

нужно еще характеризовать распределения.

(вообще как можно скорее нужно теорию вообще распределений клеить, но это позже обязательно сделаю. здесь как раз вопрос о распределениях)

Definition 16.5. *плотность состояний случайной матрицы*

в общем, на этом этапе мы должны просто понять постановку задачи и все такое и короче говоря мы задачи не умеем на это решать

16.2 diagramming Techniques

здесь типа самое главное эдванст методов

16.3 Quantum Probability

(тут от Амосова прописываю введение в его случпроцы. потом в 1 части также добавлю теорию)

16.4 Special Models and Paradoxes

(пока тут, особо углубляться нет в планах)

16.4.1 Special Patterns of Ordinary Life

Парадокс дней рождений

(должно хватить этого раздела)

Модель

ИнтуИтИвное восприяИе

В группе из 23 человек вероятность совпадения дней рождения у двух человек столь высока, потому что рассматривается вероятность совпадения дней рождения у любых двух человек в группе. Эта вероятность определяется количеством пар людей, которые можно составить из 23 человек. Так как порядок людей в парах не имеет значения, общее число таких пар равно числу сочетаний из 23 по 2, то есть $(23 \times 22)/2 = 253$ пары.

В формулировке парадокса речь идёт именно о совпадении дней рождения у каких-либо двух членов группы. Одно из распространённых заблуждений состоит в том, что этот случай путают с другим случаем, на первый взгляд похожим, когда из группы выбирается

один человек и оценивается вероятность того, что день рождения каких-либо других членов группы совпадёт с днём рождения выбранного человека. В последнем случае вероятность совпадения значительно ниже.

Расчёт вероятности

Требуется определить вероятность того, что в группе из n человек как минимум у двух из них дни рождения совпадут. Пусть дни рождения распределены равномерно, то есть примем, что: - в году 365 дней (нет високосных лет); - в группе нет людей, заведомо родившихся в один день (например, близнецов); - рождаемость не зависит от дня недели, времени года и других факторов. В действительности это не совсем так - в частности, в некоторых странах из-за особенностей работы больниц больше детей рождается в определённые дни недели. Однако неравномерность распределения может лишь увеличить вероятность совпадения дней рождения, но не уменьшить: если бы все люди рождались только в 3 дня из 365, то вероятность совпадения дней рождения была бы очень высокой.

Возьмём наугад одного человека из группы и запомним его день рождения. Затем возьмём наугад второго человека, при этом вероятность того, что у него день рождения не совпадёт с днем рождения первого человека, равна $1 - \frac{1}{365}$. Затем возьмём третьего человека; при этом вероятность того, что его день рождения не совпадёт с днем рождения одного из первых двух, равна $1 - \frac{2}{365}$. Рассуждая по аналогии, мы дойдём до последнего человека, для которого вероятность несовпадения его дня рождения со всеми предыдущими будет равна $1 - \frac{n-1}{365}$. Перемножая все эти вероятности, получаем вероятность того, что все дни рождения в группе будут различными:

$$\bar{p}(n) = \left(1 - \frac{1}{365}\right) \cdot \left(1 - \frac{2}{365}\right) \cdot \dots \cdot \left(1 - \frac{n-1}{365}\right) = \frac{365 \cdot 364 \cdot \dots \cdot (365 - n + 1)}{365^n} = \frac{365!}{365^n (365 - n)!}$$

Тогда вероятность того, что хотя бы у двух человек из n дни рождения совпадут, равна

$$p(n) = 1 - \bar{p}(n).$$

Значение этой функции превосходит 1/2 при $n = 23$, при этом вероятность совпадения равна примерно 50,73 %, а $p(22) \approx 47,57\%$. Список значений n и соответствующих им вероятностей приведён в следующей таблице.

(напишу нормально, на вики куча воды)

Выводы (??)

О применениях (?)

(кст, если группа более 23 человек, то можно найти с одинаковым днем рождения. прикольно ведь)

Вероятность счастливого билета

(пока тут, напишу потом это решение.)

16.5 Other Theories

16.5.1 Decision Making Using Mathematics

(слышал, что такое есть, но это скорее бред, так что пока не трогаю.)

16.6 Applications to Financial Mathematics

связка с записями про финансы
посмотрю, как конкретно она там используется?

16.7 Game Theory

16.7.1 Lotteries: Simple Cases

какая вероятность, что 3 раза подряд 3/6 цифр совпадают?

16.8 Models of Situations in Life (!)

16.8.1 Random Meeting

Два человека живут в одном месте.

и скажем, они выходят из своих домов случайным образом и в течение 5 мин могут видеть друг друга. какая вероятность, что они встретятся?

17 Part I Methods by Alon, Spencer

17.1 1 The Basic Method

What you need is that your brain is open.

- Paul Erdős

17.1.1 1.1 The Probabilistic Method

The probabilistic method is a powerful tool for tackling many problems in discrete mathematics. Roughly speaking, the method works as follows: Trying to prove that a structure with certain desired properties exists, one defines an appropriate probability space of structures and then shows that the desired properties hold in this structures and then shows that the desired properties hold in this space with positive probability. The method is best illustrated by examples. Here is a simple one. The Ramsey number $R(k, \ell)$ is the smallest integer n such that in any two-coloring of the edges of a complete graph on n vertices K_n by red and blue, either there is a red K_k (i.e., a complete subgraph on k vertices all of whose edges are colored red) or there is a blue K_ℓ . Ramsey (1929) showed that $R(k, \ell)$ is finite for any two integers k and ℓ . Let us obtain a lower bound for the diagonal Ramsey numbers $R(k, k)$.

Proposition 1.1.1 If $\binom{n}{k} \cdot 2^{1-\binom{k}{2}} < 1$ then $R(k, k) > n$. Thus $R(k, k) > \lfloor 2^{k/2} \rfloor$ for all $k \geq 3$.

Proof. Consider a random two-coloring of the edges of K_n obtained by coloring each edge independently either red or blue, where each color is equally likely. For any fixed set R of k vertices, let A_R be the event that the induced subgraph of K_n on R is monochromatic (i.e., that either all its edges are red or they are all blue). Clearly, $\Pr[A_R] = 2^{1-\binom{k}{2}}$. Since there are $\binom{n}{k}$ possible choices for R , the probability that at least one of the events A_R occurs is at most $\binom{n}{k} 2^{1-\binom{k}{2}} < 1$. Thus, with positive probability, no event A_R occurs and there is a two-coloring of K_n without a monochromatic K_k ; that is, $R(k, k) > n$. Note that if $k \geq 3$ and we take $n = \lfloor 2^{k/2} \rfloor$ then

$$\binom{n}{k} 2^{1-\binom{k}{2}} < \frac{2^{1+\frac{k}{2}}}{k!} \cdot \frac{n^k}{2^{k^2/2}} < 1$$

and hence $R(k, k) > \lfloor 2^{k/2} \rfloor$ for all $k \geq 3$. $\square$

This simple example demonstrates the essence of the probabilistic method. To prove the existence of a good coloring we do not present one explicitly, but rather show, in a nonconstructive way, that it exists. This example appeared in a paper of P. Erdős from 1947. Although Szele had applied the probabilistic method to another combinatorial problem, mentioned in Chapter 2, already in 1943, Erdős was certainly the first one who understood the full power of this method and applied it successfully over the years to numerous problems. One can, of course, claim that the probability is not essential in the proof given above. An equally simple proof can be described by counting; we just check that the total number of two-colorings of K_n is larger than the number of those containing a monochromatic K_k .

Moreover, since the vast majority of the probability spaces considered in the study of combinatorial problems are finite spaces, this claim applies to most of the applications of the probabilistic method in discrete mathematics. Theoretically, this is, indeed, the case. However, in practice, the probability is essential. It would be hopeless to replace the applications of many of the tools appearing in this book, including, for example, the second moment method, the Lovász Local Lemma and the concentration via martingales by counting arguments, even when these are applied to finite probability spaces.

The probabilistic method has an interesting algorithmic aspect. Consider, for example, the proof of Proposition 1.1.1 that shows that there is an edge two-coloring of K_n without a monochromatic $K_{2 \log_2 n}$. Can we actually find such a coloring? This question, as asked, may sound ridiculous; the total number of possible colorings is finite, so we can try them all until we find the desired one. However, such a procedure may require $2^{\binom{n}{2}}$ steps; an amount of time that is exponential in the size $\left[= \binom{n}{2} \right]$ of the problem. Algorithms whose running time is more than polynomial in the size of the problem are usually considered impractical. The class of problems that can be solved in polynomial time, usually denoted by **P** [see, e.g., Aho, Hopcroft and Ullman (1974)], is, in a sense, the class of all solvable problems. In this sense, the exhaustive search approach suggested above for finding a good coloring of K_n is not acceptable, and this is the reason for our remark that the proof of Proposition 1.1.1 is nonconstructive; it does not supply a constructive, efficient and deterministic way of producing a coloring with the desired properties. However, a closer look at the proof shows that, in fact, it can be used to produce, effectively, a coloring that is very likely to be good. This is because for large k , if $n = \lfloor 2^{k/2} \rfloor$ then

$$\binom{n}{k} \cdot 2^{1-\binom{k}{2}} < \frac{2^{1+\frac{k}{2}}}{k!} \left(\frac{n}{2^{k/2}} \right)^k \leq \frac{2^{1+\frac{k}{2}}}{k!} \ll 1.$$

Hence, a random coloring of K_n is very likely not to contain a monochromatic $K_{2 \log n}$. This means that if, for some reason, we must present a two-coloring of the edges of K_{1024} without a monochromatic K_{20} we can simply produce a random two-coloring by flipping a fair coin $\binom{1024}{2}$ times. We can then deliver the resulting coloring safely; the probability that it contains a monochromatic K_{20} is less than $2^{11}/20$!, probably much smaller than our chances of making a mistake in any rigorous proof that a certain coloring is good! Therefore, in some cases the probabilistic, nonconstructive method does supply effective probabilistic algorithms. Moreover, these algorithms can sometimes be converted into deterministic ones. This topic is discussed in some detail in Chapter 16.

The probabilistic method is a powerful tool in Combinatorics and in Graph Theory. It is also extremely useful in Number Theory and in Combinatorial Geometry. More recently, it has been applied in the development of efficient algorithmic techniques and in the study of various computational problems. In the rest of this chapter we present several simple examples that demonstrate some of the broad spectrum of topics in which this method is helpful. More complicated examples, involving various more delicate probabilistic arguments, appear in the

rest of the book.

17.1.2 1.2 Graph Theory

A tournament on a set V of n players is an orientation $T = (V, E)$ of the edges of the complete graph on the set of vertices V . Thus for every two distinct elements x and y of V either (x, y) or (y, x) is in E , but not both. The name "tournament" is natural, since one can think of the set V as a set of players in which each pair participates in a single match, where (x, y) is in the tournament iff x beats y . We say that T has the property S_k if for every set of k players there is one who beats them all. For example, a directed triangle $T_3 = (V, E)$, where $V = \{1, 2, 3\}$ and $E = \{(1, 2), (2, 3), (3, 1)\}$, has S_1 . Is it true that for every finite k there is a tournament T (on more than k vertices) with the property S_k ? As shown by Erdős (1963b), this problem, raised by Schütte, can be solved almost trivially by applying probabilistic arguments. Moreover, these arguments even supply a rather sharp estimate for the minimum possible number of vertices in such a tournament. The basic (and natural) idea is that if n is sufficiently large as a function of k , then a random tournament on the set $V = \{1, \dots, n\}$ of n players is very likely to have property S_k . By a random tournament we mean here a tournament T on V obtained by choosing, for each $1 \leq i < j \leq n$, independently, either the edge (i, j) or the edge (j, i) , where each of these two choices is equally likely. Observe that in this manner, all the $2^{\binom{n}{2}}$ possible tournaments on V are

equally likely; that is, the probability space considered is symmetric. It is worth noting that we often use in applications symmetric probability spaces. In these cases, we shall sometimes refer to an element of the space as a random element, without describing explicitly the probability distribution. Thus, for example, in the proof of Proposition 1.1.1 random two-colorings of K_n were considered; that is, all possible colorings were equally likely. Similarly, in the proof of the next simple result we study random tournaments on V .

Theorem 1.2.1 If $\binom{n}{k} (1 - 2^{-k})^{n-k} < 1$ then there is a tournament on n vertices that has the property S_k .

Proof. Consider a random tournament on the set $V = \{1, \dots, n\}$. For every fixed subset K of size k of V , let A_K be the event that there is no vertex that beats all the members of K . Clearly $\Pr[A_K] = (1 - 2^{-k})^{n-k}$. This is because for each fixed vertex $v \in V - K$, the probability that v does not beat all the members of K is $1 - 2^{-k}$, and all these $n - k$ events corresponding to the various possible choices of v are independent. It follows that

$$\Pr \left[\bigvee_{\substack{K \subset V \\ |K|=k}} A_K \right] \leq \sum_{\substack{K \subset V \\ |K|=k}} \Pr[A_K] = \binom{n}{k} (1 - 2^{-k})^{n-k} < 1.$$

Therefore, with positive probability, no event A_K occurs; that is, there is a tournament on n vertices that has the property S_k .

Let $f(k)$ denote the minimum possible number of vertices of a tournament that has the property S_k . Since $\binom{n}{k} < (en/k)^k$ and $(1 - 2^{-k})^{n-k} < e^{-(n-k)/2^k}$, Theorem 1.2.1 implies that $f(k) \leq k^2 \cdot 2^k \cdot (\ln 2)(1 + o(1))$. It is not too difficult to check that $f(1) = 3$ and $f(2) = 7$. As proved by Szekeres [cf. Moon (1968)], $f(k) \geq c_1 \cdot k \cdot 2^k$.

Can one find an explicit construction of tournaments with at most c_2^k vertices having property S_k ? Such a construction is known but is not trivial; it is described in Chapter 9.

A dominating set of an undirected graph $G = (V, E)$ is a set $U \subseteq V$ such that every vertex $v \in V - U$ has at least one neighbor in U .

Theorem 1.2.2 Let $G = (V, E)$ be a graph on n vertices, with minimum degree $\delta > 1$. Then G has a dominating set of at most $n[1 + \ln(\delta + 1)]/(\delta + 1)$ vertices.

Proof. Let $p \in [0, 1]$ be, for the moment, arbitrary. Let us pick, randomly and independently, each vertex of V with probability p . Let X be the (random) set of all vertices picked and let $Y = Y_X$ be the random set of all vertices in $V - X$ that do not have any neighbor in X . The expected value of $|X|$ is clearly np . For each fixed vertex $v \in V$, $\Pr[v \in Y] = \Pr[v \text{ and its neighbors are not in } X] \leq (1 - p)^{\delta+1}$. Since the expected value of a sum of random variables is the sum of their expectations (even if they are not independent) and since the random variable $|Y|$ can be written as a sum of n indicator random variables $\chi_v (v \in V)$, where $\chi_v = 1$ if $v \in Y$ and $\chi_v = 0$ otherwise, we conclude that the expected value of $|X| + |Y|$ is at most $np + n(1 - p)^{\delta+1}$. Consequently, there is at least one choice of $X \subseteq V$ such that $|X| + |Y_X| \leq np + n(1 - p)^{\delta+1}$. The set $U = X \cup Y_X$ is clearly a dominating set of G whose cardinality is at most this size.

The above argument works for any $p \in [0, 1]$. To optimize the result we use elementary calculus. For convenience we bound $1 - p \leq e^{-p}$ (this holds for all nonnegative p and is a fairly close bound when p is small) to give the simpler bound

$$|U| \leq np + ne^{-p(\delta+1)}$$

Take the derivative of the right-hand side with respect to p and set it equal to zero. The right-hand side is minimized at

$$p = \frac{\ln(\delta + 1)}{\delta + 1}$$

Formally, we set p equal to this value in the first line of the proof. We now have $|U| \leq n[1 + \ln(\delta + 1)]/(\delta + 1)$ as claimed.

Three simple but important ideas are incorporated in the last proof. The first is the linearity of expectation; many applications of this simple, yet powerful principle appear in Chapter 2. The second is perhaps more subtle and is an example of the "alteration" principle that is discussed in Chapter 3. The random choice did not supply the required dominating set U immediately; it only supplied the set X , which has to be altered a little (by adding to it the set Y_X) to provide the required dominating set. The third involves the optimal choice of p . One often wants to make a random choice but is not certain what probability p should be used. The idea is to carry out the proof with p as a parameter giving a result that is a function of p . At the end, that p is selected which gives the optimal result. There is here yet a fourth idea that might be called asymptotic calculus. We wanted the asymptotics of $\min np + n(1 - p)^{\delta+1}$, where p ranges over $[0, 1]$. The actual minimum $p = 1 - (\delta + 1)^{-1/\delta}$ is difficult to deal with and in many similar cases precise minima are impossible to find in closed form. Rather, we give away a little bit, bounding $1 - p \leq e^{-p}$, yielding a clean bound. A good part of the art of the probabilistic method lies in finding suboptimal but clean bounds. Did we give away too much in this case? The answer depends on the emphasis for the original question. For $\delta = 3$ our rough bound gives $|U| \leq 0.596n$ while the more precise calculation gives $|U| \leq 0.496n$, perhaps a substantial difference. For δ large both methods give asymptotically $n \ln \delta / \delta$.

It can easily be deduced from the results in Alon (1990b) that the bound in Theorem 1.2.2 is nearly optimal. A non probabilistic, algorithmic proof of this theorem can be obtained by choosing the vertices for the dominating set one by one, when in each step a vertex that covers the maximum number of yet uncovered vertices is picked. Indeed, for each vertex v denote by $C(v)$ the set consisting of v together with all its neighbors. Suppose that during the process of picking vertices

the number of vertices u that do not lie in the union of the sets $C(v)$ of the vertices chosen so far is r . By the assumption, the sum of the cardinalities of the sets $C(u)$ over all such uncovered vertices u is at least $r(\delta + 1)$, and hence, by averaging, there is a vertex v that belongs to at least $r(\delta + 1)/n$ such sets $C(u)$. Adding this v to the set of chosen vertices we observe that the

number of uncovered vertices is now at most $r(1 - (\delta + 1)/n)$. It follows that in each iteration of the above procedure the number of uncovered vertices decreases by a factor of $1 - (\delta + 1)/n$ and hence after $n \ln(\delta + 1)/(\delta + 1)$ steps there will be at most $n/(\delta + 1)$ yet uncovered vertices that can now be added to the set of chosen vertices to form a dominating set of size at most equal to the one in the conclusion of Theorem 1.2.2.

Combining this with some ideas of Podderyugin and Matula, we can obtain a very efficient algorithm to decide if a given undirected graph on n vertices is, say, $n/2$ edge-connected. A cut in a graph $G = (V, E)$ is a partition of the set of vertices V into two nonempty disjoint sets $V = V_1 \cup V_2$. If $v_1 \in V_1$ and $v_2 \in V_2$ we say that the cut separates v_1 and v_2 . The size of the cut is the number of edges of G having one end in V_1 and another end in V_2 . In fact, we sometimes identify the cut with the set of these edges. The edge connectivity of G is the minimum size of a cut of G . The following lemma is due to Podderyugin and Matula (independently).

Lemma 1.2.3 Let $G = (V, E)$ be a graph with minimum degree δ and let $V = V_1 \cup V_2$ be a cut of size smaller than δ in G . Then every dominating set U of G has vertices in V_1 and in V_2 .

Proof. Suppose this is false and $U \subseteq V_1$. Choose, arbitrarily, a vertex $v \in V_2$ and let $v_1, v_2, \dots, v_\delta$ be δ of its neighbors. For each $i, 1 \leq i \leq \delta$, define an edge e_i of the given cut as follows; if $v_i \in V_1$ then $e_i = \{v, v_i\}$, otherwise, $v_i \in V_2$ and since U is dominating there is at least one vertex $u \in U$ such that $\{u, v_i\}$ is an edge; take such a u and put $e_i = \{u, v_i\}$. The δ edges $e_1, \dots, e_\delta$ are all distinct and all lie in the given cut, contradicting the assumption that its size is less than δ . This completes the proof.

Let $G = (V, E)$ be a graph on n vertices, and suppose we wish to decide if G is $n/2$ edge-connected; that is, if its edge connectivity is at least $n/2$. Matula showed, by applying Lemma 1.2.3, that this can be done in time $O(n^3)$. By the remark following the proof of Theorem 1.2.2, we can slightly improve it and get an $O(n^{8/3} \log n)$ algorithm as follows. We first check if the minimum degree δ of G is at least $n/2$. If not, G is not $n/2$ edge-connected, and the algorithm ends. Otherwise, by Theorem 1.2.2 there is a dominating set $U = \{u_1, \dots, u_k\}$ of G , where $k = O(\log n)$, and it can in fact be found in time $O(n^2)$. We now find, for each $i, 2 \leq i \leq k$, the minimum size s_i of a cut that separates u_1 from u_i . Each of these problems can be solved by solving a standard network flow problem in time $O(n^{8/3})$ [see, e.g., Tarjan (1983)]. By Lemma 1.2.3 the edge connectivity of G is simply the minimum between δ and $\min_{2 \leq i \leq k} s_i$. The total time of the algorithm is $O(n^{8/3} \log n)$, as claimed.

17.1.3 1.3 Combinatorics

A hypergraph is a pair $H = (V, E)$, where V is a finite set whose elements are called vertices and E is a family of subsets of V , called edges. It is n -uniform if each of its edges contains precisely n vertices. We say that H has property B , or that it is two-colorable if there is a two-coloring of V such that no edge is monochromatic. Let $m(n)$ denote the minimum possible number of edges of an n -uniform hypergraph that does not have property B .

Proposition 1.3.1 [Erdős (1963a)] Every n -uniform hypergraph with less than 2^{n-1} edges has property B . Therefore $m(n) \geq 2^{n-1}$.

Proof. Let $H = (V, E)$ be an n -uniform hypergraph with less than 2^{n-1} edges. Color V randomly by two colors. For each edge $e \in E$, let A_e be the event that e is monochromatic. Clearly $\Pr[A_e] = 2^{1-n}$. Therefore

$$\Pr \left[\bigvee_{e \in E} A_e \right] \leq \sum_{e \in E} \Pr[A_e] < 1$$

and there is a two-coloring without monochromatic edges.

In Section 3.5 we present a more delicate argument, due to Radhakrishnan and Srinivasan, and based on an idea of Beck, that shows that

$$m(n) \geq \Omega \left(\left(\frac{n}{\ln n} \right)^{1/2} 2^n \right)$$

The best known upper bound to $m(n)$ is found by turning the probabilistic argument "on its head." Basically, the sets become random and each coloring defines an event. Fix V with v points, where we shall later optimize v . Let χ be a coloring of V with a points in one color, $b = v - a$ points in the other. Let $S \subset V$ be a uniformly selected n -set. Then

$$\Pr[S \text{ is monochromatic under } \chi] = \frac{\binom{a}{n} + \binom{b}{n}}{\binom{v}{n}}$$

Let us assume v is even for convenience. As $\binom{y}{n}$ is convex, this expression is minimized when $a = b$. Thus

$$\Pr[S \text{ is monochromatic under } \chi] \geq p,$$

where we set

$$p = \frac{2\binom{v/2}{n}}{\binom{v}{n}}$$

for notational convenience. Now let $S_1, \dots, S_m$ be uniformly and independently chosen n -sets, m to be determined. For each coloring χ let A_χ be the event that none of the S_i are monochromatic. By the independence of the S_i

$$\Pr[A_\chi] \leq (1 - p)^m.$$

There are 2^v colorings so

$$\Pr \left[\bigvee_{\chi} A_\chi \right] \leq 2^v (1 - p)^m$$

When this quantity is less than 1 there exist $S_1, \dots, S_m$ so that no A_χ holds; that is, $S_1, \dots, S_m$ is not two-colorable and hence $m(n) \leq m$.

The asymptotics provide a fairly typical example of those encountered when employing the probabilistic method. We first use the inequality $1 - p \leq e^{-p}$. This is valid for all positive p and the terms are quite close when p is small. When

$$m = \left\lceil \frac{v \ln 2}{p} \right\rceil$$

then $2^v (1 - p)^m < 2^v e^{-pm} \leq 1$ so $m(n) \leq m$. Now we need to find v to minimize v/p . We may interpret p as twice the probability of picking n white balls from an urn with $v/2$ white and $v/2$ black balls, sampling without replacement. It is tempting to estimate p by 2^{-n+1} , the probability for sampling with replacement. This approximation would yield $m \sim v 2^{n-1} (\ln 2)$. As v gets smaller, however, the approximation becomes less accurate and, as we wish to minimize m , the trade-off becomes essential. We use a second order approximation

$$p = \frac{2\binom{v/2}{n}}{\binom{v}{n}} = 2^{1-n} \prod_{i=0}^{n-1} \frac{v-2i}{v-i} \sim 2^{1-n} e^{-n^2/2v}$$

as long as $v \gg n^{3/2}$, estimating

$$\frac{v-2i}{v-i} = 1 - \frac{i}{v} + O\left(\frac{i^2}{v^2}\right) = e^{-i/v + O(i^2/v^2)}.$$

Elementary calculus gives $v = n^2/2$ for the optimal value. The evenness of v may require a change of at most 2, which turns out to be asymptotically negligible. This yields the following result of Erdős (1964).

Theorem 1.3.2 $m(n) < (1 + o(1)) \frac{e \ln 2}{4} n^2 2^n$.

Let $\mathcal{F} = \{(A_i, B_i)\}_{i=1}^h$ be a family of pairs of subsets of an arbitrary set. We call $\mathcal{F}$ a (k, ℓ) -system if $|A_i| = k$ and $|B_i| = \ell$ for all $1 \leq i \leq h$, $A_i \cap B_i = \emptyset$ and $A_i \cap B_j \neq \emptyset$ for all distinct i, j with $1 \leq i, j \leq h$. Bollobás (1965) proved the following result, which has many interesting extensions and applications.

Theorem 1.3.3 If $\mathcal{F} = \{(A_i, B_i)\}_{i=1}^h$ is a (k, ℓ) -system then $h \leq \binom{k+\ell}{k}$.

Proof. Put $X = \bigcup_{i=1}^h (A_i \cup B_i)$ and consider a random order π of X . For each i , $1 \leq i \leq h$, let X_i be the event that all the elements of A_i precede all those of B_i in this order. Clearly $\Pr[X_i] = 1/\binom{k+\ell}{k}$. It is also easy to check that the events X_i are pairwise disjoint. Indeed, assume this is false and let π be an order in which all the elements of A_i precede those of B_i and all the elements of A_j precede those of B_j . Without loss of generality we may assume that the last element of A_i does not appear after the last element of A_j . But in this case, all elements of A_i precede all those of B_j , contradicting the fact that $A_i \cap B_j \neq \emptyset$. Therefore all the events X_i are pairwise disjoint, as claimed. It follows that

$$1 \geq \Pr\left[\bigvee_{i=1}^h X_i\right] = \sum_{i=1}^h \Pr[X_i] = h/\binom{k+\ell}{k},$$

completing the proof. $\square$

Theorem 1.3.3 is sharp, as shown by the family $\mathcal{F} = \{(A, X \setminus A) : A \subset X, |A| = k\}$, where $X = \{1, 2, \dots, k + \ell\}$.

17.1.4 1.4 Combinatorial Number Theory

A subset A of an abelian group G is called sum-free if $(A + A) \cap A = \emptyset$; that is, if there are no $a_1, a_2, a_3 \in A$ such that $a_1 + a_2 = a_3$.

Theorem 1.4.1 [Erdős (1965a)] Every set $B = \{b_1, \dots, b_n\}$ of n nonzero integers contains a sum-free subset A of size $|A| > \frac{1}{3}n$.

Proof. Let $p = 3k + 2$ be a prime, which satisfies $p > 2 \max\{|b_i|\}_{i=1}^n$ and put $C = \{k + 1, k + 2, \dots, 2k + 1\}$. Observe that C is a sum-free subset of the cyclic group $\mathbb{Z}_p$ and that

$$\frac{|C|}{p-1} = \frac{k+1}{3k+1} > \frac{1}{3}.$$

Let us choose at random an integer x , $1 \leq x < p$, according to a uniform distribution on $\{1, 2, \dots, p-1\}$, and define $d_1, \dots, d_n$ by $d_i \equiv xb_i \pmod{p}$, $0 \leq d_i < p$. Trivially, for every fixed i , $1 \leq i \leq n$, as x ranges over all numbers $1, 2, \dots, p-1$, d_i ranges over all nonzero elements of $\mathbb{Z}_p$ and hence $\Pr[d_i \in C] = |C|/(p-1) > \frac{1}{3}$. Therefore the expected number of elements b_i such that $d_i \in C$ is more than $n/3$. Consequently, there is an x , $1 \leq x < p$ and a subsequence A of B of cardinality $|A| > n/3$, such that $xa \pmod{p} \in C$ for all $a \in A$. This A is clearly sum-free, since if $a_1 + a_2 = a_3$ for some $a_1, a_2, a_3 \in A$ then $xa_1 + xa_2 \equiv xa_3 \pmod{p}$, contradicting the fact that C is a sum-free subset of $\mathbb{Z}_p$. This completes the proof. $\square$

Remark. The above proof works whenever p is a prime that does not divide any of the numbers b_i . This can be used to design an efficient deterministic algorithm for finding a sum-free subset A of size bigger than $|B|/3$ in a given set B as above. In Alon and Kleitman (1990)

it is shown that every set of n nonzero elements of an arbitrary abelian group contains a sum-free subset of more than $2n/7$ elements, and that the constant $2/7$ is best possible. The best possible constant in Theorem 1.4.1 is not known.

17.1.5 1.5 Disjoint Pairs

The probabilistic method is most striking when it is applied to prove theorems whose statement does not seem to suggest at all the need for probability. Most of the examples given in the previous sections are simple instances of such statements. In this section we describe a (slightly) more complicated result, due to Alon and Frankl (1985), which solves a conjecture of Daykin and Erdős.

Let $\mathcal{F}$ be a family of m distinct subsets of $X = \{1, 2, \dots, n\}$. Let $d(\mathcal{F})$ denote the number of disjoint pairs in $\mathcal{F}$; that is,

$$d(\mathcal{F}) = |\{\{F, F'\} : F, F' \in \mathcal{F}, F \cap F' = \emptyset\}|.$$

Daykin and Erdős conjectured that if $m = 2^{(1/2+\delta)n}$, then, for every fixed $\delta > 0$, $d(\mathcal{F}) = o(m^2)$, as n tends to infinity. This result follows from the following theorem, which is a special case of a more general result.

Theorem 1.5.1 Let $\mathcal{F}$ be a family of $m = 2^{(1/2+\delta)n}$ subsets of $X = \{1, 2, \dots, n\}$, where $\delta > 0$. Then

$$d(\mathcal{F}) < m^{2-\delta^2/2} \quad (1.1)$$

Proof. Suppose (1.1) is false and pick independently t members $A_1, A_2, \dots, A_t$ of $\mathcal{F}$ with repetitions at random, where t is a large positive integer, to be chosen later. We will show that with positive probability $|A_1 \cup A_2 \cup \dots \cup A_t| > n/2$ and still this union is disjoint to more than $2^{n/2}$ distinct subsets of X . This contradiction will establish (1.1).

In fact,

$$\begin{aligned} & \Pr[|A_1 \cup A_2 \cup \dots \cup A_t| \leq n/2] \\ & \leq \sum_{\substack{S \subset X \\ |S|=n/2}} \Pr[A_i \subset S, i = 1, \dots, t] \\ & \leq 2^n \left(2^{n/2}/2^{(1/2+\delta)n}\right)^t = 2^{n(1-\delta t)}. \end{aligned} \quad (1.2)$$

Define

$$v(B) = |\{A \in \mathcal{F} : B \cap A = \emptyset\}|.$$

Clearly,

$$\sum_{B \in \mathcal{F}} v(B) = 2d(\mathcal{F}) \geq 2m^{2-\delta^2/2}$$

Let Y be a random variable whose value is the number of members $B \in \mathcal{F}$ that are disjoint to all the A_i ($1 \leq i \leq t$). By the convexity of z^t the expected value of Y satisfies

$$\begin{aligned} \mathbb{E}[Y] &= \sum_{B \in \mathcal{F}} \left(\frac{v(B)}{m}\right)^t = \frac{1}{m^t} \cdot m \left(\sum_{B \in \mathcal{F}} v(B)^t\right) \\ &\geq \frac{1}{m^t} \cdot m \left(\frac{2d(\mathcal{F})}{m}\right)^t \geq 2m^{1-t\delta^2/2} \end{aligned}$$

Since $Y \leq m$ we conclude that

$$\Pr \left[Y \geq m^{1-t\delta^2/2} \right] \geq m^{-t\delta^2/2} \quad (1.3)$$

One can check that for $t = \lceil 1 + 1/\delta \rceil$, $m^{1-t\delta^2/2} > 2^{n/2}$ and the right-hand side of (1.3) is greater than the right-hand side of (1.2). Thus, with positive probability, $|A_1 \cup A_2 \cup \dots \cup A_t| > n/2$ and still this union is disjoint to more than $2^{n/2}$ members of F . This contradiction implies inequality (1.1).

17.1.6 1.6 Exercises

1. Prove that if there is a real p , $0 \leq p \leq 1$ such that

$$\binom{n}{k} p^{\binom{k}{2}} + \binom{n}{t} (1-p)^{\binom{t}{2}} < 1,$$

then the Ramsey number $R(k, t)$ satisfies $R(k, t) > n$. Using this, show that

$$R(4, t) \geq \Omega(t^{3/2}/(\ln t)^{3/2}).$$

2. Suppose $n \geq 4$ and let H be an n -uniform hypergraph with at most $4^{n-1}/3^n$ edges. Prove that there is a coloring of the vertices of H by four colors so that in every edge all four colors are represented.
3. (*) Prove that for every two independent, identically distributed real random variables X and Y ,

$$\Pr[|X - Y| \leq 2] \leq 3 \Pr[|X - Y| \leq 1]$$

4. (*) Let $G = (V, E)$ be a graph with n vertices and minimum degree $\delta > 10$. Prove that there is a partition of V into two disjoint subsets A and B so that $|A| \leq O(n \ln \delta / \delta)$, and each vertex of B has at least one neighbor in A and at least one neighbor in B .
5. (*) Let $G = (V, E)$ be a graph on $n \geq 10$ vertices and suppose that if we add to G any edge not in G then the number of copies of a complete graph on 10 vertices in it increases. Show that the number of edges of G is at least $8n - 36$.
6. (*) Theorem 1.2.1 asserts that for every integer $k > 0$ there is a tournament $T_k = (V, E)$ with $|V| > k$ such that for every set U of at most k vertices of T_k there is a vertex v so that all directed arcs $\{(v, u) : u \in U\}$ are in E . Show that each such tournament contains at least $\Omega(k2^k)$ vertices.
7. Let $\{(A_i, B_i), 1 \leq i \leq h\}$ be a family of pairs of subsets of the set of integers such that $|A_i| = k$ for all i and $|B_i| = l$ for all i , $A_i \cap B_i = \emptyset$ and $(A_i \cap B_j) \cup (A_j \cap B_i) \neq \emptyset$ for all $i \neq j$. Prove that $h \leq (k+l)^{k+l} / (k^k l^l)$.
8. (Prefix-free codes; Kraft Inequality). Let F be a finite collection of binary strings of finite lengths and assume no member of F is a prefix of another one. Let N_i denote the number of strings of length i in F . Prove that

$$\sum_i \frac{N_i}{2^i} \leq 1.$$

9. (*) (Uniquely decipherable codes; Kraft-McMillan Inequality). Let F be a finite collection of binary strings of finite lengths and assume that no two distinct concatenations of two finite sequences of codewords result in the same binary sequence. Let N_i denote the number of strings of length i in F . Prove that

$$\sum_i \frac{N_i}{2^i} \leq 1.$$

10. Prove that there is an absolute constant $c > 0$ with the following property. Let A be an n by n matrix with pairwise distinct entries. Then there is a permutation of the rows of A so that no column in the permuted matrix contains an increasing subsequence of length at least $c\sqrt{n}$.

The Probabilistic Lens: the Erdős-ko-rado Theorem

A family $\mathcal{F}$ of sets is called intersecting if $A, B \in \mathcal{F}$ implies $A \cap B \neq \emptyset$. Suppose $n \geq 2k$ and let $\mathcal{F}$ be an intersecting family of k -element subsets of an n -set, for definiteness $\{0, \dots, n-1\}$. The Erdős-Ko-Rado Theorem is that $|\mathcal{F}| \leq \binom{n-1}{k-1}$. This is achievable by taking the family of k -sets containing a particular point. We give a short proof due to Katona (1972).

Lemma 1 For $0 \leq s \leq n-1$ set $A_s = \{s, s+1, \dots, s+k-1\}$ where addition is modulo n . Then $\mathcal{F}$ can contain at most k of the sets A_s .

Proof. Fix some $A_s \in \mathcal{F}$. All other sets A_t that intersect A_s can be partitioned into $k-1$ pairs $\{A_{s-i}, A_{s+k-i}\}$, ($1 \leq i \leq k-1$), and the members of each such pair are disjoint. The result follows, since $\mathcal{F}$ can contain at most one member of each pair. $\square$

Now we prove the Erdős-Ko-Rado Theorem. Let a permutation σ of $\{0, \dots, n-1\}$ and $i \in \{0, \dots, n-1\}$ be chosen randomly, uniformly and independently and set $A = \{\sigma(i), \sigma(i+1), \dots, \sigma(i+k-1)\}$, addition again modulo n . Conditioning on any choice of σ the lemma gives $\Pr[A \in \mathcal{F}] \leq k/n$. Hence $\Pr[A \in \mathcal{F}] \leq k/n$. But A is uniformly chosen from all k -sets so

$$\frac{k}{n} \geq \Pr[A \in \mathcal{F}] = \frac{|\mathcal{F}|}{\binom{n}{k}}$$

and

$$|\mathcal{F}| \leq \frac{k}{n} \binom{n}{k} = \binom{n-1}{k-1}.$$

17.2 2 Linearity of Expectation

The search for truth is more precious than its possession.

- Albert Einstein

17.2.1 2.1 Basics

Let $X_1, \dots, X_n$ be random variables, $X = c_1X_1 + \dots + c_nX_n$. Linearity of expectation states that

$$E[X] = c_1E[X_1] + \dots + c_nE[X_n].$$

The power of this principle comes from there being no restrictions on the dependence or independence of the X_i . In many instances $E[X]$ can easily be calculated by a judicious decomposition into simple (often indicator) random variables X_i .

Let σ be a random permutation on $\{1, \dots, n\}$, uniformly chosen. Let $X(\sigma)$ be the number of fixed points of σ . To find $E[X]$ we decompose $X = X_1 + \dots + X_n$ where X_i is the indicator random variable of the event $\sigma(i) = i$. Then

$$E[X_i] = \Pr[\sigma(i) = i] = \frac{1}{n}$$

so that

$$E[X] = \frac{1}{n} + \dots + \frac{1}{n} = 1$$

In applications we often use that there is a point in the probability space for which $X \geq E[X]$ and a point for which $X \leq E[X]$. We have selected results with a purpose of describing this basic methodology. The following result of Szele (1943) is oftentimes considered the first use of the probabilistic method.

Theorem 2.1.1 There is a tournament T with n players and at least $n!2^{-(n-1)}$ Hamiltonian paths.

Proof. In the random tournament let X be the number of Hamiltonian paths. For each permutation σ let X_σ be the indicator random variable for σ giving a Hamiltonian path; that is, satisfying $(\sigma(i), \sigma(i+1)) \in T$ for $1 \leq i < n$. Then $X = \sum X_\sigma$ and

$$E[X] = \sum E[X_\sigma] = n!2^{-(n-1)}$$

Thus some tournament has at least $E[X]$ Hamiltonian paths.

Szele conjectured that the maximum possible number of Hamiltonian paths in a tournament on n players is at most $n!/(2 - o(1))^n$. This was proved in Alon (1990a) and is presented in The Probabilistic Lens: Hamiltonian Paths (following Chapter 4).

17.2.2 2.2 Splitting Graphs

Theorem 2.2.1 Let $G = (V, E)$ be a graph with n vertices and e edges. Then G contains a bipartite subgraph with at least $e/2$ edges.

Proof. Let $T \subseteq V$ be a random subset given by $\Pr[x \in T] = 1/2$, these choices being mutually independent. Set $B = V - T$. Call an edge $\{x, y\}$ crossing if exactly one of x, y is in T . Let X be the number of crossing edges. We decompose

$$X = \sum_{\{x, y\} \in E} X_{xy},$$

where X_{xy} is the indicator random variable for $\{x, y\}$ being crossing. Then

$$E[X_{xy}] = \frac{1}{2}$$

as two fair coin flips have probability $1/2$ of being different. Then

$$E[X] = \sum_{\{x,y\} \in E} E[X_{xy}] = \frac{e}{2}$$

Thus $X \geq e/2$ for some choice of T and the set of those crossing edges form a bipartite graph.

A more subtle probability space gives a small improvement (which is tight for complete graphs).

Theorem 2.2.2 If G has $2n$ vertices and e edges then it contains a bipartite subgraph with at least $en/(2n-1)$ edges. If G has $2n+1$ vertices and e edges then it contains a bipartite subgraph with at least $e(n+1)/(2n+1)$ edges.

Proof. When G has $2n$ vertices let T be chosen uniformly from among all n -element subsets of V . Any edge $\{x, y\}$ now has probability $n/(2n-1)$ of being crossing and the proof concludes as before. When G has $2n+1$ vertices choose T uniformly from among all n -element subsets of V and the proof is similar.

Here is a more complicated example in which the choice of distribution requires a preliminary lemma. Let $V = V_1 \cup \dots \cup V_k$, where the V_i are disjoint sets of size n . Let $h : V^k \rightarrow \{\pm 1\}$ be a two-coloring of the k -sets. A k -set E is crossing if it contains precisely one point from each V_i . For $S \subseteq V$ set $h(S) = \sum h(E)$, the sum over all k -sets $E \subseteq S$.

Theorem 2.2.3 Suppose $h(E) = +1$ for all crossing k -sets E . Then there is an $S \subseteq V$ for which

$$|h(S)| \geq c_k n^k$$

Here c_k is a positive constant, independent of n .

Lemma 2.2.4 Let P_k denote the set of all homogeneous polynomials $f(p_1, \dots, p_k)$ of degree k with all coefficients having absolute value at most one and $p_1 p_2 \dots p_k$ having coefficient one. Then for all $f \in P_k$ there exist $p_1, \dots, p_k \in [0, 1]$ with

$$|f(p_1, \dots, p_k)| \geq c_k$$

Here c_k is positive and independent of f .

Proof. Set

$$M(f) = \max_{p_1, \dots, p_k \in [0, 1]} |f(p_1, \dots, p_k)|$$

For $f \in P_k$, $M(f) > 0$ as f is not the zero polynomial. As P_k is compact and $M : P_k \rightarrow R$ is continuous, M must assume its minimum c_k .

Proof [Theorem 2.2.3]. Define a random $S \subseteq V$ by setting

$$\Pr[x \in S] = p_i, \quad x \in V_i$$

these choices being mutually independent, with p_i to be determined. Set $X = h(S)$. For each k -set E set

$$X_E = \begin{cases} h(E) & \text{if } E \subseteq S \\ 0 & \text{otherwise} \end{cases}$$

Say E has type $(a_1, \dots, a_k)$ if $|E \cap V_i| = a_i$, $1 \leq i \leq k$. For these E ,

$$E[X_E] = h(E) \Pr[E \subseteq S] = h(E) p_1^{a_1} \dots p_k^{a_k}$$

Combining terms by type

$$E[X] = \sum_{a_1 + \dots + a_k = k} p_1^{a_1} \dots p_k^{a_k} \sum_{E \text{ of type } (a_1, \dots, a_k)} h(E).$$

When $a_1 = \dots = a_k = 1$ all $h(E) = 1$ by assumption so

$$\sum_{E \text{ of type } (1, \dots, 1)} h(E) = n^k.$$

For any other type there are fewer than n^k terms, each ± 1 , so

$$\left| \sum_{E \text{ of type } (a_1, \dots, a_k)} h(E) \right| \leq n^k.$$

Thus

$$E[X] = n^k f(p_1, \dots, p_k)$$

where $f \in P_k$, as defined by Lemma 2.2.4.

Now select $p_1, \dots, p_k \in [0, 1]$ with $|f(p_1, \dots, p_k)| \geq c_k$. Then

$$E[|X|] \geq |E[X]| \geq c_k n^k.$$

Some particular value of $|X|$ must exceed or equal its expectation. Hence there is a particular set $S \subseteq V$ with $|X| = |h(S)| \geq c_k n^k$.

Theorem 2.2.3 has an interesting application to Ramsey Theory. It is known [see Erdős (1965b)] that given any coloring with two colors of the k -sets of an n -set there exist k disjoint m -sets, $m = \Theta((\ln n)^{1/(k-1)})$, so that all crossing k -sets are the same color. From Theorem 2.2.3 there then exists a set of size $\Theta((\ln n)^{1/(k-1)})$, at least $\frac{1}{2} + \epsilon_k$ of whose k -sets are the same color. This is somewhat surprising since it is known that there are colorings in which the largest monochromatic set has size at most the $k - 2$ -fold logarithm of n .

17.2.3 2.3 Two Quickies

Linearity of expectation sometimes gives very quick results.

Theorem 2.3.1 There is a two-coloring of K_n with at most

$$\binom{n}{a} 2^{1 - \binom{a}{2}}$$

monochromatic K_a .

Proof [Outline]. Take a random coloring. Let X be the number of monochromatic K_a and find $E[X]$. For some coloring the value of X is at most this expectation. $\square$

In Chapter 16 it is shown how such a coloring can be found deterministically and efficiently.

Theorem 2.3.2 There is a two-coloring of $K_{m,n}$ with at most

$$\binom{m}{a} \binom{n}{b} 2^{1 - ab}$$

monochromatic $K_{a,b}$.

Proof [Outline]. Take a random coloring. Let X be the number of monochromatic $K_{a,b}$ and find $E[X]$. For some coloring the value of X is at most this expectation. $\square$

17.2.4 2.4 Balancing Vectors

The next result has an elegant non probabilistic proof, which we defer to the end of this chapter. Here $|v|$ is the usual Euclidean norm.

Theorem 2.4.1 Let $v_1, \dots, v_n \in \mathbb{R}^n$, all $|v_i| = 1$. Then there exist $\epsilon_1, \dots, \epsilon_n = \pm 1$ so that

$$|\epsilon_1 v_1 + \dots + \epsilon_n v_n| \leq \sqrt{n},$$

and also there exist $\epsilon_1, \dots, \epsilon_n = \pm 1$ so that

$$|\epsilon_1 v_1 + \dots + \epsilon_n v_n| \geq \sqrt{n}.$$

Proof. Let $\epsilon_1, \dots, \epsilon_n$ be selected uniformly and independently from $\{-1, +1\}$. Set

$$X = |\epsilon_1 v_1 + \dots + \epsilon_n v_n|^2.$$

Then

$$X = \sum_{i=1}^n \sum_{j=1}^n \epsilon_i \epsilon_j v_i \cdot v_j.$$

Thus

$$\mathbb{E}[X] = \sum_{i=1}^n \sum_{j=1}^n v_i \cdot v_j \mathbb{E}[\epsilon_i \epsilon_j].$$

When $i \neq j$, $\mathbb{E}[\epsilon_i \epsilon_j] = \mathbb{E}[\epsilon_i] \mathbb{E}[\epsilon_j] = 0$. When $i = j$, $\epsilon_i^2 = 1$ so $\mathbb{E}[\epsilon_i^2] = 1$. Thus

$$\mathbb{E}[X] = \sum_{i=1}^n v_i \cdot v_i = n.$$

Hence there exist specific $\epsilon_1, \dots, \epsilon_n = \pm 1$ with $X \geq n$ and with $X \leq n$. Taking square roots gives the theorem.

The next result includes part of Theorem 2.4.1 as a linear translation of the $p_1 = \dots = p_n = 1/2$ case.

Theorem 2.4.2 Let $v_1, \dots, v_n \in \mathbb{R}^n$, all $|v_i| \leq 1$. Let $p_1, \dots, p_n \in [0, 1]$ be arbitrary and set $w = p_1 v_1 + \dots + p_n v_n$. Then there exist $\epsilon_1, \dots, \epsilon_n \in \{0, 1\}$ so that, setting $v = \epsilon_1 v_1 + \dots + \epsilon_n v_n$,

$$|w - v| \leq \frac{\sqrt{n}}{2}$$

Proof. Pick ϵ_i independently with

$$\Pr[\epsilon_i = 1] = p_i, \quad \Pr[\epsilon_i = 0] = 1 - p_i.$$

The random choice of ϵ_i gives a random v and a random variable

$$X = |w - v|^2$$

We expand

$$X = \left| \sum_{i=1}^n (p_i - \epsilon_i) v_i \right|^2 = \sum_{i=1}^n \sum_{j=1}^n v_i \cdot v_j (p_i - \epsilon_i) (p_j - \epsilon_j)$$

so that

$$E[X] = \sum_{i=1}^n \sum_{j=1}^n v_i \cdot v_j E[(p_i - \epsilon_i)(p_j - \epsilon_j)].$$

For $i \neq j$,

$$E[(p_i - \epsilon_i)(p_j - \epsilon_j)] = E[p_i - \epsilon_i] E[p_j - \epsilon_j] = 0.$$

For $i = j$,

$$E[(p_i - \epsilon_i)^2] = p_i(p_i - 1)^2 + (1 - p_i)p_i^2 = p_i(1 - p_i) \leq \frac{1}{4}$$

($E[(p_i - \epsilon_i)^2] = \text{Var}[\epsilon_i]$, the variance to be discussed in Chapter 4.) Thus

$$E[X] = \sum_{i=1}^n p_i(1 - p_i)|v_i|^2 \leq \frac{1}{4} \sum_{i=1}^n |v_i|^2 \leq \frac{n}{4}$$

and the proof concludes as in that of Theorem 2.4.1.

17.2.5 2.5 Unbalancing Lights

Theorem 2.5.1 Let $a_{ij} = \pm 1$ for $1 \leq i, j \leq n$. Then there exist $x_i, y_j = \pm 1$, $1 \leq i, j \leq n$ so that

$$\sum_{i=1}^n \sum_{j=1}^n a_{ij} x_i y_j \geq \left(\sqrt{\frac{2}{\pi}} + o(1) \right) n^{3/2}$$

This result has an amusing interpretation. Let an $n \times n$ array of lights be given, each either on ($a_{ij} = +1$) or off ($a_{ij} = -1$). Suppose for each row and each column there is a switch so that if the switch is pulled ($x_i = -1$ for row i and $y_j = -1$ for column j) all of the lights in that line are "switched": on to off or off to on. Then for any initial configuration it is possible to perform switches so that the number of lights on minus the number of lights off is at least $(\sqrt{2/\pi} + o(1))n^{3/2}$.

Proof. Forget the x 's. Let $y_1, \dots, y_n = \pm 1$ be selected independently and uniformly and set

$$R_i = \sum_{j=1}^n a_{ij} y_j, \quad R = \sum_{i=1}^n |R_i|.$$

Fix i . Regardless of a_{ij} , $a_{ij} y_j$ is ± 1 with probability $1/2$ and their values (over j) are independent; that is, whatever the i th row is initially after random switching it becomes a uniformly distributed row, all 2^n possibilities equally likely. Thus R_i has distribution S_n - the distribution of the sum of n independent uniform $\{-1, 1\}$ random variables - and so

$$E[|R_i|] = E[|S_n|] = \left(\sqrt{\frac{2}{\pi}} + o(1) \right) \sqrt{n}$$

These asymptotics may be found by estimating S_n by $\sqrt{n}N$, where N is standard normal and using elementary calculus. Alternatively, a closed form

$$E[|S_n|] = n2^{1-n} \binom{n-1}{\lfloor (n-1)/2 \rfloor}$$

may be derived combinatorially (a problem in the 1974 Putnam competition!) and the asymptotics follows from Stirling's formula.

Now apply linearity of expectation to R :

$$\mathbb{E}[R] = \sum_{i=1}^n \mathbb{E}[|R_i|] = \left(\sqrt{\frac{2}{\pi}} + o(1) \right) n^{3/2}.$$

There exist $y_1, \dots, y_n = \pm 1$ with R at least this value. Finally, pick x_i with the same sign as R_i so that

$$\sum_{i=1}^n x_i \sum_{j=1}^n a_{ij} y_j = \sum_{i=1}^n x_i R_i = \sum_{i=1}^n |R_i| = R \geq \left(\sqrt{\frac{2}{\pi}} + o(1) \right) n^{3/2}$$

Another result on unbalancing lights appears in The Probabilistic Lens: Unbalancing Lights (following Chapter 13). The existence of Hadamard matrices and the discussion in Section 9.1 show that the estimate in the last theorem cannot be improved to anything bigger than $n^{3/2}$.

17.2.6 2.6 Without Coin Flips

A non probabilistic proof of Theorem 2.2.1 may be given by placing each vertex in either T or B sequentially. At each stage place x in either T or B so that at least half of the edges from x to previous vertices are crossing. With this effective algorithm at least half the edges will be crossing.

There is also a simple sequential algorithm for choosing signs in Theorem 2.4.1. When the sign for v_i is to be chosen, a partial sum $w = \epsilon_1 v_1 + \dots + \epsilon_{i-1} v_{i-1}$ has been calculated. Now if it is desired that the sum be small select $\epsilon_i = \pm 1$ so that $\epsilon_i v_i$ makes an obtuse (or right) angle with w . If the sum need be big make the angle acute or right. In the extreme case when all angles are right angles, Pythagoras and induction give that the final w has norm $\sqrt{n}$, otherwise it is either less than $\sqrt{n}$ or greater than $\sqrt{n}$ as desired.

For Theorem 2.4.2 a greedy algorithm produces the desired ϵ_i . Given $v_1, \dots, v_n \in \mathbb{R}^n, p_1, \dots, p_n \in [0, 1]$ suppose $\epsilon_1, \dots, \epsilon_{s-1} \in \{0, 1\}$ have already been chosen. Set $w_{s-1} = \sum_{i=1}^{s-1} (p_i - \epsilon_i) v_i$, the partial sum. Select ϵ_s so that

$$w_s = w_{s-1} + (p_s - \epsilon_s) v_s = \sum_{i=1}^s (p_i - \epsilon_i) v_i$$

has minimal norm. A random $\epsilon_s \in \{0, 1\}$ chosen with $\Pr[\epsilon_s = 1] = p_s$ gives

$$\begin{aligned} \mathbb{E}[|w_s|^2] &= |w_{s-1}|^2 + 2w_{s-1} \cdot v_s \mathbb{E}[p_s - \epsilon_s] + |v_s|^2 \mathbb{E}[p_s - \epsilon_s]^2 \\ &= |w_{s-1}|^2 + p_s (1 - p_s) |v_s|^2 \end{aligned}$$

so for some choice of $\epsilon_s \in \{0, 1\}$,

$$|w_s|^2 \leq |w_{s-1}|^2 + p_s (1 - p_s) |v_s|^2.$$

As this holds for all $1 \leq s \leq n$ (taking $w_0 = 0$), the final

$$|w_n|^2 \leq \sum_{i=1}^n p_i (1 - p_i) |v_i|^2$$

While the proofs appear similar, a direct implementation of the proof of Theorem 2.4.2 to find $\epsilon_1, \dots, \epsilon_n$ might take an exhaustive search with exponential time. In applying the greedy algorithm at the s th stage one makes two calculations of $|w_s|^2$, depending on whether $\epsilon_s = 0$ or 1 , and picks that ϵ_s giving the smaller value. Hence there are only a linear number of calculations of norms to be made and the entire algorithm takes only quadratic time. In Chapter 16 we discuss several similar examples in a more general setting.

17.2.7 2.7 Exercises

1. Suppose $n \geq 2$ and let $H = (V, E)$ be an n -uniform hypergraph with $|E| = 4^{n-1}$ edges. Show that there is a coloring of V by four colors so that no edge is monochromatic.
2. Prove that there is a positive constant c so that every set A of n nonzero reals contains a subset $B \subset A$ of size $|B| \geq cn$ so that there are no $b_1, b_2, b_3, b_4 \in B$ satisfying

$$b_1 + 2b_2 = 2b_3 + 2b_4$$

3. Prove that every set of n non zero real numbers contains a subset A of strictly more than $n/3$ numbers such that there are no $a_1, a_2, a_3 \in A$ satisfying $a_1 + a_2 = a_3$.
4. Suppose $p > n > 10m^2$, with p prime, and let $0 < a_1 < a_2 < \dots < a_m < p$ be integers. Prove that there is an integer $x, 0 < x < p$ for which the m numbers

$$(xa_i \bmod p) \bmod n, \quad 1 \leq i \leq m$$

are pairwise distinct.

5. Let H be a graph, and let $n > |V(H)|$ be an integer. Suppose there is a graph on n vertices and t edges containing no copy of H , and suppose that $tk > n^2 \log_e n$. Show that there is a coloring of the edges of the complete graph on n vertices by k colors with no monochromatic copy of H .
6. (*) Prove, using the technique shown in The Probabilistic Lens: Hamiltonian Paths, that there is a constant $c > 0$ such that for every even $n \geq 4$ the following holds: For every undirected complete graph K on n vertices whose edges are colored red and blue, the number of alternating Hamiltonian cycles in K (i.e., properly edge-colored cycles of length n) is at most

$$n^c \frac{n!}{2^n}.$$

7. Let $\mathcal{F}$ be a family of subsets of $N = \{1, 2, \dots, n\}$, and suppose there are no $A, B \in \mathcal{F}$ satisfying $A \subset B$. Let $\sigma \in S_n$ be a random permutation of the elements of N and consider the random variable

$$X = |\{i : \{\sigma(1), \sigma(2), \dots, \sigma(i)\} \in \mathcal{F}\}|.$$

By considering the expectation of X prove that $|\mathcal{F}| \leq \binom{n}{\lfloor n/2 \rfloor}$.

8. (*) Let X be a collection of pairwise orthogonal unit vectors in $\mathbb{R}^n$ and suppose the projection of each of these vectors on the first k coordinates is of Euclidean norm at least ϵ . Show that $|X| \leq k/\epsilon^2$, and this is tight for all $\epsilon^2 = k/2^r < 1$.
9. Let $G = (V, E)$ be a bipartite graph with n vertices and a list $S(v)$ of more than $\log_2 n$ colors associated with each vertex $v \in V$. Prove that there is a proper coloring of G assigning to each vertex v a color from its list $S(v)$.

The Probabilistic Lens: Brégman's Theorem

Let $A = [a_{ij}]$ be an $n \times n$ matrix with all $a_{ij} \in \{0, 1\}$. Let $r_i = \sum_{1 \leq j \leq n} a_{ij}$ be the number of ones in the i th row. Let S be the set of permutations $\sigma \in S_n$ with $a_{i, \sigma i} = 1$ for $1 \leq i \leq n$. Then the permanent $\text{per}(A)$ is simply $|S|$. The following result was conjectured by Minc and proved by Brégman (1973). The proof presented here is similar to that of Schrijver (1978).

Theorem 1 [Brégman's Theorem] $\text{per}(A) \leq \prod_{1 \leq i \leq n} (r_i!)^{1/r_i}$.

Pick $\sigma \in S$ and $\tau \in S_n$ independently and uniformly. Set $A^{(1)} = A$. Let $R_{\tau 1}$ be the number of ones in row $\tau 1$ in $A^{(1)}$. Delete row $\tau 1$ and column $\sigma \tau 1$ from $A^{(1)}$ to give $A^{(2)}$. In general, let $A^{(i)}$ denote A with rows $\tau 1, \dots, \tau(i-1)$ and columns $\sigma \tau 1, \dots, \sigma \tau(i-1)$ deleted and let $R_{\tau i}$ denote the number of ones of row τi in $A^{(i)}$. (This is nonzero as the $\sigma \tau i$ th column has a one.) Set

$$L = L(\sigma, \tau) = \prod_{1 \leq i \leq n} R_{\tau i}$$

We think, roughly, of L as Lazyman's permanent calculation. There are $R_{\tau 1}$ choices for a one in row $\tau 1$, each of which leads to a different subpermanent calculation. Instead, Lazyman takes the factor $R_{\tau 1}$, takes the one from permutation σ , and examines $A^{(2)}$. As $\sigma \in S$ is chosen uniformly Lazyman tends toward the high subpermanents and so it should not be surprising that he tends to overestimate the permanent. To make this precise we define the geometric mean $G[Y]$. If $Y > 0$ takes values $a_1, \dots, a_s$ with probabilities $p_1, \dots, p_s$, respectively, then $G[Y] = \prod a_i^{p_i}$. Equivalently, $G[Y] = e^{\mathbb{E}[\ln Y]}$. Linearity of expectation translates into the geometric mean of a product being the product of the geometric means.

Claim 2.7.1 $\text{per}(A) \leq G[L]$.

Proof. We show this for any fixed τ . Set $\tau 1 = 1$ for convenience of notation. We use induction on the size of the matrix. Reorder, for convenience, so that the first row has ones in the first r columns, where $r = r_1$. For $1 \leq j \leq r$ let t_j be the permanent of A with the first row and j th column removed or, equivalently, the number of $\sigma \in S$ with $\sigma 1 = j$. Set

$$t = \frac{t_1 + \dots + t_r}{r}$$

so that $\text{per}(A) = rt$. Conditioning on $\sigma 1 = j$, $R_2 \cdots R_n$ is Lazyman's calculation of $\text{per}(A^{(2)})$, where $A^{(2)}$ is A with the first row and j th column removed. By induction

$$G[R_2 \cdots R_n \mid \sigma 1 = j] \geq t_j$$

and so

$$G[L] \geq \prod_{j=1}^r (rt_j)^{t_j / \text{per}(A)} = r \prod_{j=1}^r t_j^{t_j / rt}$$

Lemma 2 $\left(\prod_{j=1}^r t_j^{t_j} \right)^{1/r} \geq t^t$.

Proof. Taking logarithms, this is equivalent to

$$\frac{1}{r} \sum_{j=1}^r t_j \ln t_j \geq t \ln t$$

which follows from the convexity of the function $f(x) = x \ln x$.

Applying the lemma,

$$G[L] \geq r \prod_{j=1}^r t_j^{t_j / rt} \geq r (t^t)^{1/t} = rt = \text{per}(A)$$

Now we calculate $G[L]$ conditional on a fixed σ . For convenience of notation reorder so that $\sigma i = i$, all i , and assume that the first row has ones in precisely the first r_1 columns. With τ selected uniformly the columns $1, \dots, r_1$ are deleted in order uniform over all $r_1!$ possibilities. R_1 is the number of those columns remaining when the first column is to be deleted. As the first column is equally likely to be in any position among those r_1 columns R_1 is uniformly distributed from 1 to r_1 and $G[R_1] = (r_1!)^{1/r_1}$. "Linearity" then gives

$$G[L] = G\left[\prod_{i=1}^n R_i\right] = \prod_{i=1}^n G[R_i] = \prod_{i=1}^n (r_i!)^{1/r_i}$$

The overall $G[L]$ is the geometric mean of the conditional $G[L]$ and hence has the same value. That is,

$$\text{per}(A) \leq G[L] = \prod_{i=1}^n (r_i!)^{1/r_i}$$

17.3 3 Alterations

Beauty is the first test: there is no permanent place in the world for ugly mathematics.

- G. H. Hardy

The basic probabilistic method was described in Chapter 1 as follows: Trying to prove that a structure with certain desired properties exists, one defines an appropriate probability space of structures and then shows that the desired properties hold in this space with positive probability. In this chapter we consider situations where the "random" structure does not have all the desired properties but may have a few "blemishes." With a small alteration we remove the blemishes, giving the desired structure.

17.3.1 3.1 Ramsey Numbers

Recall from Section 1.1 that $R(k, l) > n$ means there exists a two-coloring of the edges of K_n by red and blue so that there is neither a red K_k nor a blue K_l .

Theorem 3.1.1 For any integer n , $R(k, k) > n - \binom{n}{k} 2^{1-\binom{k}{2}}$.

Proof. Consider a random two-coloring of the edges of K_n obtained by coloring each edge independently either red or blue, where each color is equally likely. For any set R of k vertices let X_R be the indicator random variable for the event that the induced subgraph of K_n on R is monochromatic. Set $X = \sum X_R$, the sum over all such R . By linearity of expectation, $E[X] = \sum E[X_R] = m$ with $m = \binom{n}{k} 2^{1-\binom{k}{2}}$. Thus there exists a two-coloring for which $X \leq m$. Fix such a coloring. Remove from K_n one vertex from each monochromatic k -set. At most m vertices have been removed (we may have "removed" the same vertex more than once but this only helps) so s vertices remain with $s \geq n - m$. This coloring on these s points has no monochromatic k -set.

We are left with the "calculus" problem of finding that n which will optimize the inequality. Some analysis shows that we should take $n \sim e^{-1} k 2^{k/2} (1 + o(1))$ giving

$$R(k, k) > \frac{1}{e} (1 + o(1)) k 2^{k/2}.$$

A careful examination of Proposition 1.1.1 gives the lower bound

$$R(k, k) > \frac{1}{e\sqrt{2}} (1 + o(1)) k 2^{k/2}$$

The more powerful Lovász Local Lemma (see Chapter 5) gives

$$R(k, k) > \frac{\sqrt{2}}{e}(1 + o(1))k2^{k/2}$$

The distinctions between these bounds may be considered inconsequential since the best known upper bound for $R(k, k)$ is $(4 + o(1))^k$. The upper bounds do not involve probabilistic methods and may be found, for example, in Graham, Rothschild and Spencer (1990). We give all three lower bounds in following our philosophy of emphasizing methodologies rather than results.

In dealing with the off-diagonal Ramsey numbers the distinction between the basic method and the alteration is given in the following two results.

Theorem 3.1.2 If there exists $p \in [0, 1]$ with

$$\binom{n}{k}p^{\binom{k}{2}} + \binom{n}{l}(1-p)^{\binom{l}{2}} < 1$$

then $R(k, l) > n$.

Theorem 3.1.3 For all integers n and $p \in [0, 1]$,

$$R(k, l) > n - \binom{n}{k}p^{\binom{k}{2}} - \binom{n}{l}(1-p)^{\binom{l}{2}}.$$

Proof. In both cases we consider a random two-coloring of K_n obtained by coloring each edge independently either red or blue, where each edge is red with probability p . Let X be the number of red k -sets plus the number of blue l -sets. Linearity of expectation gives

$$E[X] = \binom{n}{k}p^{\binom{k}{2}} + \binom{n}{l}(1-p)^{\binom{l}{2}}$$

For Theorem 3.1.2, $E[X] < 1$ so there exists a two-coloring with $X = 0$. For Theorem 3.1.3 there exists a two-coloring with s "bad" sets (either red k -sets or blue l -sets), $s \leq E[X]$. Removing one point from each bad set gives a coloring of at least $n - s$ points with no bad sets.

The asymptotics of Theorems 3.1.2 and 3.1.3 can get fairly complex. Oftentimes Theorem 3.1.3 gives a substantial improvement on Theorem 3.1.2. Even further improvements may be found using the Lovász Local Lemma. These bounds have been analyzed in Spencer (1977).

17.3.2 3.2 Independent Sets

Here is a short and sweet argument that gives roughly half of the celebrated Turán's Theorem. $\alpha(G)$ is the independence number of a graph G ; $\alpha(G) \geq t$ means there exist t vertices with no edges between them.

Theorem 3.2.1 Let $G = (V, E)$ have n vertices and $nd/2$ edges, $d \geq 1$. Then $\alpha(G) \geq n/2d$.

Proof. Let $S \subseteq V$ be a random subset defined by

$$\Pr[v \in S] = p$$

p to be determined, the events $v \in S$ being mutually independent. Let $X = |S|$ and let Y be the number of edges in $G|_S$. For each $e = \{i, j\} \in E$ let Y_e be the indicator random variable for the event $i, j \in S$ so that $Y = \sum_{e \in E} Y_e$. For any such e ,

$$E[Y_e] = \Pr[i, j \in S] = p^2,$$

so by linearity of expectation,

$$E[Y] = \sum_{e \in E} E[Y_e] = \frac{nd}{2} p^2$$

Clearly $E[X] = np$, so, again by linearity of expectation,

$$E[X - Y] = np - \frac{nd}{2} p^2$$

We set $p = 1/d$ (here using $d \geq 1$) to maximize this quantity, giving

$$E[X - Y] = \frac{n}{2d}$$

Thus there exists a specific S for which the number of vertices of S minus the number of edges in S is at least $n/2d$. Select one vertex from each edge of S and delete it. This leaves a set S^* with at least $n/2d$ vertices. All edges having been destroyed, S^* is an independent set.

The full result of Turán is given in The Probabilistic Lens: Turán's Theorem (following Chapter 6).

17.3.3 3.3 Combinatorial Geometry

For a set S of n points in the unit square U , let $T(S)$ be the minimum area of a triangle whose vertices are three distinct points of S . Put $T(n) = \max T(S)$, where S ranges over all sets of n points in U . Heilbronn conjectured that $T(n) = O(1/n^2)$. This conjecture was disproved by Komlós, Pintz and Szemerédi (1982) who showed, by a rather involved probabilistic construction, that there is a set S of n points in U such that $T(S) = \Omega(\log n/n^2)$. As this argument is rather complicated, we only present here a simpler one showing that $T(n) = \Omega(1/n^2)$.

Theorem 3.3.1 There is a set S of n points in the unit square U such that $T(S) \geq 1/(100n^2)$.

Proof. We first make a calculation. Let P, Q, R be independently and uniformly selected from U and let $\mu = \mu(PQR)$ denote the area of the triangle PQR . We bound $\Pr[\mu \leq \epsilon]$ as follows. Let x be the distance from P to Q so that

$$\Pr[b \leq x \leq b + \Delta b] \leq \pi(b + \Delta b)^2 - \pi b^2$$

and in the limit $\Pr[b \leq x \leq b + db] \leq 2\pi b db$. Given P, Q at distance b , the altitude from R to the line PQ must have height $h \leq 2\epsilon/b$ and so R must lie in a strip of width $4\epsilon/b$ and length at most $\sqrt{2}$. This occurs with probability at most $4\sqrt{2}\epsilon/b$. As $0 \leq b \leq \sqrt{2}$ the total probability is bounded by

$$\int_0^{\sqrt{2}} (2\pi b)(4\sqrt{2}\epsilon/b) db = 16\pi\epsilon$$

Now let $P_1, \dots, P_{2n}$ be selected uniformly and independently in U and let X denote the number of triangles $P_i P_j P_k$ with area less than $1/(100n^2)$. For each particular i, j, k the probability of this occurring is less than $0.6n^{-2}$ and so

$$E[X] \leq \binom{2n}{3} (0.6n^{-2}) < n$$

Thus there exists a specific set of $2n$ vertices with fewer than n triangles of area less than $1/(100n^2)$. Delete one vertex from the set from each such triangle. This leaves at least n vertices and now no triangle has area less than $1/(100n^2)$.

We note the following construction of Erdős showing $T(n) \geq 1/(2(n-1)^2)$ with n prime. On $[0, n-1] \times [0, n-1]$ consider the n points (x, x^2) , where x^2 is reduced mod n (more formally,

(x, y) where $y \equiv x^2 \pmod n$ and $0 \leq y < n$). If some three points of this set were collinear they would line on a line $y = mx + b$ and m would be a rational number with denominator less than n . But then in $\mathbb{Z}_n^2$ the parabola $y = x^2$ would intersect the line $y = mx + b$ at three points, so that the quadratic $x^2 - mx - b$ would have three distinct roots, an impossibility. Triangles between lattice points in the plane have as their areas either half-integers or integers, hence the areas must be at least $1/2$. Contracting the plane by an $n - 1$ factor in both coordinates gives the desired set. While this gem does better than Theorem 3.3.1 it does not lead to the improvements of Komlós, Pintz and Szemerédi.

17.3.4 3.4 Packing

Let C be a bounded measurable subset of $\mathbb{R}^d$ and let $B(x)$ denote the cube $[0, x]^d$ of side x . A packing of C into $B(x)$ is a family of mutually disjoint copies of C , all lying inside $B(x)$. Let $f(x)$ denote the largest size of such a family. The packing constant $\delta = \delta(C)$ is defined by

$$\delta(C) = \mu(C) \lim_{x \rightarrow \infty} f(x)x^{-d}$$

where $\mu(C)$ is the measure of C . This is the maximal proportion of space that may be packed by copies of C (this limit can be proved always to exist but even without that result the following result holds with $\lim$ replaced by $\liminf$.)

Theorem 3.4.1 Let C be bounded, convex and centrally symmetric around the origin. Then $\delta(C) \geq 2^{-d-1}$.

Proof. Let P, Q be selected independently and uniformly from $B(x)$ and consider the event $(C + P) \cap (C + Q) \neq \emptyset$. For this to occur we must have, for some $c_1, c_2 \in C$,

$$P - Q = c_1 - c_2 = 2 \frac{c_1 - c_2}{2} \in 2C$$

by central symmetry and convexity. The event $P \in Q + 2C$ has probability at most $\mu(2C)x^{-d}$ for each given Q , hence

$$\Pr[(C + P) \cap (C + Q) \neq \emptyset] \leq \mu(2C)x^{-d} = 2^d x^{-d} \mu(C).$$

Now let $P_1, \dots, P_n$ be selected independently and uniformly from $B(x)$ and let X be the number of $i < j$ with $(C + P_i) \cap (C + P_j) \neq \emptyset$. From linearity of expectation,

$$\mathbb{E}[X] \leq \frac{n^2}{2} 2^d x^{-d} \mu(C)$$

Hence there exists a specific choice of n points with fewer than that many intersecting copies of C . For each P_i, P_j with $(C + P_i) \cap (C + P_j) \neq \emptyset$ remove either P_i or P_j from the set. This leaves at least $n - (n^2/2) 2^d x^{-d} \mu(C)$ nonintersecting copies of C . Set $n = x^d 2^{-d} / \mu(C)$ to maximize this quantity, so that there are at least $x^d 2^{-d-1} / \mu(C)$ nonintersecting copies of C . These do not all lie inside $B(x)$ but, letting w denote an upper bound on the absolute values of the coordinates of the points of C , they do all lie inside a cube of side $x + 2w$. Hence

$$f(x + 2w) \geq x^d 2^{-d-1} / \mu(C)$$

and so $\delta(C) \geq \lim_{x \rightarrow \infty} \mu(C) f(x + 2w)(x + 2w)^{-d} \geq 2^{-d-1}$.

A simple greedy algorithm does somewhat better. Let $P_1, \dots, P_m$ be any maximal subset of $[0, x]^d$ with the property that the sets $C + P_i$ are disjoint. We have seen that $C + P_i$ overlaps $C + P$ if and only if $P \in 2C + P_i$. Hence the sets $2C + P_i$ must cover $[0, x]^d$. As each such set

has measure $\mu(2C) = 2^d \mu(C)$ we must have $m \geq x^d 2^{-d} / \mu(C)$. As before, all sets $C + P_i$ lie in a cube of side $x + 2w$, w a constant, so that

$$f(x + 2w) \geq m \geq x^d 2^{-d} / \mu(C)$$

and so

$$\delta(C) \geq 2^{-d}$$

A still further improvement appears in *The Probabilistic Lens: Efficient Packing* (following Chapter 14).

17.3.5 3.5 Recoloring

Suppose that a random coloring leaves a set of blemishes. Here we apply a random recoloring to the blemishes to remove them. If the recoloring is too weak then not all the blemishes are removed. If the recoloring is too strong then new blemishes are created. The recoloring is given a parameter p and these two possibilities are decreasing and increasing functions of p . Calculus then points us to the optimal p .

We use the notation of Section 1.3 on property B: $m(n) > m$ means that given any n -uniform hypergraph $H = (V, E)$ with m edges there exists a two-coloring of V so that no edge is monochromatic. Beck (1978) improved Erdős' 1963 bound to $m(n) = \Omega(2^n n^{1/3})$. Building on his methods, Radhakrishnan and Srinivasan (2000) proved $m(n) = \Omega(2^n (n/\ln n)^{1/2})$ and it is that proof we shall give. While this proof is neither long nor technically complex it has a number of subtle and beautiful steps and it is not surprising that it took more than thirty-five years to find it. That said, the upper and lower bounds on $m(n)$ remain quite far apart!

Theorem 3.5.1 If there exists $p \in [0, 1]$ with $k(1-p)^n + k^2 p < 1$ then $m(n) > 2^{n-1} k$.

Corollary 3.5.2 $m(n) = \Omega(2^n (n/\ln n)^{1/2})$.

Proof. Bound $1 - p \leq e^{-p}$. The function $ke^{-pn} + k^2 p$ is minimized at $p = \ln(n/k)/n$. Substituting back in, if

$$\frac{k^2}{n}(1 + \ln(n/k)) < 1$$

then the condition of Theorem 3.5.1 holds. This inequality is true when $k = c(n/\ln n)^{1/2}$ for any $c < \sqrt{2}$ with n sufficiently large.

The condition of Theorem 3.5.1 is somewhat typical; one wants the total failure probability to be less than 1 and there are two types of failure. Oftentimes one finds reasonable bounds by requiring the stronger condition that each failure type has probability less than one-half. Here $k^2 p \leq \frac{1}{2}$ gives $p \leq \frac{1}{2} k^{-2}$. Plugging the maximal possible p into the second inequality $k(1-p)^n \leq \frac{1}{2}$ gives $2k^2 \ln(2k) \leq n$. This again holds when $k = c(n/\ln n)^{1/2}$ though now we have the weaker condition $c < 1$. We recommend this rougher approach as a first attempt at a problem, when the approximate range of the parameters is still in doubt. The refinements of calculus can be placed in the published work!

Proof [Theorem 3.5.1]. Fix $H = (V, E)$ with $m = 2^{n-1} k$ edges and p satisfying the condition. We describe a randomized algorithm that yields a coloring of V . It is best to preprocess the randomness: Each $v \in V$ flips a first coin, which comes up heads with probability $\frac{1}{2}$ and a second coin, which comes up heads (representing potential recoloration) with probability p . In addition (and importantly), the vertices of V are ordered randomly.

Step 1. Color each $v \in V$ red if its first coin was heads, otherwise blue. Call this the first coloring. Let D (for dangerous) denote the set of $v \in V$ that lie in some (possibly many) monochromatic $e \in E$.

Step 2. Consider the elements of D sequentially in the (random) order of V . When d is being considered call it still dangerous if there is some (possibly many) $e \in H$ containing d that was monochromatic in the first coloring and for which no vertices have yet changed color. If d is not still dangerous then do nothing. But if it is still dangerous then check its second coin. If it is heads then change the color of d , otherwise do nothing. We call the coloring at the time of termination the final coloring.

We say the algorithm fails if some $e \in H$ is monochromatic in the final coloring. We shall bound the failure probability by $k(1-p)^n + k^2p$. The assumption of Theorem 3.5.1 then assures us that with positive probability the algorithm succeeds. This, by our usual magic, means that there is some running of the algorithm which yields a final coloring with no monochromatic e ; that is, there exists a two-coloring of V with no monochromatic edge. For convenience, we bound the probability that some $e \in H$ is red in the final coloring; the failure probability for the algorithm is at most twice that.

An $e \in E$ can be red in the final coloring in two ways. Either e was red in the first coloring and remained red through to the final coloring or e was not red in the first coloring but was red in the final coloring (the structure of the algorithm assures us that vertices cannot change color more than once). Let A_e be the first event and C_e the second. Then

$$\Pr[A_e] = 2^{-n}(1-p)^n.$$

The first factor is the probability e is red in the first coloring, that all first coins of e came up heads. The second factor is the probability that all second coins came up tails. If they all did, then no $v \in e$ would be recolored in Step 2. Inversely, if any second coins of $v \in e$ came up heads there would be a first v (in the ordering) that came up heads. When it did v was still dangerous as e was still monochromatic and so v does look at its second coin and change its color. We have

$$2 \sum_{e \in H} \Pr[A_e] = k(1-p)^n$$

giving the first addend of our failure probability.

In Beck's 1978 proof, given in our first edition, there was no notion of "still dangerous" - every $d \in D$ changed its color if and only if its second coin was heads. The values $\Pr[A_e] = 2^{-n}(1-p)^n$ are the same in both arguments. Beck's had bounded $\Pr[C_e] \leq k^2pe^{pn}$. The new argument avoids excessive recoloration and leads to a better bound on $\Pr[C_e]$. We turn to the ingenious bounding of $\Pr[C_e]$.

For distinct $e, f \in E$ we say e blames f if:

- e, f overlap in precisely one element. Call it v .
- In the first coloring f was blue and in the final coloring e was red.
- In Step 2 v was the last vertex of e that changed color from blue to red.
- When v changed its color f was still entirely blue.

Suppose C_e holds. Some points of e changed color from blue to red so there is a last point v that did so. But why did v flip its coin? It must have been still dangerous. That is, v must be in some (perhaps many) set f that was blue in the first coloring and was still blue when v was considered. Can e, f overlap in another vertex v' ? No! For such a v' would necessarily have been blue in the first coloring (as $v' \in f$) and red in the final coloring (as $v' \in e$), but then v' changed color before v . Hence f was no longer entirely blue when v was considered,

contradicting the assumption on f . Therefore, when C_e holds, e blames some f . Let B_{ef} be the event that e blames f . Then $\sum_e \Pr[C_e] \leq \sum_{e \neq f} \Pr[B_{ef}]$. As there are less than $(2^{n-1}k)^2$ pairs $e \neq f$ it now suffices to bound $\Pr[B_{ef}] \leq 2^{1-2n}p$.

Let e, f with $e \cap f = \{v\}$ (otherwise B_{ef} cannot occur) be fixed. The random ordering of V induces a random ordering σ of $e \cup f$. Let $i = i(\sigma)$ denote the number of $v' \in e$ coming before v in the ordering and let $j = j(\sigma)$ denote the number of $v' \in f$ coming before v in the ordering. Fixing σ we claim

$$\Pr[B_{ef} \mid \sigma] \leq \frac{p}{2} 2^{-n+1} (1-p)^j 2^{-n+1+i} \left(\frac{1+p}{2} \right)^i$$

Let's take the factors one at a time. First, v itself must start blue and turn red. Second, all other $v' \in f$ must start blue. Third, all $v' \in f$ coming before v must have second coin tails. Fourth, all $v' \in e$ coming after v must start red (since v is the last point of e to change color). Finally, all $v' \in e$ coming before v must either start red or start blue and turn red. [The final factor may well be a substantial overestimate. Those $v' \in e$ coming before v which start blue must not only have second coin heads but must themselves lie in an $e' \in H$ monochromatic under the first coloring. Attempts to further improve bounds on $m(n)$ have often centered on this overestimate but (thus far!) to no avail.]

We can then write

$$\Pr[B_{ef}] \leq 2^{1-2n} p \mathbb{E}[(1+p)^i (1-p)^j],$$

where the expectation is over the uniform choice of σ . The following gem therefore completes the argument.

Lemma 3.5.3 $\mathbb{E}[(1+p)^i (1-p)^j] \leq 1$.

Proof. Fix a matching between $e - \{v\}$ and $f - \{v\}$; think of Mr. & Mrs. Jones, Mr. & Mrs. Smith, and so on. Condition on how many of each pair (two Joneses, one Smith, no Taylors, etc.) come before v . This splits the space into 3^{n-1} parts, and it suffices to show that the conditional expectation in each of them is at most 1. Indeed, the factor contributed to $(1+p)^i (1-p)^j$ from each pair is at most 1, as follows: when there is no Taylor there is no factor. When there are two Joneses there is a factor $(1+p)(1-p) \leq 1$. When there is one Smith the factor is equally likely to be $1+p$ (Brad) or $1-p$ (Angelina), giving a factor of one. Moreover, these factors are independent for different pairs (given the above conditioning). All factors are at most one, and hence so is their product. $\square$

The desired result follows. $\square$

17.3.6 3.6 Continuous Time

Discrete random processes can sometimes be analyzed by placing them in a continuous time framework. This allows the powerful methods of analysis (such as integration!) to be applied. The approach seems most effective when dealing with random orderings. We give two examples.

Property B. We modify the proof that $m(n) = \Omega(2^n n^{1/2} \ln^{-1/2} n)$ of the previous section. We assign to each vertex $v \in V$ a "birth time" x_v . The x_v are independent real variables, each uniform in $[0, 1]$. The ordering of V is then the ordering (under less than) of the x_v . We now claim

$$\Pr[B_{ef}] \leq \sum_{l=0}^{n-1} \binom{n-1}{l} 2^{1-2n} \int_0^1 x^l p^{l+1} (1-xp)^{n-1} dx.$$

For $T \subseteq e - \{v\}$ let B_{efT} be the event that B_{ef} and in the first coloring e had precisely $T \cup \{v\}$ blue. There are $\binom{n-1}{l}$ choices for an l -set T , with l ranging from 0 to $n-1$. The first

coloring on $e \cup f$ is then determined and has probability 2^{1-2n} of occurring. Suppose v has birth time $x_v = x$. All $w \in T \cup \{v\}$ must have second coin flip heads - probability p^{l+1} . All $w \in T$ must be born before v - so that $x_w < x$, which has probability x^l . No $w \in f - \{v\}$ can be born before v and have coin flip heads. Each such w has probability xp of doing that so there is probability $(1 - xp)^{n-1}$ that no w does. As $x_v = x$ was uniform in $[0, 1]$ we integrate over x . Recombining terms,

$$\Pr[B_{ef}] \leq 2^{1-2n} p \int_0^1 (1 + xp)^{n-1} (1 - xp)^{n-1} dx$$

The integrand is always at most one so $\Pr[B_{ef}] \leq 2^{1-2n} p$. The remainder of the proof is unchanged.

Random Greedy Packing. Let H be a $(k + 1)$ -uniform hypergraph on a vertex set V of size N . The $e \in H$, which we call edges, are simply subsets of V of size $k + 1$. We assume:

Degree Condition: Every $v \in V$ is in precisely D edges.

Codegree Condition: Every distinct pair $v, v' \in V$ have only $o(D)$ edges in common.

We think of k fixed ($k = 2$ being an illustrative example) and the asymptotics as $N, D \rightarrow \infty$, with no set relationship between N and D .

A packing is a family P of vertex disjoint edges $e \in H$. Clearly $|P| \leq N/(k + 1)$. We define a randomized algorithm to produce a (not necessarily optimal) packing. Assign to each $e \in H$ uniformly and independently a birth time $x_e \in [0, D)$. [The choice of $[0, D)$ rather than $[0, 1]$ proves to be a technical convenience. Note that as the x_e are real variables with probability one there are no ties.] At time zero $P \leftarrow \emptyset$. As time progresses from 0 to D when an edge e is born it is added to P if possible that is, unless there is already some $e' \in P$ that overlaps e . Let P_c denote the value of P just before time c - when all e with birth times $t_e < c$ have been examined. Set $P^{\text{FINAL}} = P_D$. Note that by time D all edges have been born and their births were in random order. Thus P^{FINAL} is identical to the discrete process - often called the random greedy algorithm - in which H is first randomly ordered and then the $e \in H$ are considered sequentially.

Theorem 3.6.1 [Spencer (1995)] The expected value of $|P^{\text{FINAL}}|$ is asymptotic to $N/(k + 1)$.

We say $v \in V$ survives at time c if no $e \in P_c$ contains v and we let S_c denote the set of $v \in V$ so surviving. Rather than looking at P^{FINAL} we shall examine P_c , where c is an arbitrary fixed nonnegative real. Let

$$f(c) = \lim^* \Pr[v \in S_c],$$

where, formally, we mean here that for all $\epsilon > 0$ there exist D_0, N_0 and $\delta > 0$ so that if H is $(k + 1)$ -uniform on $N > N_0$ vertices with each v in $D > D_0$ edges and every distinct pair $v, v' \in V$ has less than δD , common edges then $|f(c) - \Pr[v \in S_c]| < \epsilon$ for all $v \in V$.

The heart of the argument lies in showing that $f(c)$ exists by defining a continuous time birth process yielding that value. We now describe the birth process, omitting some of the epsilon-delta-manship needed to formally show the limit.

Our birth process starts at time c and time goes backwards to 0. It begins with root Eve, our anthropomorphized v . Eve has births in time interval $[0, c)$. The number of births is given by a Poisson distribution with mean c and given their number their times are uniformly and independently distributed. [This is a standard Poisson process with intensity one. Equivalently, on any infinitesimal time interval $[x, x + dx)$, Eve has probability dx of giving birth and these events are independent over disjoint intervals.] Our fertile Eve always gives birth to k -tuplets. Each child is born fertile under the same rules, so if Alice is born at time x she (in our unisexual model) has a Poisson distribution with mean x of births, uniformly distributed in $[0, x)$.

The resulting random tree $T = T_c$ can be shown to be finite (note the time interval is finite) with probability 1. Given a finite T we say for each vertex Alice that Alice survives or dies according to the following scheme.

Menendez Rule: If Alice has given birth to a set (or possibly several sets) of k -tuplets all of whom survived then she dies; otherwise she survives.

In particular, if Alice is childless she survives. We can then work our way up the tree to determine of each vertex whether she survives or dies.

Example. $c = 10, k = 2$. Eve gives birth to Alice, Barbara at time 8.3 and then to Rachel, Siena at time 4.3. Alice gives birth to Nancy, Olive at time 5.7 and Rachel gives birth to Linda, Mayavati at time 0.4. There are no other births. Leaves Nancy, Olive, Linda, Mayavati, Barbara and Siena then survive. Working up the tree Alice and Rachel die. In neither of Eve's births did both children survive and therefore Eve survives.

We define $f(c)$ to be the probability that the root Eve survives in the random birth tree $T = T_c$.

We outline the equivalence by defining a tree $T = T_c(v)$ for $v \in H$. For each edge e containing v with birth time $t = t_e < c$ we say that $e - \{v\}$ is a set of k -tuplets born to v at time t . We work recursively; if w is born at time t then for each e' containing w with birth time $t' = t_{e'} < t$ we say that $e' - \{w\}$ is a set of k -tuplets born to w at time t' . Possibly this process does not give a tree since the same vertex w may be reached in more than one way - the simplest example is if $v \in e, e'$ where both have birth times less than c and e, e' share another common vertex w . Then the process is stillborn and $T_c(v)$ is not defined. We'll argue that for any particular tree T ,

$$\lim^* \Pr [T_c(v) \cong T] = \Pr [T_c = T] \quad (3.1)$$

As $\sum_T \Pr [T_c = T] = 1$ this gives a rather roundabout argument that the process defining $T_c(v)$ is almost never stillborn.

We find $T_c(v)$ in stages. First consider the D edges e containing v . The number of them with birth time $t_e < c$ has binomial distribution $\text{BIN}[D, c/D]$ which approaches (critically) the Poisson distribution with mean c . Given that there are l such e their birth times t_e are uniformly distributed. There are (by the codegree condition) $o(D^2)$ pairs e, e' containing v and also some other vertex so there is probability $o(1)$ that two such e, e' have birth time less than c . Now suppose $T_c(v)$ has been built out to a certain level and a vertex w has been born at time t . There are only $o(D)$ common edges between w and any of the finite number of w' already born, so there are still about D edges e containing w and no other such w' . We now examine their birth times, the number with $t_e < x$ has binomial distribution $\text{BIN}[D - o(D), x/D]$ which approaches the Poisson distribution with mean x . As above, almost surely no two such e, e' will have a common vertex other than w itself. For any fixed T the calculation of $\Pr [T_c(v) \cong T]$ involves a finite number of these limits, which allows us to conclude (3.1).

With $c < d$ the random tree T_d includes T_c as a subtree by considering only those births of Eve occurring in $[0, c)$. If Eve survives in T_d she must survive in T_c . Hence $f(d) \leq f(c)$. We now claim

$$\lim_{c \rightarrow \infty} f(c) = 0$$

If not, the nondecreasing f would have a limit $L > 0$ and all $f(x) \geq L$. Suppose in T_c Eve had i births. In each birth there would be probability at least L^k that all k children survived. The probability that Eve survived would then be at most $(1 - L^k)^i$. Since the number of Eve's births is Poisson with mean c ,

$$f(c) \leq \sum_{i=0}^{\infty} e^{-c} \frac{c^i}{i!} (1 - L^k)^i = e^{-L^k c}$$

but then $\lim_{c \rightarrow \infty} f(c) = 0$, a contradiction.

By linearity of expectation $E[|S_c|] \rightarrow f(c)n$. As $(k+1)|P_c| + |S_c| = n$, $E[|P_c|] \rightarrow (1 - f(c))n/(k+1)$. But $E[|P^{\text{FINAL}}|] \geq E[|P_c|]$. We make $f(c)$ arbitrarily small by taking c appropriately big, so that $E[|P^{\text{FINAL}}|] \geq (1 - o(1))n/(k+1)$. As $|P^{\text{FINAL}}| \leq n/(k+1)$ always, the theorem follows.

Remark. We can actually say more about $f(c)$. For Δc small, $f(c + \Delta c) - f(c) \sim -(\Delta c)f(c)^{k+1}$ as, roughly, an Eve starting at time $c + \Delta c$ might have a birth in time interval $[c, c + \Delta c)$, all of whose children survive, while Eve has no births in $[0, c)$, all of whose children survive. Letting $\Delta c \rightarrow 0$ yields the differential equation $f'(c) = -f(c)^{k+1}$. The initial value $f(0) = 1$ gives a unique solution $f(c) = (1 + ck)^{-1/k}$. It is intriguing to plug in $c = D$. This is not justified as our limit arguments were for c fixed and $N, D \rightarrow \infty$. Nonetheless, that would yield $E[|S_D|] = O(ND^{-1/k})$, that the random greedy algorithm would leave $O(ND^{-1/k})$ vertices uncovered. Suppose we replace the codegree condition by the stronger condition that every distinct pair $v, v' \in V$ have at most one edge in common. There is computer simulation data that in those cases the random greedy algorithm does leave $O(ND^{-1/k})$ vertices uncovered. This remains an open question, though it is shown in Alon, Kim and Spencer (1997) that this is the case for a modified version of the greedy algorithm.

Corollary 3.6.2 Under the assumptions of the theorem there exists a packing P of size $\sim N/(k+1)$.

Proof. We have defined a random process that gives a packing with expected size $\sim N/(k+1)$ and our usual magic implies such a P must exist.

In particular, this gives an alternate proof to the Erdős-Hanani conjecture, first proved by Rödl as given in Section 4.7. We use the notation of that section and define the packing number $m(n, k, l)$ as the maximal size of a family F of k -element subsets of $[n] = \{1, \dots, n\}$ such that no l -set is contained in more than one k -set. Define a hypergraph $H = H(n, k, l)$ as follows: The vertices of H are the l -element subsets of $[n]$. For each k -element $A \subset [n]$ we define an edge e_A as the set of l -element subsets of A . A family F satisfying the above conditions then corresponds to a packing $P = \{e_A : A \in F\}$ in H . H has $N = \binom{n}{l}$ vertices. Each edge e_A has size $K + 1 = \binom{k}{l}$. Each vertex is in $D = \binom{n-l}{k-l}$ edges. The number of edges containing two vertices v, v' depends on their intersection. It is largest (given $v \neq v'$) when v, v' (considered as l -sets) overlap in $l - 1$ points and then it is $\binom{n-l-1}{k-l-1}$. We assume (as in Section 4.7) that k, l are fixed and $n \rightarrow \infty$ so this number of common edges is $o(D)$. The assumptions of Section 4.7 give $K + 1$ fixed, $N, D \rightarrow \infty$ so that there exists P with

$$m(n, k, l) = |P| \sim N/(K + 1) \sim \binom{n}{l} / \binom{k}{l}.$$

17.3.7 3.7 Exercises

1. As shown in Section 3.1, the Ramsey number $R(k, k)$ satisfies

$$R(k, k) > n - \binom{n}{k} 2^{1 - \binom{k}{2}}$$

for every integer n . Conclude that

$$R(k, k) \geq (1 - o(1)) \frac{k}{e} 2^{k/2}$$

2. Prove that the Ramsey number $R(4, k)$ satisfies

$$R(4, k) \geq \Omega((k/\ln k)^2)$$

3. Prove that every three-uniform hypergraph with n vertices and $m \geq n/3$ edges contains an independent set (i.e., a set of vertices containing no edges) of size at least

$$\frac{2n^{3/2}}{3\sqrt{3}\sqrt{m}}$$

4. (*) Show that there is a finite n_0 such that any directed graph on $n > n_0$ vertices in which each outdegree is at least $\log_2 n - \frac{1}{10} \log_2 \log_2 n$ contains an even simple directed cycle.

The Probabilistic Lens: High Girth and High Chromatic Number

Many consider this one of the most pleasing uses of the probabilistic method, as the result is surprising and does not appear to call for nonconstructive techniques. The girth of a graph G is the size of its shortest cycle, $\alpha(G)$ is the size of the largest independent set in G and $\chi(G)$ denotes its chromatic number.

Theorem 1 [Erdős (1959)] For all k, l there exists a graph G with girth $(G) > l$ and $\chi(G) > k$.

Proof. Fix $\theta < 1/l$ and let $G \sim G(n, p)$ with $p = n^{\theta-1}$; that is, G is a random graph on n vertices chosen by picking each pair of vertices as an edge randomly and independently with probability p . Let X be the number of cycles of size at most l . Then

$$\mathbb{E}[X] = \sum_{i=3}^l \frac{\binom{n}{i}}{2i} p^i \leq \sum_{i=3}^l \frac{n^{\theta i}}{2i} = o(n)$$

as $\theta l < 1$. In particular,

$$\Pr[X \geq n/2] = o(1).$$

Set $x = \lceil (3/p) \ln n \rceil$ so that

$$\Pr[\alpha(G) \geq x] \leq \binom{n}{x} (1-p)^{\binom{x}{2}} < [ne^{-p(x-1)/2}]^x = o(1).$$

Let n be sufficiently large so that both these events have probability less than 0.5. Then there is a specific G with less than $n/2$ cycles of length at most l and with $\alpha(G) < 3n^{1-\theta} \ln n$. Remove from G a vertex from each cycle of length at most l . This gives a graph G^* with at least $n/2$ vertices. G^* has girth greater than l and $\alpha(G^*) \leq \alpha(G)$. Thus

$$\chi(G^*) \geq \frac{|G^*|}{\alpha(G^*)} \geq \frac{n/2}{3n^{1-\theta} \ln n} = \frac{n^\theta}{6 \ln n}.$$

To complete the proof, let n be sufficiently large so that this is greater than k . $\square$

17.4 4 The Second Moment

You don't have to believe in God but you should believe in The Book.

- Paul Erdős

17.4.1 4.1 Basics

After the expectation the most vital statistic for a random variable X is the variance. We denote it $\text{Var}[X]$. It is defined by

$$\text{Var}[X] = \text{E}[(X - \text{E}[X])^2]$$

and measures how spread out X is from its expectation. We shall generally, following standard practice, let μ denote expectation and σ^2 denote variance. The positive square root σ of the variance is called the standard deviation. With this notation, here is our basic tool.

Theorem 4.1.1 [Chebyshev's Inequality] For any positive λ ,

$$\Pr[|X - \mu| \geq \lambda\sigma] \leq \frac{1}{\lambda^2}.$$

Proof. $\sigma^2 = \text{Var}[X] = \text{E}[(X - \mu)^2] \geq \lambda^2\sigma^2 \Pr[|X - \mu| \geq \lambda\sigma]$. $\square$

The use of Chebyshev's Inequality is called the second moment method.

Chebyshev's Inequality is most possible when no additional restrictions are placed on X as X may be $\mu + \lambda\sigma$ and $\mu - \lambda\sigma$ with probability $1/2\lambda^2$ and otherwise μ . Note, however, that when X is a normal distribution with mean μ and standard deviation σ then

$$\Pr[|X - \mu| \geq \lambda\sigma] = 2 \int_{\lambda}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt$$

and for λ large this quantity is asymptotically $\sqrt{2/\pi} e^{-\lambda^2/2}/\lambda$, which is significantly smaller than $1/\lambda^2$. In Chapters 7 and 8 we shall see examples where X is the sum of "nearly independent" random variables and these better bounds can apply.

Suppose we have a decomposition

$$X = X_1 + \cdots + X_m.$$

Then $\text{Var}[X]$ may be computed by the formula

$$\text{Var}[X] = \sum_{i=1}^m \text{Var}[X_i] + \sum_{i \neq j} \text{Cov}[X_i, X_j]$$

Here the second sum is over ordered pairs and the covariance $\text{Cov}[Y, Z]$ is defined by

$$\text{Cov}[Y, Z] = \text{E}[YZ] - \text{E}[Y]\text{E}[Z]$$

In general, if Y, Z are independent then $\text{Cov}[Y, Z] = 0$. This often simplifies variance calculations considerably. Now suppose further, as will generally be the case in our applications, that the X_i are indicator random variables; that is, $X_i = 1$ if a certain event A_i holds and otherwise $X_i = 0$. If X_i is one with probability $p_i = \Pr[A_i]$ then

$$\text{Var}[X_i] = p_i(1 - p_i) \leq p_i = \text{E}[X_i],$$

and so

$$\text{Var}[X] \leq \text{E}[X] + \sum_{i \neq j} \text{Cov}[X_i, X_j].$$

17.4.2 4.2 Number Theory

The second moment method is an effective tool in number theory. Let $\nu(n)$ denote the number of primes p dividing n . (We do not count multiplicity though it would make little difference.) The following result says, roughly, that "almost all" n have "very close to" $\ln \ln n$ prime factors. This was first shown by Hardy and Ramanujan in 1920 by a quite complicated argument. We give a remarkably simple proof of Turán (1934), a proof that played a key role in the development of probabilistic methods in number theory.

Theorem 4.2.1 Let $\omega(n) \rightarrow \infty$ arbitrarily slowly. Then the number of x in $\{1, \dots, n\}$ such that

$$|\nu(x) - \ln \ln n| > \omega(n) \sqrt{\ln \ln n}$$

is $o(n)$.

Proof. Let x be randomly chosen from $\{1, \dots, n\}$. For p prime set

$$X_p = \begin{cases} 1 & \text{if } p \mid x \\ 0 & \text{otherwise} \end{cases}$$

Set $M = n^{1/10}$ and set $X = \sum X_p$, the summation over all primes $p \leq M$. As no $x \leq n$ can have more than ten prime factors larger than M we have $\nu(x) - 10 \leq X(x) \leq \nu(x)$ so that large deviation bounds on X will translate into asymptotically similar bounds for ν . [Here 10 could be any (large) constant.] Now

$$\mathbb{E}[X_p] = \frac{\lfloor n/p \rfloor}{n}$$

As $y - 1 < \lfloor y \rfloor \leq y$,

$$\mathbb{E}[X_p] = 1/p + O(1/n)$$

By linearity of expectation,

$$\mathbb{E}[X] = \sum_{p \leq M} \left(\frac{1}{p} + O\left(\frac{1}{n}\right) \right) = \ln \ln n + O(1)$$

where here we used the well-known fact that $\sum_{p \leq x} (1/p) = \ln \ln x + O(1)$, which can be proved by combining Stirling's formula with Abel summation.

Now we find an asymptotic expression for

$$\text{Var}[X] = \sum_{p \leq M} \text{Var}[X_p] + \sum_{p \neq q} \text{Cov}[X_p, X_q].$$

As $\text{Var}[X_p] = (1/p)(1 - 1/p) + O(1/n)$,

$$\sum_{p \leq M} \text{Var}[X_p] = \left(\sum_{p \leq M} \frac{1}{p} \right) + O(1) = \ln \ln n + O(1).$$

With p, q distinct primes, $X_p X_q = 1$ if and only if $p \mid x$ and $q \mid x$, which occurs if and only if $pq \mid x$. Hence

$$\begin{aligned}
\text{Cov}[X_p, X_q] &= E[X_p X_q] - E[X_p] E[X_q] \\
&= \frac{\lfloor n/pq \rfloor}{n} - \frac{\lfloor n/p \rfloor}{n} \frac{\lfloor n/q \rfloor}{n} \\
&\leq \frac{1}{pq} - \left(\frac{1}{p} - \frac{1}{n} \right) \left(\frac{1}{q} - \frac{1}{n} \right) \\
&\leq \frac{1}{n} \left(\frac{1}{p} + \frac{1}{q} \right).
\end{aligned}$$

Thus

$$\sum_{p \neq q} \text{Cov}[X_p, X_q] \leq \frac{1}{n} \sum_{p \neq q} \left(\frac{1}{p} + \frac{1}{q} \right) \leq \frac{2M}{n} \sum \frac{1}{p}.$$

Thus

$$\sum_{p \neq q} \text{Cov}[X_p, X_q] \leq O(n^{-9/10} \ln \ln n) = o(1)$$

and similarly

$$\sum_{p \neq q} \text{Cov}[X_p, X_q] \geq -o(1)$$

That is, the covariances do not affect the variance, $\text{Var}[X] = \ln \ln n + O(1)$ and Chebyshev's Inequality actually gives

$$\Pr[|X - \ln \ln n| > \lambda \sqrt{\ln \ln n}] < \lambda^{-2} + o(1)$$

for any constant $\lambda > 0$. As $|X - \nu| \leq 10$ the same holds for ν .

In a classic paper Erdős and Kac (1940) showed, essentially, that ν does behave like a normal distribution with mean and variance $\ln \ln n$. Here is their precise result.

Theorem 4.2.2 Let λ be fixed, positive, negative or zero. Then

$$\lim_{n \rightarrow \infty} \frac{1}{n} |\{x : 1 \leq x \leq n, \nu(x) \geq \ln \ln n + \lambda \sqrt{\ln \ln n}\}| = \int_{\lambda}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt$$

Proof. We outline the argument, emphasizing the similarities to Turán's proof. Fix a function $s(n)$ with $s(n) \rightarrow \infty$ and $s(n) = o((\ln \ln n)^{1/2})$ - for example, $s(n) = \ln \ln \ln n$. Set $M = n^{1/s(n)}$. Set $X = \sum X_p$, the summation over all primes $p \leq M$. As no $x \leq n$ can have more than $s(n)$ prime factors greater than M we have $\nu(x) - s(n) \leq X(x) \leq \nu(x)$ so that it suffices to show Theorem 4.2.2 with ν replaced by X . Let Y_p be independent random variables with $\Pr[Y_p = 1] = 1/p$, $\Pr[Y_p = 0] = 1 - 1/p$ and set $Y = \sum Y_p$, the summation over all primes $p \leq M$. This Y represents an idealized version of X . Set

$$\mu = E[Y] = \sum_{p \leq M} \frac{1}{p} = \ln \ln n + o((\ln \ln n)^{1/2})$$

and

$$\sigma^2 = \text{Var}[Y] = \sum_{p \leq M} \frac{1}{p} \left(1 - \frac{1}{p} \right) \sim \ln \ln n$$

and define the normalized $\tilde{Y} = (Y - \mu)/\sigma$. From the Central Limit Theorem $\tilde{Y}$ approaches the standard normal N and $E[\tilde{Y}^k] \rightarrow E[N^k]$ for every positive integer k . Set $\tilde{X} = (X - \mu)/\sigma$. We compare $\tilde{X}, \tilde{Y}$.

For any distinct primes $p_1, \dots, p_s \leq M$,

$$E[X_{p_1} \cdots X_{p_s}] - E[Y_{p_1} \cdots Y_{p_s}] = \frac{1}{n} \left\lfloor \frac{n}{p_1 \cdots p_s} \right\rfloor - \frac{1}{p_1 \cdots p_s} = O\left(\frac{1}{n}\right)$$

We let k be an arbitrary fixed positive integer and compare $E[\tilde{X}^k]$ and $E[\tilde{Y}^k]$. Expanding, $\tilde{X}^k$ is a polynomial in X with coefficients $n^{o(1)}$. Further expanding each $X^j = (\sum X_p)^j$ - always reducing X_p^a to X_p when $a \geq 2$ - gives the sum of $O(M^k) = n^{o(1)}$ terms of the form $X_{p_1} \cdots X_{p_s}$. The same expansion applies to $\tilde{Y}$. As the corresponding terms have expectations within $O(1/n)$ the total difference

$$E[\tilde{X}^k] - E[\tilde{Y}^k] = n^{-1+o(1)} = o(1)$$

Hence each moment of $\tilde{X}$ approaches that of the standard normal N . A standard, though nontrivial, theorem in probability theory gives that $\tilde{X}$ must therefore approach N in distribution.

We recall the famous quotation of G. H. Hardy:

317 is a prime, not because we think so, or because our minds are shaped in one way rather than another, but because it is so, because mathematical reality is built that way.

How ironic - though not contradictory - that the methods of probability theory can lead to a greater understanding of the prime factorization of integers. Additional results applying information about the moments of a distribution in order to determine it appear in Chapter 8; see also Billingsley (1995).

17.4.3 4.3 More Basics

Let X be a nonnegative integral valued random variable and suppose we want to bound $\Pr[X = 0]$ given the value $\mu = E[X]$. If $\mu < 1$ we may use the inequality

$$\Pr[X > 0] \leq E[X]$$

so that if $E[X] \rightarrow 0$ then $X = 0$ almost always. (Here we are imagining an infinite sequence of X dependent on some parameter n going to infinity.) But now suppose $E[X] \rightarrow \infty$. It does not necessarily follow that $X > 0$ almost always. For example, let X be the number of deaths due to nuclear war in the twelve months after reading this paragraph. Calculation of $E[X]$ can make for lively debate but few would deny that it is quite large. Yet we may believe - or hope - that $\Pr[X \neq 0]$ is very close to zero. We can sometimes deduce $X > 0$ almost always if we have further information about $\text{Var}[X]$.

Theorem 4.3.1 $\Pr[X = 0] \leq \frac{\text{Var}[X]}{E[X]^2}$.

Proof. Set $\lambda = \mu/\sigma$ in Chebyshev's Inequality. Then

$$\Pr[X = 0] \leq \Pr[|X - \mu| \geq \lambda\sigma] \leq \frac{1}{\lambda^2} = \frac{\sigma^2}{\mu^2}$$

We generally apply this result in asymptotic terms.

Corollary 4.3.2 If $\text{Var}[X] = o(E[X]^2)$ then $X > 0$ almost always.

The proof of Theorem 4.3.1 actually gives that, for any $\epsilon > 0$,

$$\Pr[|X - E[X]| \geq \epsilon E[X]] \leq \frac{\text{Var}[X]}{\epsilon^2 E[X]^2}$$

and thus in asymptotic terms we actually have the following stronger assertion.

Corollary 4.3.3 If $\text{Var}[X] = o(E[X]^2)$ then $X \sim E[X]$ almost always.

Suppose again $X = X_1 + \cdots + X_m$, where X_i is the indicator random variable for event A_i . For indices i, j write $i \sim j$ if $i \neq j$ and the events A_i, A_j are not independent. We set (the sum is over ordered pairs)

$$\Delta = \sum_{i \sim j} \Pr[A_i \wedge A_j]$$

Note that when $i \sim j$,

$$\text{Cov}[X_i, X_j] = E[X_i X_j] - E[X_i] E[X_j] \leq E[X_i X_j] = \Pr[A_i \wedge A_j]$$

and that when $i \neq j$ and not $i \sim j$ then $\text{Cov}[X_i, X_j] = 0$. Thus

$$\text{Var}[X] \leq E[X] + \Delta.$$

Corollary 4.3.4 If $E[X] \rightarrow \infty$ and $\Delta = o(E[X]^2)$ then $X > 0$ almost always. Furthermore $X \sim E[X]$ almost always.

Let us say $X_1, \dots, X_m$ are symmetric if for every $i \neq j$ there is a measure preserving mapping of the underlying probability space that sends event A_i to event A_j . Examples will appear in the next section. In this instance we write

$$\Delta = \sum_{i \sim j} \Pr[A_i \wedge A_j] = \sum_i \Pr[A_i] \sum_{j \sim i} \Pr[A_j | A_i]$$

and note that the inner summation is independent of i . We set

$$\Delta^* = \sum_{j \sim i} \Pr[A_j | A_i]$$

where i is any fixed index. Then

$$\Delta = \sum_i \Pr[A_i] \Delta^* = \Delta^* \sum_i \Pr[A_i] = \Delta^* E[X].$$

Corollary 4.3.5 If $E[X] \rightarrow \infty$ and $\Delta^* = o(E[X])$ then $X > 0$ almost always. Furthermore $X \sim E[X]$ almost always.

The condition of Corollary 4.3.5 has the intuitive sense that conditioning on any specific A_i holding does not substantially increase the expected number $E[X]$ of events holding.

17.4.4 4.4 Random Graphs

The random graph $G(n, p)$ is, informally, the graph on n labeled vertices, obtained by selecting each pair of vertices to be an edge, randomly and independently, with probability p . A property of graphs is a family of graphs closed under isomorphism. A function $r(n)$ is a threshold function for some property P , if whenever $p = p(n) \ll r(n)$ then $G(n, p)$ does not satisfy P almost always, and whenever $p \gg r(n)$ then $G(n, p)$ satisfies P almost always. For

more precise definitions of the random graph $G(n, p)$ and of threshold functions, see Section 10.1.

The results of this section are generally surpassed by those of Chapter 10 but they were historically the first results and provide a good illustration of the second moment. We begin with a particular example. By $\omega(G)$ we denote here and in the rest of the book the number of vertices in the maximum clique of the graph G .

Theorem 4.4.1 The property $\omega(G) \geq 4$ has threshold function $n^{-2/3}$.

Proof. For every 4-set S of vertices in $G(n, p)$ let A_S be the event " S is a clique" and X_S its indicator random variable. Then

$$E[X_S] = \Pr[A_S] = p^6$$

as six different edges must all lie in $G(n, p)$. Set

$$X = \sum_{|S|=4} X_S$$

so that X is the number of 4-cliques in G and $\omega(G) \geq 4$ if and only if $X > 0$. Linearity of expectation gives

$$E[X] = \sum_{|S|=4} E[X_S] = \binom{n}{4} p^6 \sim \frac{n^4 p^6}{24}.$$

When $p(n) \ll n^{-2/3}$, $E[X] = o(1)$ and so $X = 0$ almost surely.

Now suppose $p(n) \gg n^{-2/3}$ so that $E[X] \rightarrow \infty$ and consider the Δ^* of Corollary 4.3.5. (All 4-sets "look the same" so that the X_S are symmetric.) Here $S \sim T$ if and only if $S \neq T$ and S, T have common edges; that is, if and only if $|S \cap T| = 2$ or 3 . Fix S . There are $O(n^2)$ sets T with $|S \cap T| = 2$ and for each of these $\Pr[A_T | A_S] = p^5$. There are $O(n)$ sets T with $|S \cap T| = 3$ and for each of these $\Pr[A_T | A_S] = p^3$. Thus


$$\Delta^* = O(n^2 p^5) + O(np^3) = o(n^4 p^6) = o(E[X])$$

since $p \gg n^{-2/3}$. Corollary 4.3.5 therefore applies and $X > 0$; that is, there does exist a clique of size 4, almost always. $\square$


The proof of Theorem 4.4.1 appears to require a fortuitous calculation of Δ^* . The following definitions pave the way for the more general Theorem 4.4.2.

Definition 1 Let H be a graph with v vertices and e edges. We call $\rho(H) = e/v$ the density of H . We call H balanced if every subgraph H' has $\rho(H') \leq \rho(H)$. We call H strictly balanced if every proper subgraph H' has $\rho(H') < \rho(H)$.

Examples. K_4 and, in general, K_k are strictly balanced. The graph



is not balanced as it has density $7/5$ while the subgraph K_4 has density $3/2$. The graph



is balanced but not strictly balanced as it and its subgraph K_4 have density $3/2$.

Theorem 4.4.2 Let H be a balanced graph with v vertices and e edges. Let $A(G)$ be the event that H is a subgraph (not necessarily induced) of G . Then $p = n^{-v/e}$ is the threshold function for A .

Proof. We follow the argument of Theorem 4.4.1. For each v -set S let A_S be the event that $G|_S$ contains H as a subgraph. Then

$$p^e \leq \Pr[A_S] \leq v!p^e$$

(Any particular placement of H has probability p^e of occurring and there are at most $v!$ possible placements. The precise calculation of $\Pr[A_S]$ is, in general, complicated due to the overlapping of potential copies of H .) Let X_S be the indicator random variable for A_S and

$$X = \sum_{|S|=v} X_S$$

so that A holds if and only if $X > 0$. Linearity of expectation gives

$$\mathbb{E}[X] = \sum_{|S|=v} \mathbb{E}[X_S] = \binom{n}{v} \Pr[A_S] = \Theta(n^v p^e).$$

If $p \ll n^{-v/e}$ then $\mathbb{E}[X] = o(1)$, so $X = 0$ almost always.

Now assume $p \gg n^{-v/e}$ so that $\mathbb{E}[X] \rightarrow \infty$ and consider the Δ^* of Corollary 4.3.5. (All v -sets look the same so the X_S are symmetric.) Here $S \sim T$ if and only if $S \neq T$ and S, T have common edges; that is, if and only if $|S \cap T| = i$ with $2 \leq i \leq v-1$. Let S be fixed. We split

$$\Delta^* = \sum_{T \sim S} \Pr[A_T | A_S] = \sum_{i=2}^{v-1} \sum_{|T \cap S|=i} \Pr[A_T | A_S].$$

For each i there are $O(n^{v-i})$ choices of T . Fix S, T and consider $\Pr[A_T | A_S]$. There are $O(1)$ possible copies of H on T . Each has - since, critically, H is balanced - at most ie/v edges with both vertices in S and thus at least $e - (ie/v)$ other edges. Hence

$$\Pr[A_T | A_S] = O(p^{e-(ie/v)})$$

and

$$\begin{aligned} \Delta^* &= \sum_{i=2}^{v-1} O(n^{v-i} p^{e-(ie/v)}) \\ &= \sum_{i=2}^{v-1} O((n^v p^e)^{1-i/v}) \\ &= \sum_{i=2}^{v-1} o(n^v p^e) \\ &= o(\mathbb{E}[X]) \end{aligned}$$

since $n^v p^e \rightarrow \infty$. Hence Corollary 4.3.5 applies. $\square$

Theorem 4.4.3 In the notation of Theorem 4.4.2 if H is not balanced then $p = n^{-v/e}$ is not the threshold function for A .

Proof. Let H_1 be a subgraph of H with v_1 vertices, e_1 edges and $e_1/v_1 > e/v$. Let α satisfy $v_1/e_1 < \alpha < v/e$ and set $p = n^{-\alpha}$. The expected number of copies of H_1 is then $o(1)$ so almost always $G(n, p)$ contains no copy of H_1 . But if it contains no copy of H_1 then it surely can contain no copy of H .

The threshold function for the property of containing a copy of H , for general H , was examined in the original papers of Erdős and Rényi (1960). It still provides an excellent introduction to the theory of random graphs. Let H_1 be that subgraph with maximal density $\rho(H_1) = e_1/v_1$. (When H is balanced we may take $H_1 = H$.) They showed that $p = n^{-v_1/e_1}$ is the threshold function. We do not show this here though it follows fairly straightforwardly from these methods.

We finish this section with two strengthenings of Theorem 4.4.2.

Theorem 4.4.4 Let H be strictly balanced with v vertices, e edges and a automorphisms. Let X be the number of copies of H in $G(n, p)$. Assume $p \gg n^{-v/e}$. Then almost always

$$X \sim \frac{n^v p^e}{a}$$

Proof. Label the vertices of H by $1, \dots, v$. For each ordered $x_1, \dots, x_v$ let $A_{x_1, \dots, x_v}$ be the event that $x_1, \dots, x_v$ provides a copy of H in that order. Specifically we define

$$A_{x_1, \dots, x_v} : \{i, j\} \in E(H) \Rightarrow \{x_i, x_j\} \in E(G).$$

We let $I_{x_1, \dots, x_v}$ be the corresponding indicator random variable. We define an equivalence class on v -tuples by setting $(x_1, \dots, x_v) \equiv (y_1, \dots, y_v)$ if there is an automorphism σ of $V(H)$ so that $y_{\sigma(i)} = x_i$ for $1 \leq i \leq v$. Then

$$X = \sum I_{x_1, \dots, x_v}$$

gives the number of copies of H in G where the sum is taken over one entry from each equivalence class. As there are $(n)_v/a$ terms,

$$E[X] = \frac{(n)_v}{a} E[I_{x_1, \dots, x_v}] = \frac{(n)_v p^e}{a} \sim \frac{n^v p^e}{a}$$

Our assumption $p \gg n^{-v/e}$ implies $E[X] \rightarrow \infty$. It suffices therefore to show $\Delta^* = o(E[X])$. Fixing $x_1, \dots, x_v$,

$$\Delta^* = \sum_{(y_1, \dots, y_v) \sim (x_1, \dots, x_v)} \Pr[A_{(y_1, \dots, y_v)} \mid A_{(x_1, \dots, x_v)}]$$

There are $v!/a = O(1)$ terms with $\{y_1, \dots, y_v\} = \{x_1, \dots, x_v\}$ and for each the conditional probability is at most 1 (actually, at most p), thus contributing $O(1) = o(E[X])$ to Δ^* . When $\{y_1, \dots, y_v\} \cap \{x_1, \dots, x_v\}$ has i elements, $2 \leq i \leq v-1$ the argument of Theorem 4.4.2 gives that the contribution to Δ^* is $o(E[X])$. Altogether $\Delta^* = o(E[X])$ and we apply Corollary 4.3.5.

Theorem 4.4.5 Let H be any fixed graph. For every subgraph H' of H (including H itself) let $X_{H'}$ denote the number of copies of H' in $G(n, p)$. Assume p is such that $E[X_{H'}] \rightarrow \infty$ for every H' . Then

$$X_H \sim E[X_H]$$

almost always.

Proof. Let H have v vertices and e edges. As in Theorem 4.4.4 it suffices to show $\Delta^* = o(\mathbb{E}[X])$. We split Δ^* into a finite number of terms. For each H' with w vertices and f edges we have those $(y_1, \dots, y_v)$ that overlap with the fixed $(x_1, \dots, x_v)$ in a copy of H' . These terms contribute, up to constants,

$$n^{v-w} p^{e-f} = \Theta \left(\frac{\mathbb{E}[X_H]}{\mathbb{E}[X_{H'}]} \right) = o(\mathbb{E}[X_H])$$

to Δ^* . Hence Corollary 4.3.5 does apply.

17.4.5 4.5 Clique Number

Now we fix edge probability $p = \frac{1}{2}$ and consider the clique number $\omega(G)$. We set

$$f(k) = \binom{n}{k} 2^{-\binom{k}{2}},$$

the expected number of k -cliques. The function $f(k)$ drops under one at $k \sim 2 \log_2 n$. [Very roughly, $f(k)$ is like $n^k 2^{-k^2/2}$.]

Theorem 4.5.1 Let $k = k(n)$ satisfy $k \sim 2 \log_2 n$ and $f(k) \rightarrow \infty$. Then almost always $\omega(G) \geq k$.

Proof. For each k -set S let A_S be the event " S is a clique" and X_S the corresponding indicator random variable. We set

$$X = \sum_{|S|=k} X_S$$

so that $\omega(G) \geq k$ if and only if $X > 0$. Then $\mathbb{E}[X] = f(k) \rightarrow \infty$ and we examine the Δ^* of Corollary 4.3.5. Fix S and note that $T \sim S$ if and only if $|T \cap S| = i$, where $2 \leq i \leq k-1$. Hence

$$\Delta^* = \sum_{i=2}^{k-1} \binom{k}{i} \binom{n-k}{k-i} 2^{\binom{i}{2} - \binom{k}{2}}$$

and so

$$\frac{\Delta^*}{\mathbb{E}[X]} = \sum_{i=2}^{k-1} g(i)$$

where we set

$$g(i) = \frac{\binom{k}{i} \binom{n-k}{k-i}}{\binom{n}{k}} 2^{\binom{i}{2}}$$

Observe that $g(i)$ may be thought of as the probability that a randomly chosen T will intersect a fixed S in i points times the factor increase in $\Pr[A_T]$ when it does. Setting $i = 2$,

$$g(2) = 2 \frac{\binom{k}{2} \binom{n-k}{k-2}}{\binom{n}{k}} \sim \frac{k^4}{n^2} \leq o(1/n)$$

At the other extreme $i = k-1$,

$$g(k-1) = \frac{k(n-k) 2^{-\binom{k-1}{2}}}{\binom{n}{k} 2^{-\binom{k}{2}}} \sim \frac{2kn 2^{-k}}{\mathbb{E}[X]}.$$

As $k \sim 2 \log_2 n$, the numerator is $n^{-1+o(1)}$. The denominator approaches infinity and so $g(k-1) \leq o(1/n)$. Some detailed calculation (which we omit) gives that the remaining $g(i)$ and their sum are also negligible so that Corollary 4.3.5 applies.

Theorem 4.5.1 leads to a strong concentration result for $\omega(G)$. For $k \sim 2 \log_2 n$,

$$\frac{f(k+1)}{f(k)} = \frac{n-k}{k+1} 2^{-k} = n^{-1+o(1)} = o(1).$$

Let $k_0 = k_0(n)$ be that value with $f(k_0) \geq 1 > f(k_0+1)$. For "most" n the function $f(k)$ will jump from a large $f(k_0)$ to a small $f(k_0+1)$. The probability that G contains a clique of size k_0+1 is at most $f(k_0+1)$, which will be very small. When $f(k_0)$ is large, Theorem 4.5.1 implies that G contains a clique of size k_0 with probability nearly 1. Together, with very high probability $\omega(G) = k_0$. For some n one of the values $f(k_0), f(k_0+1)$ may be of moderate size so this argument does not apply. Still one may show a strong concentration result found independently by Bollobás and Erdős (1976) and Matula (1976).

Corollary 4.5.2 There exists $k = k(n)$ so that

$$\Pr[\omega(G) = k \text{ or } k+1] \rightarrow 1.$$

We give yet stronger results on the distribution of $\omega(G)$ in Section 10.2.

17.4.6 4.6 Distinct Sums

A set $x_1, \dots, x_k$ of positive integers is said to have distinct sums if all sums

$$\sum_{i \in S} x_i, \quad S \subseteq \{1, \dots, k\}$$

are distinct. Let $f(n)$ denote the maximal k for which there exists a set

$$\{x_1, \dots, x_k\} \subset \{1, \dots, n\}$$

with distinct sums. The simplest example of a set with distinct sums is $\{2^i : i \leq \log_2 n\}$. This example shows

$$f(n) \geq 1 + \lfloor \log_2 n \rfloor$$

Erdős offered \$300 for a proof or disproof that

$$f(n) \leq \log_2 n + C$$

for some constant C . From above, as all $2^{f(n)}$ sums are distinct and less than nk ,

$$2^{f(n)} < nk = nf(n),$$

and so

$$f(n) < \log_2 n + \log_2 \log_2 n + O(1).$$

Examination of the second moment gives a modest improvement. Fix $\{x_1, \dots, x_k\} \subset \{1, \dots, n\}$ with distinct sums. Let $\epsilon_1, \dots, \epsilon_k$ be independent with

$$\Pr[\epsilon_i = 1] = \Pr[\epsilon_i = 0] = 1/2$$

and set

$$X = \epsilon_1 x_1 + \cdots + \epsilon_k x_k.$$

(We may think of X as a random sum.) Set

$$\mu = E[X] = \frac{x_1 + \cdots + x_k}{2}$$

and $\sigma^2 = \text{Var}[X]$. We bound

$$\sigma^2 = \frac{x_1^2 + \cdots + x_k^2}{4} \leq \frac{n^2 k}{4}$$

so that $\sigma \leq n\sqrt{k}/2$. By Chebyshev's Inequality for any $\lambda > 1$,

$$\Pr[|X - \mu| \geq \lambda n\sqrt{k}/2] \leq \lambda^{-2}.$$

Reversing,

$$1 - \frac{1}{\lambda^2} \leq \Pr[|X - \mu| < \lambda n\sqrt{k}/2].$$

But X has any particular value with probability either zero or 2^{-k} since, critically, a sum can be achieved in at most one way. Thus

$$\Pr[|X - \mu| < \lambda n\sqrt{k}/2] \leq 2^{-k}(\lambda n\sqrt{k} + 1)$$

and

$$n \geq \frac{2^k(1 - \lambda^{-2}) - 1}{\sqrt{k}\lambda}.$$

While $\lambda = \sqrt{3}$ gives optimal results any choice of $\lambda > 1$ gives the following.

Theorem 4.6.1 $f(n) \leq \log_2 n + (1/2) \log_2 \log_2 n + O(1)$.

17.4.7 4.7 The Rödl Nibble

For $2 \leq l < k < n$ let $M(n, k, l)$, the covering number, denote the minimal size of a family $\mathcal{K}$ of k -element subsets of $\{1, \dots, n\}$ having the property that every l -element set is contained in at least one $A \in \mathcal{K}$. Clearly $M(n, k, l) \geq \binom{n}{l} / \binom{k}{l}$ since each k -set covers $\binom{k}{l}$ l -sets and every l -set must be covered. Equality holds if and only if the family $\mathcal{K}$ has the property that every l -set is contained in exactly one $A \in \mathcal{K}$. This is called an (n, k, l) tactical configuration (or block design). For example, $(n, 3, 2)$ tactical configurations are better known as Steiner Triple Systems. The question of the existence of tactical configurations is a central one for combinatorics but one for which probabilistic methods (at least so far!) play little role. In 1963 Paul Erdős and Haim Hanani conjectured that for fixed $2 \leq l < k$,

$$\lim_{n \rightarrow \infty} \frac{M(n, k, l)}{\binom{n}{l} / \binom{k}{l}} = 1.$$

Their conjecture was, roughly, that one can get asymptotically close to a tactical configuration. While this conjecture seemed ideal for a probabilistic analysis it was a full generation before Rödl (1985) found the proof, which we describe in this section. [One may similarly define the packing number $m(n, k, l)$ as the maximal size of a family $\mathcal{K}$ of k -element subsets of $\{1, \dots, n\}$ having the property that every l -element set is contained in at most one $A \in \mathcal{K}$. Erdős and Hanani noticed from elementary arguments that

$$\lim_{n \rightarrow \infty} \frac{M(n, k, l)}{\binom{n}{l} / \binom{k}{l}} = 1 \iff \lim_{n \rightarrow \infty} \frac{m(n, k, l)}{\binom{n}{l} / \binom{k}{l}} = 1.$$

While the Rödl result may be formulated in terms of either packing or covering here we deal only with the covering problem.]

Several researchers realized that the Rödl method applies in a much more general setting, dealing with covers in uniform hypergraphs. This was first observed by Frankl and Rödl and has been simplified and extended by Pippenger and Spencer (1989) as well as by Kahn (1996). Our treatment here follows the one in Pippenger and Spencer (1989) and is based on the description of Füredi (1988), where the main tool is the second moment method.

For an r -uniform hypergraph $H = (V, E)$ and for a vertex $x \in V$, we let $d_H(x)$ [or simply $d(x)$, when there is no danger of confusion] denote the degree of x in H , that is, the number of edges containing x . Similarly, for $x, y \in V$, $d(x, y) = d_H(x, y)$ is the number of edges of H containing both x and y . A covering of H is a set of edges whose union contains all vertices. In what follows, whenever we write $\pm\delta$ we mean a quantity between $-\delta$ and δ . The following theorem is due to Pippenger, following Frankl and Rödl.

Theorem 4.7.1 For every integer $r \geq 2$ and reals $k \geq 1$ and $a > 0$, there are $\gamma = \gamma(r, k, a) > 0$ and $d_0 = d_0(r, k, a)$ such that for every $n \geq D \geq d_0$ the following holds.

Every r -uniform hypergraph $H = (V, E)$ on a set V of n vertices in which all vertices have positive degrees and which satisfies the following conditions:

- (1) For all vertices $x \in V$ but at most γn of them, $d(x) = (1 \pm \gamma)D$.
- (2) For all $x \in V$, $d(x) < kD$.
- (3) For any two distinct $x, y \in V$, $d(x, y) < \gamma D$.

contains a cover of at most $(1 + a)(n/r)$ edges.

The basic idea in the proof is simple. Fixing a small $\epsilon > 0$ one shows that a random set of roughly $\epsilon n/r$ edges has, with high probability, only some $O(\epsilon^2 n)$ vertices covered more than once, and hence covers at least $\epsilon n - O(\epsilon^2 n)$ vertices. Moreover, after deleting the vertices covered, the induced hypergraph on the remaining vertices still satisfies the properties described in (1), (2) and (3) above (for some other values of n, γ, k and D). Therefore one can choose again a random set of edges of this hypergraph, covering roughly an ϵ -fraction of its vertices with nearly no overlaps. Proceeding in this way for a large number of times we are finally left with at most ϵn uncovered vertices, and we then cover them trivially, by taking for each of them an arbitrarily chosen edge containing it. Since ϵ is sufficiently small, although this last step is very inefficient, it can be tolerated.

The technical details require a careful application of the second moment method, used several times in the proof of the following lemma.

Lemma 4.7.2 For every integer $r \geq 2$ and reals $K \geq 1$ and $\epsilon > 0$, and for every real $\delta' > 0$, there are $\delta = \delta(r, K, \epsilon, \delta') > 0$ and $D_0 = D_0(r, K, \epsilon, \delta')$ such that for every $n \geq D \geq D_0$ the following holds.

Every r -uniform hypergraph $H = (V, E)$ on a set V of n vertices which satisfies the following conditions:

- (i) For all vertices $x \in V$ but at most δn of them, $d(x) = (1 \pm \delta)D$.
- (ii) For all $x \in V$, $d(x) < KD$.
- (iii) For any two distinct $x, y \in V$, $d(x, y) < \delta D$.

contains a set E' of edges with the following properties:

- (iv) $|E'| = (\epsilon n/r)(1 \pm \delta')$.
- (v) The set $V' = V - \cup_{e \in E'} e$ is of cardinality $|V'| = ne^{-\epsilon}(1 \pm \delta')$.
- (vi) For all vertices $x \in V'$ but at most $\delta'|V'|$ of them, the degree $d'(x)$ of x in the induced hypergraph of H on V' satisfies $d'(x) = De^{-\epsilon(r-1)}(1 \pm \delta')$.

Proof. Throughout the proof we assume, whenever this is needed, that D (and hence n)

is sufficiently large. We denote by $\delta_1, \delta_2, \dots$ positive constants (that can be explicitly estimated) that tend to 0 when δ tends to 0 and D tends to infinity (for fixed r, K, ϵ). Therefore, by choosing δ and D_0 appropriately, we can ensure that each of those will be smaller than δ' .

Let E' be a random subset of E obtained by picking, randomly and independently, each edge in E to be a member of E' with probability $p = \epsilon/D$. We have to show that with positive probability, the properties (iv), (v) and (vi) hold.

The proof that (iv) holds is easy. Note that by the assumptions H has at least $(1 - \delta)n$ vertices of degree at least $(1 - \delta)D$, showing that its number of edges is at least $(1 - \delta)^2 nD/r$. Similarly, the number of edges of H does not exceed $[(1 + \delta)Dn + \delta nKD]/r$. Therefore $|E| = (1 \pm \delta_1) Dn/r$. It follows that the expected value of the size of E' satisfies $E[|E'|] = |E|p = (1 \pm \delta_1)(\epsilon n/r)$ and its variance is $\text{Var}[|E'|] = |E|p(1 - p) \leq (1 \pm \delta_1)(\epsilon n/r)$. Therefore, by Chebyshev's Inequality, for an appropriately chosen $\delta_2 > 0$,

$$\Pr \left[|E'| = (1 \pm \delta_2) \frac{\epsilon n}{r} \right] > 0.99$$

say, giving (iv).

To prove (v), define for each vertex $x \in V$ an indicator random variable I_x , where $I_x = 1$ if $x \notin \cup_{e \in E'} e$ and $I_x = 0$ otherwise. Note that $|V'| = \sum_{x \in V} I_x$. Call a vertex $x \in V$ good if $d(x) = (1 \pm \delta)D$; otherwise call it bad. If x is good, then

$$E[I_x] = \Pr[I_x = 1] = (1 - p)^{d(x)} = \left(1 - \frac{\epsilon}{D}\right)^{(1 \pm \delta)D} = e^{-\epsilon} (1 \pm \delta_3)$$

If x is bad then, clearly, $0 \leq E[I_x] \leq 1$. Since there are at most δn bad vertices it follows, by linearity of expectation, that the expected value of $|V'|$ is $ne^{-\epsilon} (1 \pm \delta_4)$. To compute the variance of $|V'| = \sum_{x \in V} I_x$, note that

$$\begin{aligned} \text{Var}[|V'|] &= \sum_{x \in V} \text{Var}[I_x] + \sum_{x, y \in V, x \neq y} \text{Cov}[I_x, I_y] \\ &\leq E[|V'|] + \sum_{x, y \in V, x \neq y} \text{Cov}[I_x, I_y] \end{aligned}$$

However,

$$\begin{aligned} \text{Cov}[I_x, I_y] &= E[I_x I_y] - E[I_x] E[I_y] \\ &= (1 - p)^{d(x) + d(y) - d(x, y)} - (1 - p)^{d(x) + d(y)} \\ &\leq (1 - p)^{-d(x, y)} - 1 \leq \left(1 - \frac{\epsilon}{D}\right)^{-\delta D} - 1 \leq \delta_5 \end{aligned}$$

It follows that

$$\text{Var}[|V'|] \leq E[|V'|] + \delta_5 n^2 \leq \delta_6 (E[|V'|])^2,$$

which, by Chebyshev, implies that with probability at least 0.99

$$|V'| = (1 \pm \delta_7) E[|V'|] = (1 \pm \delta_8) ne^{-\epsilon},$$

as claimed in (v).

It remains to prove (vi). To do so note, first, that all but at most $\delta_9 n$ vertices x satisfy the following two conditions:

- (A) $d(x) = (1 \pm \delta)D$, and
- (B) all but at most $\delta_{10}D$ edges $e \in E$ with $x \in e$ satisfy

$$|\{f \in E : x \notin f, f \cap e \neq \emptyset\}| = (1 \pm \delta_{11})(r-1)D. \quad (4.1)$$

Indeed, (A) holds for all but $\delta n < \delta_9 n/2$ vertices, by assumption. Moreover, the total number of edges containing vertices whose degrees are not $(1 \pm \delta)D$ is at most $\delta n K D$ and hence the number of vertices contained in more than $\delta_{10} D$ such edges is at most $\delta n K D r / (\delta_{10} D) \leq \delta_9 n/2$ for an appropriate choice of δ_9, δ_{10} . Note, next, that if $x \in e$ and e contains no vertex of degree which is not $(1 \pm \delta)D$ then, since $d(y, z) < \delta D$ for all y, z , the number of edges f not containing x that intersect e is at most $(r-1)(1 \pm \delta)D$ and at least $(r-1)(1 \pm \delta)D - \binom{r-1}{2} \delta D$ and hence e satisfies (4.1).

It thus suffices to show that for most of the vertices x satisfying (A) and (B), $d'(x)$ satisfies (iv). Fix such a vertex x . Call an edge e with $x \in e$ good if it satisfies (4.1). Conditioning on $x \in V'$, the probability that a good edge containing x stays in the hypergraph on V' is $(1-p)^{(1 \pm \delta_{11})(r-1)D}$. Therefore the expected value of $d'(x)$ is

$$\mathbb{E}[d'(x)] = (1 \pm \delta_{10} \pm \delta) D (1-p)^{(1 \pm \delta_{11})(r-1)D} \pm \delta_{10} D = e^{-\epsilon(r-1)} D (1 \pm \delta_{12}).$$

For each edge e containing x , let I_e denote the indicator random variable whose value is 1 iff e is contained in V' . Then, the degree $d'(x)$ is simply the sum of these indicator random variables, conditioned on $x \in V'$. It follows that

$$\begin{aligned} \text{Var}[d'(x)] &\leq \mathbb{E}[d'(x)] + \sum_{x \in e, x \in f} \text{Cov}[I_e, I_f] \\ &\leq \mathbb{E}[d'(x)] + 2\delta_{10} D^2 (1 \pm \delta) + \sum_{x \in e, x \in f, e, f \text{ good}} \text{Cov}[I_e, I_f]. \end{aligned}$$

It remains to bound the sum $\sum_{x \in e, x \in f, e, f \text{ good}} \text{Cov}[I_e, I_f]$. For each fixed good e this sum is a sum of the form $\sum_{x \in f, f \text{ good}} \text{Cov}[I_e, I_f]$. There are at most $(r-1)\delta D$ edges f in the last sum for which $|e \cap f| > 1$, and their contribution to the sum cannot exceed $(r-1)\delta D$. If $e \cap f = \{x\}$ then let $t(e, f)$ denote the number of edges of H that intersect both e and f and do not contain x . Clearly, in this case, $t(e, f) \leq (r-1)^2 \delta D$. It follows that for such e and f , $\text{Cov}[I_e, I_f] \leq (1-p)^{-t(e, f)} - 1 \leq \delta_{13}$, implying that for each fixed good edge e ,

$$\sum_{x \in f, f \text{ good}} \text{Cov}[I_e, I_f] \leq (r-1)\delta D + D(1+\delta)\delta_{13} \leq \delta_{14} D.$$

As the sum $\sum_{x \in e, x \in f, e, f \text{ good}} \text{Cov}[I_e, I_f]$ is the sum of at most $D(1+\delta)$ such quantities, we conclude that

$$\text{Var}[d'(x)] \leq \mathbb{E}[d'(x)] + \delta_{15} D^2 \leq \delta_{16} (\mathbb{E}[d'(x)])^2.$$

It thus follows, by Chebyshev, that with probability at most δ_{17} , $d'(x)$ is not $(1 \pm \delta_{18}) D e^{-\epsilon(r-1)}$, and therefore, by Markov, that with probability at least, say, 0.99, for all but at most $\delta_{19} n$ vertices, $d'(x) = (1 \pm \delta_{18}) D e^{-\epsilon(r-1)}$. This completes the proof of the lemma.

Proof [Theorem 4.7.1]. Fix $\epsilon > 0$ such that

$$\frac{\epsilon}{1 - e^{-\epsilon}} + r\epsilon < 1 + a$$

and fix $1/10 > \delta > 0$ such that

$$(1 + 4\delta) \frac{\epsilon}{1 - e^{-\epsilon}} + r\epsilon < 1 + a$$

Fix an integer t so that $e^{-\epsilon t} < \epsilon$. The theorem is proved by applying the lemma t times. Put $\delta = \delta_t$ and then define, by reverse induction $\delta_t > \delta_{t-1} > \dots > \delta_0$ such that

$\delta_i \leq \delta_{i+1}e^{-\epsilon(r-1)}$, $\prod_{i=0}^t (1 + \delta_i) < 1 + 2\delta$, and for $n \geq D \geq R_i$ one can apply the lemma with $r, K = ke^{\epsilon i(r-1)}$, $\epsilon, \delta' = \delta_{i+1}$ and $\delta = \delta_i$. This will give the assertion of the theorem with $\gamma = \delta_0, d_0 = \max R_i$. Indeed, by applying the lemma repeatedly we obtain a decreasing sequence of sets of vertices $V = V_0, V_1, \dots, V_t$, each contained in the previous one, and a sequence of sets of edges $E_1, E_2, \dots, E_t$, where E_i is the set of edges E' obtained in the application of the lemma to the hypergraph induced on V_{i-1} . Here

$$|V_i| = |V_{i-1}| e^{-\epsilon} (1 \pm \delta_i) \quad (= |V_0| e^{-i\epsilon} (1 \pm 2\delta)),$$

$$|E_i| = \frac{\epsilon |V_{i-1}|}{r} (1 \pm \delta_i) \leq (1 + 4\delta) \frac{\epsilon n}{r} e^{-(i-1)\epsilon},$$

and

$$D_i = D_{i-1}e^{-\epsilon(r-1)} = De^{-\epsilon i(r-1)}$$

By covering each vertex of V_t separately by an edge containing it we conclude that the total number of edges in the cover obtained is at most

$$\begin{aligned} (1 + 4\delta) \sum_{i=0}^{t-1} \frac{\epsilon n}{r} e^{-i\epsilon} + |V_t| &\leq (1 + 4\delta) \frac{\epsilon n}{r} \frac{1}{1 - e^{-\epsilon}} + (1 + 2\delta) n e^{-\epsilon t} \\ &\leq \frac{n}{r} (1 + 4\delta) \left(\frac{\epsilon}{1 - e^{-\epsilon}} + r\epsilon \right) \\ &< (1 + a) \frac{n}{r} \end{aligned}$$

This completes the proof.

We conclude the section by showing how the theorem quickly implies Rödl solution of the Erdős-Hanani problem mentioned at the beginning of the section.

Theorem 4.7.3 [Rödl] For k, l fixed,

$$M(n, k, l) \leq (1 + o(1)) \binom{n}{l} / \binom{k}{l}$$

where the $o(1)$ term tends to zero as n tends to infinity.

Proof. Put $r = \binom{k}{l}$ and let H be the r -uniform hypergraph whose vertices are all l -subsets of $\{1, 2, \dots, n\}$, and whose edges are all collections of $\binom{k}{l} l$ -tuples that lie in a k -set. H has $\binom{n}{l}$ vertices, each of its vertices has degree $D = \binom{n-l}{k-l}$, and every two distinct vertices lie in at most $\binom{n-l-1}{k-l-1} = o(D)$ common edges. Therefore, by Theorem 4.7.1, H has a cover of size at most $(1 + o(1)) \binom{n}{l} / \binom{k}{l}$, as needed.

17.4.8 4.8 Exercises

1. Let X be a random variable taking integral nonnegative values, let $E[X^2]$ denote the expectation of its square, and let $\text{Var}[X]$ denote its variance. Prove that

$$\Pr[X = 0] \leq \frac{\text{Var}[X]}{E[X^2]}$$

2. (*) Show that there is a positive constant c such that the following holds. For any n reals $a_1, a_2, \dots, a_n$ satisfying $\sum_{i=1}^n a_i^2 = 1$, if $(\epsilon_1, \dots, \epsilon_n)$ is a $\{-1, 1\}$ random vector obtained by choosing each ϵ_i randomly and independently with uniform distribution to be either -1 or 1, then

$$\Pr \left[\left| \sum_{i=1}^n \epsilon_i a_i \right| \leq 1 \right] \geq c.$$

3. (*) Show that there is a positive constant c such that the following holds. For any n vectors $a_1, a_2, \dots, a_n \in \mathbb{R}^2$ satisfying $\sum_{i=1}^n \|a_i\|^2 = 1$ and $\|a_i\| \leq 1/10$, where $\|\cdot\|$ denotes the usual Euclidean norm, if $(\epsilon_1, \dots, \epsilon_n)$ is a $\{-1, 1\}$ random vector obtained by choosing each ϵ_i randomly and independently with uniform distribution to be either -1 or 1, then

$$\Pr \left[\left\| \sum_{i=1}^n \epsilon_i a_i \right\| \leq 1/3 \right] \geq c.$$

4. Let X be a random variable with expectation $E[X] = 0$ and variance σ^2 . Prove that for all $\lambda > 0$,

$$\Pr[X \geq \lambda] \leq \frac{\sigma^2}{\sigma^2 + \lambda^2}$$

5. Let $v_1 = (x_1, y_1), \dots, v_n = (x_n, y_n)$ be n two-dimensional vectors, where each x_i and each y_i is an integer whose absolute value does not exceed $2^{n/2}/(100\sqrt{n})$. Show that there are two disjoint sets $I, J \subset \{1, 2, \dots, n\}$ such that

$$\sum_{i \in I} v_i = \sum_{j \in J} v_j.$$

6. (*) Prove that for every set X of at least $4k^2$ distinct residue classes modulo a prime p , there is an integer a such that the set $\{ax \pmod{p} : x \in X\}$ intersects every interval in $\{0, 1, \dots, p-1\}$ of length at least p/k .

The Probabilistic Lens: Hamiltonian Paths

What is the maximum possible number of directed Hamiltonian paths in a tournament on n vertices? Denote this number by $P(n)$. The first application of the probabilistic method in combinatorics is the result of Szele (1943) described in Chapter 2, which states that $P(n) \geq n!/2^{n-1}$. This bound follows immediately from the observation that the right-hand side is the expected number of such paths in a random tournament on n vertices. In the same paper Szele shows that

$$\frac{1}{2} \leq \lim_{n \rightarrow \infty} \left(\frac{P(n)}{n!} \right)^{1/n} \leq \frac{1}{2^{3/4}}$$

proves that this limit does exist, and conjectures that its correct value is $1/2$. This conjecture is proved in Alon (1990a). The proof is given below. The main tool is the Brégman proof of the Minc Conjecture for the permanent of a $(0, 1)$ -matrix, described in The Probabilistic Lens: Brégman Theorem (following Chapter 2).

Theorem 1 There exists a positive constant c such that for every n ,

$$P(n) \leq cn^{3/2} \frac{n!}{2^{n-1}}$$

Proof. For a tournament T , denote by $P(T)$ the number of directed Hamiltonian paths of T . Similarly, $C(T)$ denotes the number of directed Hamiltonian cycles of T , and $F(T)$ denotes the number of spanning subgraphs of T in which the indegree and the outdegree of every vertex is exactly 1. Clearly,

$$C(T) \leq F(T) \quad (1)$$

If $T = (V, E)$ is a tournament on a set $V = \{1, 2, \dots, n\}$ of n vertices, the adjacency matrix of T is the n by n $(0, 1)$ -matrix $A_T = (a_{ij})$ defined by $a_{ij} = 1$ if $(i, j) \in E$ and $a_{ij} = 0$ otherwise. Let r_i denote the number of ones in row i . Clearly,

$$\sum_i^n r_i = \binom{n}{2} \quad (2)$$

By interpreting combinatorially the terms in the expansion of the permanent $\text{per}(A_T)$, it follows that

$$\text{per}(A_T) = F(T) \quad (3)$$

We need the following technical lemma.

Lemma 2 For every two integers a, b satisfying $b \geq a + 2 > a \geq 1$ the inequality

$$(a!)^{1/a} \cdot (b!)^{1/b} < ((a+1)!)^{1/(a+1)} \cdot ((b-1)!)^{1/(b-1)}$$

holds.

Proof. The assertion is simply that $f(a) < f(b-1)$, for the function f defined by $f(a) = (a!)^{1/a} / ((a+1)!)^{1/(a+1)}$. Thus it suffices to show that for every integer $x \geq 2$, $f(x-1) < f(x)$. Substituting the expression for f and raising both sides to the power $x(x-1)(x+1)$ it follows that it suffices to show that for all $x \geq 2$,

$$((x-1)!)^{x(x+1)} \cdot ((x+1)!)^{x(x-1)} < (x!)^{2(x^2-1)}$$

that is,

$$\left(\frac{x^x}{x!}\right)^2 > \left(\frac{x+1}{x}\right)^{x(x-1)}$$

This is certainly true for $x = 2$. For $x \geq 3$ it follows from the facts that $4^x > e^{x+1}$, that $x! < ((x+1)/2)^x$ and that $e^{x-1} > ((x+1)/x)^{x(x-1)}$.

Corollary 3 Define $g(x) = (x!)^{1/x}$. For every integer $S \geq n$ the maximum of the function $\prod_{i=1}^n g(x_i)$ subject to the constraints $\sum_{i=1}^n x_i = S$ and $x_i \geq 1$ are integers, is obtained iff the variables x_i are as equal as possible (i.e., if each x_i is either $\lfloor S/n \rfloor$ or $\lceil S/n \rceil$.)

Proof. If there are indices i and j such that $x_i \geq x_j + 2$ then, by Lemma 2, the value of the product would increase once we add one to x_j and subtract one from x_i .

Returning to our tournament T we observe that the numbers r_i defined above are precisely the outdegrees of the vertices of T . If at least one of these is 0, then clearly $C(T) = F(T) = 0$. Otherwise, by Brégman's Theorem, by Corollary 3 and by (2) and (3), $F(T)$ is at most the value of the function $\prod_{i=1}^n (r_i!)^{1/r_i}$, where the integral variables r_i satisfy (2) and are as equal as possible. By a straightforward (though somewhat tedious) derivation of the asymptotics using Stirling's formula this gives the following.

Proposition 4 For every tournament T on n vertices,

$$C(T) \leq F(T) \leq (1 + o(1)) \frac{\sqrt{\pi}}{\sqrt{2e}} n^{3/2} \frac{(n-1)!}{2^n}$$

To complete the proof of the theorem, we have to derive a bound for the number of Hamiltonian paths in a tournament from the above result. Given a tournament S on n vertices, let T be the random tournament obtained from S by adding to it a new vertex y and by orienting each edge connecting y with one of the vertices of S , randomly and independently. For every fixed Hamiltonian path in S , the probability that it can be extended to a Hamiltonian cycle in T is precisely $1/4$. Thus the expected number of Hamiltonian cycles in T is $\frac{1}{4}P(S)$ and hence there is a specific T for which $C(T) \geq \frac{1}{4}P(S)$. However, by Proposition 4,

$$C(T) \leq (1 + o(1)) \frac{\sqrt{\pi}}{\sqrt{2e}} (n+1)^{3/2} \frac{n!}{2^{n+1}}$$

and thus

$$P(S) \leq O\left(n^{3/2} \frac{n!}{2^{n-1}}\right)$$

completing the proof of Theorem 1.

17.5 5 The Local Lemma

It's a thing that non-mathematicians don't realize. Mathematics is actually an esthetic subject almost entirely.

- John Conway

17.5.1 5.1 The Lemma

In a typical probabilistic proof of a combinatorial result, one usually has to show that the probability of a certain event is positive. However, many of these proofs actually give more and show that the probability of the event considered is not only positive but is large. In fact, most probabilistic proofs deal with events that hold with high probability; that is, a probability that tends to 1 as the dimensions of the problem grow. For example, consider the proof given in Chapter 1 that for each $k \geq 1$ there are tournaments in which for every set of k players there is one who beats them all. The proof actually shows that for every fixed k if the number n of players is sufficiently large then almost all tournaments with n players satisfy this property; that is, the probability that a random tournament with n players has the desired property tends to 1 as n tends to infinity.

On the other hand, there is a trivial case in which one can show that a certain event holds with positive, though very small, probability. Indeed, if we have n mutually independent events and each of them holds with probability at least $p > 0$, then the probability that all events hold simultaneously is at least p^n , which is positive, although it may be exponentially small in n .

It is natural to expect that the case of mutual independence can be generalized to that of rare dependencies and provide a more general way of proving that certain events hold with positive, though small, probability. Such a generalization is, indeed, possible and is stated in the following lemma, known as the Lovász Local Lemma. This simple lemma, first proved in Erdős and Lovász (1975), is an extremely powerful tool, as it supplies a way for dealing with rare events.

Lemma 5.1.1 [The Local Lemma; General Case] Let $A_1, A_2, \dots, A_n$ be events in an arbitrary probability space. A directed graph $D = (V, E)$ on the set of vertices $V = \{1, 2, \dots, n\}$ is called a dependency digraph for the events $A_1, \dots, A_n$ if for each $i, 1 \leq i \leq n$, the event A_i is mutually independent of all the events $\{A_j : (i, j) \notin E\}$. Suppose that $D = (V, E)$ is a dependency

digraph for the above events and suppose there are real numbers $x_1, \dots, x_n$ such that $0 \leq x_i < 1$ and $\Pr[A_i] \leq x_i \prod_{(i,j) \in E} (1 - x_j)$ for all $1 \leq i \leq n$. Then

$$\Pr \left[\bigwedge_{i=1}^n \overline{A_i} \right] \geq \prod_{i=1}^n (1 - x_i)$$

In particular, with positive probability no event A_i holds.

Proof. We first prove, by induction on s , that for any $S \subset \{1, \dots, n\}$, $|S| = s < n$, and any $i \notin S$,

$$\Pr \left[A_i \mid \bigwedge_{j \in S} \overline{A_j} \right] \leq x_i \quad (5.1)$$

This is certainly true for $s = 0$. Assuming it holds for all $s' < s$, we prove it for s . Put $S_1 = \{j \in S : (i, j) \in E\}$, $S_2 = S \setminus S_1$. Then

$$\Pr \left[A_i \mid \bigwedge_{j \in S} \overline{A_j} \right] = \frac{\Pr \left[A_i \wedge \left(\bigwedge_{j \in S_1} \overline{A_j} \right) \mid \bigwedge_{\ell \in S_2} \overline{A_\ell} \right]}{\Pr \left[\bigwedge_{j \in S_1} \overline{A_j} \mid \bigwedge_{\ell \in S_2} \overline{A_\ell} \right]} \quad (5.2)$$

To bound the numerator observe that since A_i is mutually independent of the events $\{A_\ell : \ell \in S_2\}$,

$$\begin{aligned} \Pr \left[A_i \wedge \left(\bigwedge_{j \in S_1} \overline{A_j} \right) \mid \bigwedge_{\ell \in S_2} \overline{A_\ell} \right] &\leq \Pr \left[A_i \mid \bigwedge_{\ell \in S_2} \overline{A_\ell} \right] \\ &= \Pr[A_i] \leq x_i \prod_{(i,j) \in E} (1 - x_j). \end{aligned} \quad (5.3)$$

The denominator, on the other hand, can be bounded by the induction hypothesis. Indeed, suppose $S_1 = \{j_1, j_2, \dots, j_r\}$. If $r = 0$ then the denominator is 1, and (5.1) follows. Otherwise

$$\begin{aligned} &\Pr \left[\overline{A_{j_1}} \wedge \overline{A_{j_2}} \wedge \dots \wedge \overline{A_{j_r}} \mid \bigwedge_{\ell \in S_2} \overline{A_\ell} \right] \\ &= \left(1 - \Pr \left[A_{j_1} \mid \bigwedge_{\ell \in S_2} \overline{A_\ell} \right] \right) \cdot \left(1 - \Pr \left[A_{j_2} \mid \overline{A_{j_1}} \wedge \bigwedge_{\ell \in S_2} \overline{A_\ell} \right] \right) \cdots \\ &\quad \cdots \left(1 - \Pr \left[A_{j_r} \mid \overline{A_{j_1}} \wedge \dots \wedge \overline{A_{j_{r-1}}} \wedge \bigwedge_{\ell \in S_2} \overline{A_\ell} \right] \right) \\ &\geq (1 - x_{j_1}) (1 - x_{j_2}) \cdots (1 - x_{j_r}) \geq \prod_{(i,j) \in E} (1 - x_j) \end{aligned} \quad (5.4)$$

Substituting (5.3) and (5.4) into (5.2) we conclude that $\Pr[A_i \mid \bigwedge_{j \in S} \overline{A_j}] \leq x_i$, completing the proof of the induction.

The assertion of Lemma 5.1.1 now follows easily, as

$$\begin{aligned} \Pr \left[\bigwedge_{i=1}^n \overline{A_i} \right] &= (1 - \Pr[A_1]) \cdot (1 - \Pr[A_2 \mid \overline{A_1}]) \cdots \\ &\quad \cdots \left(1 - \Pr \left[A_n \mid \bigwedge_{i=1}^{n-1} \overline{A_i} \right] \right) \geq \prod_{i=1}^n (1 - x_i), \end{aligned}$$

completing the proof.

Corollary 5.1.2 [The Local Lemma; Symmetric Case] Let $A_1, A_2, \dots, A_n$ be events in an arbitrary probability space. Suppose that each event A_i is mutually independent of a set of all the other events A_j but at most d , and that $\Pr[A_i] \leq p$ for all $1 \leq i \leq n$. If

$$ep(d+1) \leq 1 \quad (5.5)$$

then $\Pr\left[\bigwedge_{i=1}^n \overline{A_i}\right] > 0$.

Proof. If $d = 0$ the result is trivial. Otherwise, by the assumption there is a dependency digraph $D = (V, E)$ for the events $A_1, \dots, A_n$ in which for each i , $|\{j : (i, j) \in E\}| \leq d$. The result now follows from Lemma 5.1.1 by taking $x_i = 1/(d+1) (< 1)$ for all i and using the fact that for any $d \geq 1$,

$$\left(1 - \frac{1}{d+1}\right)^d > \frac{1}{e}.$$

It is worth noting that, as shown by Shearer (1985), the constant " e " is the best possible constant in inequality (5.5). Note also that the proof of Lemma 5.1.1 indicates that the conclusion remains true even when we replace the two assumptions that each A_i is mutually independent of $\{A_j : (i, j) \notin E\}$ and that for each i

$$\Pr[A_i] \leq x_i \prod_{(ij) \in E} (1 - x_j)$$

by the weaker assumption that for each i and each $S_2 \subset \{1, \dots, n\} \setminus \{j : (i, j) \in E\}$,

$$\Pr\left[A_i \mid \bigwedge_{j \in S_2} \overline{A_j}\right] \leq x_i \prod_{(i,j) \in E} (1 - x_j)$$

This turns out to be useful in certain applications.

In the next few sections we present various applications of the Local Lemma for obtaining combinatorial results. There is no known proof of any of these results, which does not use the Local Lemma. Additional applications of the Local Lemma for coloring problems, and much much more, can be found in Molloy and Reed (1999).

17.5.2 5.2 Property b and Multicolored Sets of Real Numbers

Recall that a hypergraph $H = (V, E)$ has property B (i.e., is two-colorable), if there is a coloring of V by two colors so that no edge $f \in E$ is monochromatic.

Theorem 5.2.1 Let $H = (V, E)$ be a hypergraph in which every edge has at least k elements, and suppose that each edge of H intersects at most d other edges. If $e(d+1) \leq 2^{k-1}$ then H has property B .

Proof. Color each vertex v of H , randomly and independently, either blue or red (with equal probability). For each edge $f \in E$, let A_f be the event that f is monochromatic. Clearly $\Pr[A_f] = 2/2^{|f|} \leq 1/2^{k-1}$. Moreover, each event A_f is clearly mutually independent of all the other events $A_{f'}$ for all edges f' that do not intersect f . The result now follows from Corollary 5.1.2.

A special case of Theorem 5.2.1 is that for any $k \geq 9$, any k -uniform k -regular hypergraph H has property B . Indeed, since any edge f of such an H contains k vertices, each of which is incident with k edges (including f), it follows that f intersects at most $d = k(k-1)$ other edges. The desired result follows, since $e(k(k-1)+1) < 2^{k-1}$ for each $k \geq 9$.

The next result we consider, which appeared in the original paper of Erdős and Lovász, deals with k -colorings of the real numbers. For a k -coloring $c : \mathbb{R} \rightarrow \{1, 2, \dots, k\}$ of the real

numbers by the k colors $1, 2, \dots, k$, and for a subset $T \subset \mathbb{R}$, we say that T is multicolored (with respect to c) if $c(T) = \{1, 2, \dots, k\}$; that is, if T contains elements of all colors.

Theorem 5.2.2 Let m and k be two positive integers satisfying

$$e(m(m-1) + 1)k \left(1 - \frac{1}{k}\right)^m \leq 1 \quad (5.6)$$

Then, for any set S of m real numbers there is a k -coloring so that each translation $x + S$ (for $x \in \mathbb{R}$) is multicolored.

Note that (5.6) holds whenever $m > (3 + o(1))k \log k$.

Proof. We first fix a finite subset $X \subseteq \mathbb{R}$ and show the existence of a k -coloring so that each translation $x + S$ (for $x \in X$) is multicolored. This is an easy consequence of the Local Lemma. Indeed, put $Y = \bigcup_{x \in X} (x + S)$ and let $c : Y \rightarrow \{1, 2, \dots, k\}$ be a random k -coloring of Y obtained by choosing, for each $y \in Y$, randomly and independently, $c(y) \in \{1, 2, \dots, k\}$ according to a uniform distribution on $\{1, 2, \dots, k\}$. For each $x \in X$, let A_x be the event that $x + S$ is not multicolored (with respect to c). Clearly $\Pr[A_x] \leq k(1 - 1/k)^m$. Moreover, each event A_x is mutually independent of all the other events $A_{x'}$ but those for which $(x + S) \cap (x' + S) \neq \emptyset$. As there are at most $m(m-1)$ such events, the desired result follows from Corollary 5.1.2.

We can now prove the existence of a coloring of the set of all reals with the desired properties, by a standard compactness argument. Since the discrete space with k points is (trivially) compact, Tikhonov's Theorem (which is equivalent to the axiom of choice) implies that an arbitrary product of such spaces is compact. In particular, the space of all functions from $\mathbb{R}$ to $\{1, 2, \dots, k\}$, with the usual product topology, is compact. In this space for every fixed $x \in \mathbb{R}$, the set C_x of all colorings c , such that $x + S$ is multicolored, is closed. (In fact, it is both open and closed, since a basis to the open sets is the set of all colorings whose values are prescribed in a finite number of places). As we proved above, the intersection of any finite number of sets C_x is nonempty. It thus follows, by compactness, that the intersection of all sets C_x is nonempty. Any coloring in this intersection has the properties in the conclusion of Theorem 5.2.2.

Note that it is impossible, in general, to apply the Local Lemma to an infinite number of events and conclude that in some point of the probability space none of them holds. In fact, there are trivial examples of countably many mutually independent events A_i , satisfying $\Pr[A_i] = 1/2$ and $\bigwedge_{i \geq 1} \overline{A_i} = \emptyset$. Thus the compactness argument is essential in the above proof.

17.5.3 5.3 Lower Bounds for Ramsey Numbers

The derivation of lower bounds for Ramsey numbers by Erdős in 1947 was one of the first applications of the probabilistic method. The Local Lemma provides a simple way of improving these bounds. Let us obtain, first, a lower bound for the diagonal Ramsey number $R(k, k)$. Consider a random two-coloring of the edges of K_n . For each set S of k vertices of K_n , let A_S be the event that the complete graph on S is monochromatic. Clearly $\Pr[A_S] = 2^{1-\binom{k}{2}}$. It is obvious that each event A_S is mutually independent of all the events A_T , but those which satisfy $|S \cap T| \geq 2$, since this is the only case in which the corresponding complete graphs share an edge. We can therefore apply Corollary 5.1.2 with $p = 2^{1-\binom{k}{2}}$ and $d < \binom{k}{2} \binom{n-2}{k-2}$ to conclude the following.

Proposition 5.3.1 If $e \binom{k}{2} \binom{n-2}{k-2} \cdot 2^{1-\binom{k}{2}} < 1$ then $R(k, k) > n$.

A short computation shows that this gives $R(k, k) > (\sqrt{2}/e)(1 + o(1))k2^{k/2}$, only a factor 2 improvement on the bound obtained by the straightforward probabilistic method. Although this minor improvement is somewhat disappointing it is certainly not surprising; the Local Lemma is most powerful when the dependencies between events are rare, and this is not the case here.

Indeed, there is a total number of $K = \binom{n}{k}$ events considered, and the maximum outdegree d in the dependency digraph is roughly $\binom{k}{2} \binom{n}{k-2}$. For large k and much larger n (which is the case of interest for us) we have $d > K^{1-O(1/k)}$, that is, quite a lot of dependencies. On the other hand, if we consider small sets S (e.g., sets of size 3) we observe that out of the total $K = \binom{n}{3}$ of them each shares an edge with only $3(n-3) \approx K^{1/3}$. This suggests that the Local Lemma may be much more significant in improving the off-diagonal Ramsey numbers $R(k, \ell)$, especially if one of the parameters, say, ℓ , is small. Let us consider, for example, following Spencer (1977), the Ramsey number $R(k, 3)$. Here, of course, we have to apply the nonsymmetric form of the Local Lemma. Let us two-color the edges of K_n randomly and independently, where each edge is colored blue with probability p . For each set of three vertices T , let A_T be the event that the triangle on T is blue. Similarly, for each set of k vertices S , let B_S be the event that the complete graph on S is red. Clearly $\Pr[A_T] = p^3$ and $\Pr[B_S] = (1-p)^{\binom{k}{2}}$. Construct a dependency digraph for the events A_T and B_S by joining two vertices by edges (in both directions) iff the corresponding complete graphs share an edge. Clearly each A_T -node of the dependency graph is adjacent to $3(n-3) < 3nA_{T'}$ -nodes and to at most $\binom{n}{k}B_{S'}$ -nodes. Similarly, each B_S -node is adjacent to at most $\binom{k}{2}(n-2) < k^2n/2A_{T'}$ -nodes and to at most $\binom{n}{k}B_{S'}$ -nodes. It follows from the general case of the Local Lemma (Lemma 5.1.1) that if we can find a $0 < p < 1$ and two real numbers $0 \leq x < 1$ and $0 \leq y < 1$ such that

$$p^3 \leq x(1-x)^{3n}(1-y)^{\binom{n}{k}}$$

and

$$(1-p)^{\binom{k}{2}} \leq y(1-x)^{k^2n/2}(1-y)^{\binom{n}{k}}$$

then $R(k, 3) > n$.

Our objective is to find the largest possible $k = k(n)$ for which there is such a choice of p, x and y . An elementary (but tedious) computation shows that the best choice is when $p = c_1n^{-1/2}$, $k = c_2n^{1/2} \log n$, $x = c_3/n^{3/2}$ and y so that $\binom{n}{k}y = c_4$. This gives $R(k, 3) > c_5k^2/\log^2 k$. A similar argument gives $R(k, 4) > k^{5/2+o(1)}$. In both cases the amount of computation required is considerable. However, the hard work does pay; the bound $R(k, 3) > c_5k^2/\log^2 k$ matches a lower bound of Erdős proved in 1961 by a highly complicated probabilistic argument. This was improved to $R(k, 3) > c_6k^2/\log k$ by Kim(1995). The bound above for $R(k, 4)$ is better than any bound for $R(k, 4)$ known to be proved without the Local Lemma.

17.5.4 5.4 A Geometric Result

A family of open unit balls F in the three-dimensional Euclidean space $\mathbb{R}^3$ is called a k -fold covering of $\mathbb{R}^3$ if any point $x \in \mathbb{R}^3$ belongs to at least k balls. In particular, a 1 fold covering is simply called a covering. A k -fold covering $\mathcal{F}$ is called decomposable if there is a partition of $\mathcal{F}$ into two pairwise disjoint families $\mathcal{F}_1$ and $\mathcal{F}_2$, each being a covering of $\mathbb{R}^3$. Mani-Levitska and Pach (1988) constructed, for any integer $k \geq 1$, a non decomposable k -fold covering of $\mathbb{R}^3$ by open unit balls. On the other hand, they proved that any k -fold covering of $\mathbb{R}^3$ in which no point is covered by more than $c2^{k/3}$ balls is decomposable. This reveals a somewhat surprising phenomenon: it is more difficult to decompose coverings that cover some of the points of $\mathbb{R}^3$ too often than to decompose coverings that cover every point about the same number of times. The exact statement of the Mani-Levitska-Pach Theorem is the following.

Theorem 5.4.1 Let $\mathcal{F} = \{B_i\}_{i \in I}$ be a k -fold covering of the three-dimensional Euclidean space by open unit balls. Suppose, further, that no point of $\mathbb{R}^3$ is contained in more than t members of $\mathcal{F}$. If

$$e \cdot t^3 2^{18} / 2^{k-1} \leq 1$$

then $\mathcal{F}$ is decomposable.

Proof. Define an infinite hypergraph $H = (V(H), E(H))$ as follows. The set of vertices of H , $V(H)$, is simply $\mathcal{F} = \{B_i\}_{i \in I}$. For each $x \in \mathbb{R}^3$ let E_x be the set of balls $B_i \in \mathcal{F}$ that contain x . The set of edges of H , $E(H)$, is simply the set of E_x , with the understanding that when $E_x = E_y$ the edge is taken only once. We claim each edge E_x intersects less than $t^3 2^{18}$ other edges E_y of H . If $x \in B_i$ the center of B_i is within distance 1 of x . If now $B_j \cap B_i \neq \emptyset$ the center of B_j is within distance three of x and so B_j lies entirely inside the ball of radius four centered at x . Such a B_j covers precisely $4^{-3} = 2^{-6}$ of the volume of that ball. As no vertex is covered more than t times there can be at most $2^6 t$ such balls. It is not too difficult to check that m balls in $\mathbb{R}^3$ cut $\mathbb{R}^3$ into less than m^3 connected components so that there are at most $(2^6 t)^3$ distinct E_y overlapping E_x .

Consider, now, any finite subhypergraph L of H . Each edge of L has at least k vertices, and it intersects at most $d < t^3 2^{18}$ other edges of L . Since, by assumption, $e(d+1) \leq 2^{k-1}$, Theorem 5.2.1 (which is a simple corollary of the Local Lemma), implies that L is two-colorable. This means that one can color the vertices of L blue and red so that no edge of L is monochromatic. Since this holds for any finite L , a compactness argument, analogous to the one used in the proof of Theorem 5.2.2, shows that H is two-colorable. Given a two-coloring of H with no monochromatic edges, we simply let $\mathcal{F}_1$ be the set of all blue balls, and $\mathcal{F}_2$ be the set of all red ones. Clearly each $\mathcal{F}_i$ is a covering of $\mathbb{R}^3$, completing the proof of the theorem.

It is worth noting that Theorem 5.4.1 can easily be generalized to higher dimensions. We omit the detailed statement of this generalization.

17.5.5 5.5 The Linear Arboricity of Graphs

A linear forest is a forest (i.e., an acyclic simple graph) in which every connected component is a path. The linear arboricity $\text{la}(G)$ of a graph G is the minimum number of linear forests in G , whose union is the set of all edges of G . This notion was introduced by Harary as one of the covering invariants of graphs. The following conjecture, known as the Linear Arboricity Conjecture, was raised in Akiyama, Exoo and Harary (1981).

Conjecture 5.5.1 [The Linear Arboricity Conjecture] The linear arboricity of every d -regular graph is $\lceil (d+1)/2 \rceil$.

Note that since every d -regular graph G on n vertices has $nd/2$ edges, and every linear forest in it has at most $n-1$ edges, the inequality

$$\text{la}(G) \geq \frac{nd}{2(n-1)} > \frac{d}{2}$$

is immediate. Since $\text{la}(G)$ is an integer this gives $\text{la}(G) \geq \lceil (d+1)/2 \rceil$. The difficulty in Conjecture 5.5.1 lies in proving the converse inequality: $\text{la}(G) \leq \lceil (d+1)/2 \rceil$. Note also that since every graph G with maximum degree Δ is a subgraph of a Δ -regular graph (which may have more vertices, as well as more edges than G), the Linear Arboricity Conjecture is equivalent to the statement that the linear arboricity of every graph G with maximum degree Δ is at most $\lceil (\Delta+1)/2 \rceil$.

Although this conjecture received a considerable amount of attention, the best general result concerning it, proved without any probabilistic arguments, is that $\text{la}(G) \leq \lceil 3\Delta/5 \rceil$ for even Δ and that $\text{la}(G) \leq \lceil (3\Delta+2)/5 \rceil$ for odd Δ . In this section we prove that for every $\varepsilon > 0$ there is a $\Delta_0 = \Delta_0(\varepsilon)$ such that for every $\Delta \geq \Delta_0$, the linear arboricity of every graph with maximum degree Δ is less than $(\frac{1}{2} + \varepsilon)\Delta$. This result (with a somewhat more complicated

proof) appears in Alon (1988) and its proof relies heavily on the Local Lemma. We note that this proof is more complicated than the other proofs given in this chapter and requires certain preparations, some of which are of independent interest.

It is convenient to deduce the result for undirected graphs from its directed version. A d -regular digraph is a directed graph in which the indegree and the outdegree of every vertex is precisely d . A linear directed forest is a directed graph in which every connected component is a directed path. The dilinear arboricity $\text{dla}(G)$ of a directed graph G is the minimum number of linear directed forests in G whose union covers all edges of G . The directed version of the Linear Arboricity Conjecture, first stated in Nakayama and Peroche (1987), is the following.

Conjecture 5.5.2 For every d -regular digraph D ,

$$\text{dla}(D) = d + 1.$$

Note that since the edges of any (connected) undirected $2d$ -regular graph G can be oriented along a Euler cycle, so that the resulting oriented digraph is d -regular, the validity of Conjecture 5.5.2 for d implies that of Conjecture 5.5.1 for $2d$.

It is easy to prove that any graph with n vertices and maximum degree d contains an independent set of size at least $n/(d+1)$. The following proposition shows that at the price of decreasing the size of such a set by a constant factor we can guarantee that it has a certain structure.

Proposition 5.5.3 Let $H = (V, E)$ be a graph with maximum degree d , and let $V = V_1 \cup V_2 \cup \dots \cup V_r$ be a partition of V into r pairwise disjoint sets. Suppose each set V_i is of cardinality $|V_i| \geq 2ed$, where e is the basis of the natural logarithm. Then there is an independent set of vertices $W \subseteq V$ that contains a vertex from each V_i .

Proof. Clearly we may assume that each set V_i is of cardinality precisely $g = \lceil 2ed \rceil$ (otherwise, simply replace each V_i by a subset of cardinality g of it, and replace H by its induced subgraph on the union of these r new sets). Let us pick from each set V_i randomly and independently a single vertex according to a uniform distribution. Let W be the random set of the vertices picked. To complete the proof we show that with positive probability W is an independent set of vertices in H .

For each edge f of H , let A_f be the event that W contains both ends of f . Clearly, $\Pr[A_f] \leq 1/g^2$. Moreover, if the endpoints of f are in V_i and in V_j , then the event A_f is mutually independent of all the events corresponding to edges whose endpoints do not lie in $V_i \cup V_j$. Thus there is a dependency digraph for the events in which the maximum degree is less than $2gd$, and since $e \cdot 2gd \cdot 1/g^2 = 2ed/g < 1$ we conclude, by Corollary 5.1.2, that with positive probability none of the events A_f holds. But this means that W is an independent set containing a vertex from each V_i , completing the proof. $\square$

Proposition 5.5.3 suffices to prove Conjecture 5.5.2 for digraphs with no short directed cycle. Recall that the directed girth of a digraph is the minimum length of a directed cycle in it.

Theorem 5.5.4 Let $G = (U, F)$ be a d -regular digraph with directed girth $g \geq 8ed$. Then

$$\text{dla}(G) = d + 1.$$

Proof. As is well known, F can be partitioned into d pairwise disjoint 1-regular spanning subgraphs $F_1, \dots, F_d$ of G . [This is an easy consequence of the HallKönig Theorem; let H be the bipartite graph whose two classes of vertices A and B are copies of U , in which $u \in A$ is joined to $v \in B$ iff $(u, v) \in F$. Since H is d -regular its edges can be decomposed into d perfect matchings, which correspond to d 1-regular spanning subgraphs of G .] Each F_i is a union of vertex disjoint directed cycles $C_{i1}, C_{i2}, \dots, C_{ir_i}$. Let $V_1, V_2, \dots, V_r$ be the sets of edges of all the cycles $\{C_{ij} : 1 \leq i \leq d, 1 \leq j \leq r_i\}$. Clearly $V_1, V_2, \dots, V_r$ is a partition of the set F of all edges

of G , and by the girth condition, $|V_i| \geq g \geq 8ed$ for all $1 \leq i \leq r$. Let H be the line graph of G , that is, the graph whose set of vertices is the set F of edges of G in which two edges are adjacent iff they share a common vertex in G . Clearly H is $4d - 2$ regular. As the cardinality of each V_i is at least $8ed \geq 2e(4d - 2)$, there is, by Proposition 5.5.3, an independent set of H containing a member from each V_i . But this means that there is a matching M in G , containing at least one edge from each cycle C_{ij} of the 1-factors $F_1, \dots, F_d$. Therefore $M, F_1 \setminus M, F_2 \setminus M, \dots, F_d \setminus M$ are $d + 1$ directed forests in G (one of which is a matching) that cover all its edges. Hence

$$\text{dla}(G) \leq d + 1$$

As G has $|U| \cdot d$ edges and each directed linear forest can have at most $|U| - 1$ edges,

$$\text{dla}(G) \geq |U|d / (|U| - 1) > d$$

Thus $\text{dla}(G) = d + 1$, completing the proof.

The last theorem shows that the assertion of Conjecture 5.5.2 holds for digraphs with sufficiently large (directed) girth. In order to deal with digraphs with small girth, we show that most of the edges of each regular digraph can be decomposed into a relatively small number of almost regular digraphs with high girth. To do this, we need the following statement, which is proved using the Local Lemma.

Lemma 5.5.5 Let $G = (V, E)$ be a d -regular directed graph, where d is sufficiently large, and let p be an integer satisfying $10\sqrt{d} \leq p \leq 20\sqrt{d}$. Then, there is a p -coloring of the vertices of G by the colors $0, 1, 2, \dots, p - 1$ with the following property; for each vertex $v \in V$ and each color i , the numbers

$$N^+(v, i) = |\{u \in V : (v, u) \in E \text{ and } u \text{ is colored } i\}|$$

and

$$N^-(v, i) = |\{u \in V : (u, v) \in E \text{ and } u \text{ is colored } i\}|$$

satisfy

$$|N^-(v, i) - d/p|, |N^+(v, i) - d/p| \leq 3\sqrt{d/p} \sqrt{\log d} \quad (5.7)$$

Proof. Let $f : V \rightarrow \{0, 1, \dots, p - 1\}$ be a random vertex coloring of V by p colors, where for each $v \in V$, $f(v) \in \{0, 1, \dots, p - 1\}$ is chosen according to a uniform distribution. For every vertex $v \in V$ and every color i , $0 \leq i < p$, let $A_{v,i}^+$ be the event that the number $N^+(v, i)$ of neighbors of v in G whose color is i does not satisfy inequality (5.7). Clearly $N^+(v, i)$ is a binomial random variable with expectation d/p and standard deviation $\sqrt{(d/p)(1 - 1/p)} < \sqrt{d/p}$. Hence, by the standard estimates for binomial distribution given in Appendix A, for every $v \in V$ and $0 \leq i < p$,

$$\Pr[A_{v,i}^+] < 1/d^4.$$

Similarly, if $A_{v,i}^-$ is the event that the number $N^-(v, i)$ violates (5.7) then

$$\Pr[A_{v,i}^-] < 1/d^4$$

Clearly each of the events $A_{v,i}^+$ or $A_{v,i}^-$ is mutually independent of all the events $A_{u,j}^+$ or $A_{u,j}^-$ for all vertices $u \in V$ that do not have a common neighbor with v in G . Thus there is a dependency digraph for all our events with maximum degree $\leq (2d)^2 \cdot p$. Since $e \cdot (1/d^4) ((2d)^2 p + 1) < 1$, Corollary 5.1.2 (i.e., the symmetric form of the Local Lemma) implies that with positive

probability no event $A_{v,i}^+$ or $A_{v,i}^-$ occurs. Hence there is a coloring f that satisfies (5.7) for all $v \in V$ and $0 \leq i < p$, completing the proof.

We are now ready to deal with general regular digraphs. Let $G = (V, E)$ be an arbitrary d -regular digraph. Throughout the argument we assume, whenever it is needed, that d is sufficiently large. Let p be a prime satisfying $10d^{1/2} \leq p \leq 20d^{1/2}$ (it is well known that for every n there is a prime between n and $2n$). By Lemma 5.5.5 there is a vertex coloring $f : V \rightarrow \{0, 1, \dots, p-1\}$ satisfying (5.7). For each $i, 0 \leq i < p$, let $G_i = (V, E_i)$ be the spanning subdigraph of G defined by $E_i = \{(u, v) \in E : f(v) \equiv f(u) + i \pmod{p}\}$. By inequality (5.7) the maximum indegree Δ_i^- and the maximum outdegree Δ_i^+ in each G_i is at most $(d/p) + 3\sqrt{d/p}\sqrt{\log d}$. Moreover, for each $i > 0$, the length of every directed cycle in G_i is divisible by p . Thus the directed girth g_i of G_i is at least p . Since each G_i can be completed, by adding vertices and edges, to a Δ_i -regular digraph with the same girth g_i and with $\Delta_i = \max(\Delta_i^+, \Delta_i^-)$, and since $g_i > 8e\Delta_i$ (for all sufficiently large d), we conclude, by Theorem 5.5.4, that $\text{dla}(G_i) \leq \Delta_i + 1 \leq (d/p) + 3\sqrt{d/p}\sqrt{\log d} + 1$ for all $1 \leq i < p$. For G_0 , we only apply the trivial inequality

$$\text{dla}(G_0) \leq 2\Delta_0 \leq 2\frac{d}{p} + 6\sqrt{\frac{d}{p}}\sqrt{\log d}$$

obtained, for example, by embedding G_0 as a subgraph of a Δ_0 -regular graph, splitting the edges of this graph into $\Delta_0/2$ -regular spanning subgraphs, and breaking each of these 1 -regular spanning subgraphs into two linear directed forests. The last two inequalities, together with the fact that $10\sqrt{d} \leq p \leq 20\sqrt{d}$, imply

$$\text{dla}(G) \leq d + 2\frac{d}{p} + 3\sqrt{pd}\sqrt{\log d} + 3\sqrt{\frac{d}{p}}\sqrt{\log d} + p - 1 \leq d + c \cdot d^{3/4}(\log d)^{1/2}.$$

We have thus proved the following.

Theorem 5.5.6 There is an absolute constant $c > 0$ such that for every d -regular digraph G

$$\text{dla}(G) \leq d + cd^{3/4}(\log d)^{1/2}$$

We note that by being a little more careful, we can improve the error term to $c'd^{2/3}(\log d)^{1/3}$. Since the edges of any undirected $d = 2f$ -regular graph can be oriented so that the resulting digraph is f -regular, and since any $(2f-1)$ -regular undirected graph is a subgraph of a $2f$ -regular graph, the last theorem implies the following.

Theorem 5.5.7 There is an absolute constant $c > 0$ such that for every undirected d -regular graph G

$$\text{la}(G) \leq \frac{d}{2} + cd^{3/4}(\log d)^{1/2}$$

17.5.6 5.6 Latin Transversals

Following the proof of the Local Lemma we noted that the mutual independency assumption in this lemma can be replaced by the weaker assumption that the conditional probability of each event, given the mutual nonoccurrence of an arbitrary set of events, each nonadjacent to it in the dependency digraph, is sufficiently small. In this section we describe an application, from Erdős and Spencer (1991), of this modified version of the lemma. Let $A = (a_{ij})$ be an $n \times n$ matrix with, say, integer entries. A permutation π is called a Latin transversal (of A) if the entries $a_{i\pi(i)}$ ($1 \leq i \leq n$) are all distinct.

Theorem 5.6.1 Suppose $k \leq (n-1)/(4e)$ and suppose that no integer appears in more than k entries of A . Then A has a Latin transversal.

Proof. Let π be a random permutation of $\{1, 2, \dots, n\}$, chosen according to a uniform distribution among all possible $n!$ permutations. Denote by T the set of all ordered four-tuples

(i, j, i', j') satisfying $i < i', j \neq j'$ and $a_{ij} = a_{i'j'}$. For each $(i, j, i', j') \in T$, let $A_{ijj'i'j'}$ denote the event that $\pi(i) = j$ and $\pi(i') = j'$. The existence of a Latin transversal is equivalent to the statement that with positive probability none of these events hold. Let us define a symmetric digraph (i.e., a graph) G on the vertex set T by making (i, j, i', j') adjacent to (p, q, p', q') if and only if $\{i, i'\} \cap \{p, p'\} \neq \emptyset$ or $\{j, j'\} \cap \{q, q'\} \neq \emptyset$. Thus these two four-tuples are not adjacent iff the four cells $(i, j), (i', j'), (p, q)$ and (p', q') occupy four distinct rows and columns of A . The maximum degree of G is less than $4nk$; indeed, for a given $(i, j, i', j') \in T$ there are at most $4n$ choices of (s, t) with either $s \in \{i, i'\}$ or $t \in \{j, j'\}$, and for each of these choices of (s, t) there are less than k choices for $(s', t') \neq (s, t)$ with $a_{st} = a_{s't'}$. Each such four-tuple (s, t, s', t') can be uniquely represented as (p, q, p', q') with $p < p'$. Since $e \cdot 4nk \cdot [1/n(n-1)] \leq 1$, the desired result follows from the above mentioned strengthening of the symmetric version of the Local Lemma, if we can show that

$$\Pr \left[A_{ijj'i'j'} \mid \bigwedge_S \overline{A_{pp'q'q'}} \right] \leq \frac{1}{n(n-1)} \quad (5.8)$$

for any $(i, j, i', j') \in T$ and any set S of members of T that are nonadjacent in G to (i, j, i', j') . By symmetry, we may assume that $i = j = 1, i' = j' = 2$ and that hence none of the p 's or q 's are either 1 or 2. Let us call a permutation π good if it satisfies $\bigwedge_S \overline{A_{pp'q'q'}}$, and let S_{ij} denote the set of all good permutations π satisfying $\pi(1) = i$ and $\pi(2) = j$. We claim that $|S_{12}| \leq |S_{ij}|$ for all $i \neq j$. Indeed, suppose first that $i, j > 2$. For each good $\pi \in S_{12}$ define a permutation π^* as follows. Suppose $\pi(x) = i, \pi(y) = j$. Then define $\pi^*(1) = i, \pi^*(2) = j, \pi^*(x) = 1, \pi^*(y) = 2$ and $\pi^*(t) = \pi(t)$ for all $t \neq 1, 2, x, y$. One can easily check that π^* is good, since the cells $(1, i), (2, j), (x, 1), (y, 2)$ are not part of any $(p, q, p', q') \in S$. Thus $\pi^* \in S_{ij}$, and since the mapping $\pi \rightarrow \pi^*$ is injective $|S_{12}| \leq |S_{ij}|$, as claimed. Similarly one can define injective mappings showing that $|S_{12}| \leq |S_{ij}|$ even when $\{i, j\} \cap \{1, 2\} \neq \emptyset$. It follows that

$$\Pr \left[A_{1122} \wedge \bigwedge_S \overline{A_{pp'q'q'}} \right] \leq \Pr \left[A_{1i2j} \wedge \bigwedge_S \overline{A_{pp'q'q'}} \right]$$

for all $i \neq j$ and hence that

$$\Pr \left[A_{1122} \mid \bigwedge_S \overline{A_{pp'q'q'}} \right] \leq \frac{1}{n(n-1)}$$

By symmetry, this implies (5.8) and completes the proof.

17.5.7 5.7 The Algorithmic Aspect

When the probabilistic method is applied to prove that a certain event holds with high probability, it often supplies an efficient deterministic, or at least randomized, algorithm for the corresponding problem.

By applying the Local Lemma we often manage to prove that a given event holds with positive probability, although this probability may be exponentially small in the dimensions of the problem. Consequently, it is not clear if any of these proofs can provide polynomial algorithms for the corresponding algorithmic problems. For many years there was no known method of converting the proofs of any of the examples discussed in this chapter into an efficient algorithm. In 1991 J. Beck found such a method that works for some of these examples, with a little loss in the constants.

Beck (1991) demonstrated his method by considering the problem of hypergraph two-coloring. For simplicity we only describe here the case of fixed edge-size in which each edge intersects a fixed number of other edges.

Let n, d be fixed positive integers. By the (n, d) problem we mean the following: Given sets $A_1, \dots, A_N \subseteq \Omega$ with all $|A_i| = n$, such that no set A_i intersects more than d other sets A_j , find a two-coloring of Ω so that no A_i is monochromatic. When $e(d+1) < 2^{n-1}$, Theorem 5.2.1 assures us that this problem always does have a solution. Can we find the coloring in polynomial (in N for fixed n, d) time? Beck has given an affirmative answer under somewhat more restrictive assumptions. We assume Ω is of the form $\Omega = \{1, \dots, m\}$, $m \leq Nn$ and the initial data structure consists of a list of the elements of the sets A_i and a list giving for each element j those i for which $j \in A_i$. We let G denote the dependency graph with vertices the sets A_i and A_i, A_j adjacent if they overlap.

Theorem 5.7.1 Let n, d be such that, setting $D = d(d-1)^3$, there exists a decomposition $n = n_1 + n_2 + n_3$ with

$$16D(1+d) < 2^{n_1}$$

$$16D(1+d) < 2^{n_2}$$

$$2e(1+d) < 2^{n_3}$$

Then there is a randomized algorithm with expected running time $O(N(\ln N)^c)$ for the (n, d) problem, where c is a constant (depending only on n and d).

For $\epsilon < 1/11$, fixed, we note that the above conditions are satisfied, for n sufficiently large, when $d < 2^{n\epsilon}$ by taking $n_1 = n_2 \sim 5n/11$ and $n_3 \sim n/11$. We emphasize again that the algorithmic analysis here is for fixed n, d and N approaching infinity, although the argument can be extended to the nonfixed case as well.

Beck has given a deterministic algorithm for the (n, d) problem. The randomized algorithm we give may be derandomized using the techniques of Chapter 16. The running time remains polynomial but seemingly no longer $N^{1+o(1)}$. Moreover, the algorithm can even be parallelized using some of the techniques in Chapter 16 together with a certain modification in the algorithm.

Proof. The First Pass. During this pass, points will be either red, blue, uncolored or saved. We move through the points $j \in \Omega$ sequentially, coloring them red or blue at random, flipping a fair coin. After each j is colored we check all $A_i \ni j$. If A_i now has n_1 points in one color and no points in the other color we call A_i dangerous. All uncolored $k \in A_i$ are now considered saved. When saved points k are reached in the sequential coloring they are not colored but simply skipped over. At the conclusion of the First Pass points are red, blue or saved. We say a set A_i survives if it does not have both red and blue points. Let $S \subseteq G$ denote the (random) set of surviving sets.

Claim 5.7.2 Almost surely all components C of $G|_S$ have size $O(\ln N)$.

Proof. An $A_i \in S$ may be dangerous or, possibly, many of its points were saved because neighboring (in G) sets were dangerous. The probability of a particular A_i becoming dangerous is at most 2^{1-n_1} since for this to occur the first n_1 coin flips determining colors of $j \in A_i$ must come up the same. (We only have inequality since in addition n_1 points of A_i must be reached before being saved.) Let V be an independent set in G ; that is, the $A_i \in V$ are mutually disjoint. Then the probability that all $A_i \in V$ become dangerous is at most $(2^{1-n_1})^{|V|}$ as the coin flips involve disjoint sets. Now let $V \subseteq G$ be such that all distances between the $A_i \in V$ are at least 4, distance being the length of the shortest path in G . We claim that

$$\Pr[V \subseteq S] \leq (d+1)^{|V|} (2^{1-n_1})^{|V|}$$

This is because for each $A_i \in V$ there are at most $d+1$ choices for a dangerous neighbor $A_{i'}$, giving $(d+1)^{|V|}$ choices for the $A_{i'}$. As the A_i are at least four apart the $A_{i'}$ cannot be adjacent and so the probability that they are all dangerous is at most $(2^{1-n_1})^{|V|}$, as claimed.

Call $T \subseteq G$ a 4-tree if the $A_i \in T$ are such that all their mutual distances in G are at least four and so that, drawing an arc between $A_i, A_j \in T$ if their distance is precisely four, the resulting graph is connected. We first bound the number of 4-trees of size u . The "distance-four" graph defined on T must contain a tree. There are less than 4^j trees (up to isomorphism) on j vertices, now fix one. We can label the tree $1, \dots, u$ so that each $j > 1$ is adjacent to some $i < j$. Now consider the number of $(A^1, \dots, A^u)$ whose distance-four graph corresponds to this tree. There are N choices for A^1 . Having chosen A^i for all $i < j$ the set A^j must be at distance four from A^i in G and there are at most D such points. Hence the number of 4-trees of size u is at most $4^u N D^{u-1} < N(4D)^u$. For any particular 4-tree T we have already that $\Pr[T \subseteq S] \leq [(d+1)2^{1-n_1}]^u$. Hence the expected number of 4-trees $T \subseteq S$ is at most

$$N [8D(d+1)2^{-n_1}]^u$$

As the bracketed term is less than $1/2$ by assumption, for $u = c_1 \ln N$ this term is $o(1)$. Thus almost surely $G|_S$ will contain no 4-tree of size bigger than $c_1 \ln N$. We actually want to bound the size of the components C of $G|_S$. A maximal 4-tree T in a component C must have the property that every $A_i \in C$ lies within three of an $A_j \in T$. There are less than d^3 (a constant) A_i within three of any given A_j so that $c_1 \ln N \geq |T| \geq |C|d^{-3}$ and so (since d is a constant)

$$|C| \leq c_2 \ln N,$$

proving the claim.

If the First Pass leaves components of size larger than $c_2 \ln N$ we simply repeat the entire procedure. In expected linear time the First Pass is successful. The points that are red or blue are now fixed. The sets A_i with both red and blue points can now be ignored. For each surviving A_i fix a subset B_i of $n - n_1$ saved points. It now suffices to color the saved points so that no B_i is monochromatic. B_i are split into components of size $O(\ln N)$ and it suffices to color each component separately. On the Second Pass we apply the method of the First Pass to each component of the B_i . Now we call a set B_i dangerous if it receives n_2 points of one color and none of the other. The Second Pass takes expected time $O(M)$ to color a component of size M ,

hence an expected time $O(N)$ to color all the components. (For success we require that a component of size M is broken into components of size at most $c_2 \ln M$. To avoid trivialities, if $M < \ln \ln N$ we skip the Second Pass for the corresponding component.) At the end of the Second Pass (still in linear time!) there is a family of twice surviving sets $C_i \subset B_i \subset A_i$ of size n_3 , the largest component of which has size $O(\ln \ln N)$.

We still need to color these $O(N)$ components of sets of size n_3 , each component of size $O(\ln \ln N)$. By the Local Lemma (or directly by Theorem 5.2.1), each of these components can be two-colored. We now find the two-coloring by brute force! Examining all two-colorings of a component of size M takes time $O(M2^{n_M})$, which is $O((\ln N)^c)$ in our case. Doing this for all components takes time $O(N(\ln N)^c)$. This completes the coloring.

We note that with slightly more restrictions on n, d , a Third Pass could be made and then the total time would be $O(N(\ln N)^c)$. We note also that a similar technique can be applied for converting several other applications of the Local Lemma into efficient algorithms.

17.5.8 5.8 Exercises

- (*) Prove that for every integer $d > 1$ there is a finite $c(d)$ such that the edges of any bipartite graph with maximum degree d in which every cycle has at least $c(d)$ edges can be colored by $d+1$ colors so that there are no two adjacent edges with the same color and there is no two-colored cycle.

2. (*) Prove that for every $\epsilon > 0$ there is a finite $l_0 = l_0(\epsilon)$ and an infinite sequence of bits $a_1, a_2, a_3, \dots$, where $a_i \in \{0, 1\}$, such that for every $l > l_0$ and every $i \geq 1$ the two binary vectors $u = (a_i, a_{i+1}, \dots, a_{i+l-1})$ and $v = (a_{i+l}, a_{i+l+1}, \dots, a_{i+2l-1})$ differ in at least $(\frac{1}{2} - \epsilon)l$ coordinates.
3. Let $G = (V, E)$ be a simple graph and suppose each $v \in V$ is associated with a set $S(v)$ of colors of size at least $10d$, where $d \geq 1$. Suppose, in addition, that for each $v \in V$ and $c \in S(v)$ there are at most d neighbors u of v such that c lies in $S(u)$. Prove that there is a proper coloring of G assigning to each vertex v a color from its class $S(v)$.
4. Let $G = (V, E)$ be a cycle of length $4n$ and let $V = V_1 \cup V_2 \cup \dots \cup V_n$ be a partition of its $4n$ vertices into n pairwise disjoint subsets, each of cardinality 4. Is it true that there must be an independent set of G containing precisely one vertex from each V_i ? (Prove or supply a counter example.)
5. (*) Prove that there is an absolute constant $c > 0$ such that for every k there is a set S_k of at least $ck \ln k$ integers, such that for every coloring of the integers by k colors there is an integer x for which the set $x + S$ does not intersect all color classes.

The Probabilistic Lens: Directed Cycles

Let $D = (V, E)$ be a simple directed graph with minimum outdegree δ and maximum indegree Δ .

Theorem 1 [Alon and Linial (1989)] If $e(\Delta\delta + 1)(1 - 1/k)^\delta < 1$ then D contains a (directed, simple) cycle of length $0 \pmod k$.

Proof. Clearly we may assume that every outdegree is precisely δ , since otherwise we can consider a subgraph of D with this property.

Let $f : V \rightarrow \{0, 1, \dots, k-1\}$ be a random coloring of V , obtained by choosing, for each $v \in V$, $f(v) \in \{0, \dots, k-1\}$ independently, according to a uniform distribution. For each $v \in V$, let A_v denote the event that there is no $u \in V$, with $(v, u) \in E$ and $f(u) \equiv f(v) + 1 \pmod k$. Clearly $\Pr[A_v] = (1 - 1/k)^\delta$. One can easily check that each event A_v is mutually independent of all the events A_u but those satisfying

$$N^+(v) \cap \left(\{u\} \cup \bigcup N^+(u) \right) \neq \emptyset,$$

where here $N^+(v) = \{w \in V : (v, w) \in E\}$. The number of such u 's is at most $\Delta\delta$ and hence, by our assumption and by the Local Lemma (Corollary 5.1.2), $\Pr[\bigwedge_{v \in V} A_v] > 0$. Thus there is an $f : V \rightarrow \{0, 1, \dots, k-1\}$ such that for every $v \in V$ there is a $u \in V$ with

$$(v, u) \in E \quad \text{and} \quad f(u) \equiv f(v) + 1 \pmod k. \quad (1)$$

Starting at an arbitrary $v = v_0 \in V$ and applying (1) repeatedly we obtain a sequence $v_0, v_1, v_2, \dots$ of vertices of D so that $(v_i, v_{i+1}) \in E$ and $f(v_{i+1}) \equiv f(v_i) + 1 \pmod k$ for all $i \geq 0$. Let j be the minimum integer so that there is an $\ell < j$ with $v_\ell = v_j$. The cycle $v_\ell v_{\ell+1} v_{\ell+2} \dots v_j = v_\ell$ is a directed simple cycle of D whose length is divisible by k .

17.6 6 Correlation Inequalities

Let $G = (V, E)$ be a random graph on the set of vertices $V = \{1, 2, \dots, n\}$ generated by choosing, for each $i, j \in V, i \neq j$ independently, the pair $\{i, j\}$ to be an edge with probability p , where $0 < p < 1$. Let H be the event that G is Hamiltonian and let P be the event that G is planar. Suppose one wants to compare the two quantities $\Pr[P \wedge H]$ and $\Pr[P] \cdot \Pr[H]$.

Intuitively, knowing that G is Hamiltonian suggests that it has many edges and hence seems to indicate that G is less likely to be planar. Therefore it seems natural to expect that $\Pr[P \mid H] \leq \Pr[P]$ implying

$$\Pr[P \wedge H] \leq \Pr[H] \cdot \Pr[P].$$

This inequality, which is, indeed, correct, is a special case of the FKG inequality of Fortuin, Kasteleyn and Ginibre (1971). In this chapter we present the proof of this inequality and several related results, which deal with the correlation between certain events in probability spaces. The proofs of all these results are rather simple, and still they supply many interesting consequences. The first inequality of this type is due to Harris (1960). A result closer to the ones considered here is a lemma of Kleitman (1966a), stating that if $\mathcal{A}$ and $\mathcal{B}$ are two monotone decreasing families of subsets of $\{1, 2, \dots, n\}$ (i.e., $A \in \mathcal{A}$ and $A' \subseteq A \Rightarrow A' \in \mathcal{A}$ and, similarly $B \in \mathcal{B}$ and $B' \subseteq B \Rightarrow B' \in \mathcal{B}$) then

$$|\mathcal{A} \cap \mathcal{B}| \cdot 2^n \geq |\mathcal{A}| \cdot |\mathcal{B}|.$$

This lemma was followed by many extensions and generalizations until Ahlswede and Daykin (1978) obtained a very general result, which implies all these extensions. In the next section we present this result and its proof. Some of its many applications are discussed in the rest of the chapter.

17.6.1 6.1 The Four Functions Theorem of Ahlswede and Daykin

Suppose $n \geq 1$ and put $N = \{1, 2, \dots, n\}$. Let $P(N)$ denote the set of all subsets of N , and let $\mathbb{R}^+$ denote the set of nonnegative real numbers. For a function $\varphi : P(N) \rightarrow \mathbb{R}^+$ and for a family $\mathcal{A}$ of subsets of N denote $\varphi(\mathcal{A}) = \sum_{A \in \mathcal{A}} \varphi(A)$. For two families $\mathcal{A}$ and $\mathcal{B}$ of subsets of N define $\mathcal{A} \cup \mathcal{B} = \{A \cup B : A \in \mathcal{A}, B \in \mathcal{B}\}$ and $\mathcal{A} \cap \mathcal{B} = \{A \cap B : A \in \mathcal{A}, B \in \mathcal{B}\}$.

Theorem 6.1.1 [The Four Functions Theorem] Let $\alpha, \beta, \gamma, \delta : P(N) \rightarrow \mathbb{R}^+$ be four functions from the set of all subsets of N to the nonnegative reals. If, for every two subsets $A, B \subseteq N$ the inequality

$$\alpha(A)\beta(B) \leq \gamma(A \cup B)\delta(A \cap B) \tag{6.1}$$

holds, then, for every two families of subsets $\mathcal{A}, \mathcal{B} \subseteq P(N)$,

$$\alpha(\mathcal{A})\beta(\mathcal{B}) \leq \gamma(\mathcal{A} \cup \mathcal{B})\delta(\mathcal{A} \cap \mathcal{B}). \tag{6.2}$$

Proof. Observe, first, that we may modify the four functions $\alpha, \beta, \gamma, \delta$ by defining $\alpha(A) = 0$ for all $A \notin \mathcal{A}$, $\beta(B) = 0$ for all $B \notin \mathcal{B}$, $\gamma(C) = 0$ for all $C \notin \mathcal{A} \cup \mathcal{B}$, and $\delta(D) = 0$ for all $D \notin \mathcal{A} \cap \mathcal{B}$. Clearly (6.1) still holds for the modified functions and in inequality (6.2) we may assume now that $\mathcal{A} = \mathcal{B} = \mathcal{A} \cup \mathcal{B} = \mathcal{A} \cap \mathcal{B} = P(N)$.

To prove this inequality we apply induction on n . The only step that requires some computation is $n = 1$. In this case $P(N) = \{\emptyset, N\}$. For each function $\varphi \in \{\alpha, \beta, \gamma, \delta\}$ define $\varphi_0 = \varphi(\emptyset)$ and $\varphi_1 = \varphi(N)$. By (6.1) we have

$$\begin{aligned} \alpha_0\beta_0 &\leq \gamma_0\delta_0, \\ \alpha_0\beta_1 &\leq \gamma_1\delta_0, \\ \alpha_1\beta_0 &\leq \gamma_1\delta_0, \\ \alpha_1\beta_1 &\leq \gamma_1\delta_1. \end{aligned} \tag{6.3}$$

By the above paragraph we only have to prove inequality (6.2), where $\mathcal{A} = \mathcal{B} = P(N)$, that is, to prove that

$$(\alpha_0 + \alpha_1)(\beta_0 + \beta_1) \leq (\gamma_0 + \gamma_1)(\delta_0 + \delta_1). \quad (6.4)$$

If either $\gamma_1 = 0$ or $\delta_0 = 0$ this follows immediately from (6.3). Otherwise, by (6.3), $\gamma_0 \geq \alpha_0\beta_0/\delta_0$ and $\delta_1 \geq \alpha_1\beta_1/\gamma_1$. It thus suffices to show that

$$\left(\frac{\alpha_0\beta_0}{\delta_0} + \gamma_1\right) \left(\delta_0 + \frac{\alpha_1\beta_1}{\gamma_1}\right) \geq (\alpha_0 + \alpha_1)(\beta_0 + \beta_1),$$

or, equivalently, that

$$(\alpha_0\beta_0 + \gamma_1\delta_0)(\delta_0\gamma_1 + \alpha_1\beta_1) \geq (\alpha_0 + \alpha_1)(\beta_0 + \beta_1)\delta_0\gamma_1.$$

The last inequality is equivalent to

$$(\gamma_1\delta_0 - \alpha_0\beta_1)(\gamma_1\delta_0 - \alpha_1\beta_0) \geq 0,$$

which follows from (6.3), as both factors on the left-hand side are nonnegative. This completes the proof for $n = 1$.

Suppose, now, that the theorem holds for $n - 1$ and let us prove it for $n (\geq 2)$. Put $N' = N \setminus \{n\}$ and define for each $\varphi \in \{\alpha, \beta, \gamma, \delta\}$ and each $A \subseteq N'$, $\varphi'(A) = \varphi(A) + \varphi(A \cup \{n\})$. Clearly, for each function $\varphi \in \{\alpha, \beta, \gamma, \delta\}$, $\varphi'(P(N')) = \varphi(P(N))$. Therefore the desired inequality (6.3) would follow from applying the induction hypothesis to the functions $\alpha', \beta', \gamma', \delta' : P(N') \rightarrow \mathbb{R}^+$. However, in order to apply this hypothesis we have to check that these new functions satisfy the assumption of Theorem 6.1.1 on N' ; that is, that for every $A', B' \subseteq N'$,

$$\alpha'(A')\beta'(B') \leq \gamma'(A' \cup B')\delta'(A' \cap B'). \quad (6.5)$$

Not surprisingly, this last inequality follows easily from the case $n = 1$, which we have already proved. Indeed, let T be a 1-element set and define

$$\begin{aligned} \bar{\alpha}(\emptyset) &= \alpha(A'), & \bar{\alpha}(T) &= \alpha(A' \cup \{n\}), \\ \bar{\beta}(\emptyset) &= \beta(B'), & \bar{\beta}(T) &= \beta(B' \cup \{n\}), \\ \bar{\gamma}(\emptyset) &= \gamma(A' \cup B'), & \bar{\gamma}(T) &= \gamma(A' \cup B' \cup \{n\}), \\ \bar{\delta}(\emptyset) &= \delta(A' \cap B'), & \bar{\delta}(T) &= \delta((A' \cap B') \cup \{n\}). \end{aligned}$$

By the assumption (6.1), $\bar{\alpha}(S)\bar{\beta}(R) \leq \bar{\gamma}(S \cup R)\bar{\delta}(S \cap R)$ for all $S, R \subseteq T$ and hence, by the case $n = 1$ already proved,

$$\alpha'(A')\beta'(B') = \bar{\alpha}(P(T))\bar{\beta}(P(T)) \leq \bar{\gamma}(P(T))\bar{\delta}(P(T)) = \gamma'(A' \cup B')\delta'(A' \cap B'),$$

which is the desired inequality (6.5). Therefore inequality (6.2) holds, completing the proof. $\square$

The Ahlswede-Daykin Theorem can be extended to arbitrary finite distributive lattices. A lattice is a partially ordered set in which every two elements, x and y , have a unique minimal upper bound, denoted by $x \vee y$ and called the join of x and y and a unique maximal lower bound, denoted by $x \wedge y$ and called the meet of x and y . A lattice L is distributive if for all $x, y, z \in L$,

$$x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$$

or, equivalently if for all $x, y, z \in L$,

$$x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z)$$

For two sets $X, Y \subseteq L$ define

$$X \vee Y = \{x \vee y : x \in X, y \in Y\}$$

and

$$X \wedge Y = \{x \wedge y : x \in X, y \in Y\}$$

Any subset L of $P(N)$, where $N = \{1, 2, \dots, n\}$, ordered by inclusion, which is closed under the union and intersection operations is a distributive lattice. Here, the join of two members $A, B \in L$ is simply their union $A \cup B$ and their meet is the intersection $A \cap B$. It is somewhat more surprising (but easy to check) that every finite distributive lattice L is isomorphic to a sublattice of $P(\{1, 2, \dots, n\})$ for some n . [To see this, call an element $x \in L$ join-irreducible if whenever $x = y \vee z$ then either $x = y$ or $x = z$. Let $x_1, x_2, \dots, x_n$ be the set of all join-irreducible elements in L and associate each element $x \in L$ with the set $A = A(x) \subseteq N$, where $x = \bigvee_{i \in A} x_i$ and $\{x_i : i \in A\}$ are all the join-irreducibles y satisfying $y \leq x$. The mapping $x \rightarrow A(x)$ is the desired isomorphism.] This fact enables us to generalize Theorem 6.1.1 to arbitrary finite distributive lattices as follows.

Corollary 6.1.2 Let L be a finite distributive lattice and let α, β, γ and δ be four functions from L to $\mathbb{R}^+$. If

$$\alpha(x)\beta(y) \leq \gamma(x \vee y)\delta(x \wedge y)$$

for all $x, y \in L$ then for every $X, Y \subseteq L$,

$$\alpha(X)\beta(Y) \leq \gamma(X \vee Y)\delta(X \wedge Y)$$

The simplest case in the last corollary is the case where all the four functions α, β, γ and δ are identically 1, stated below.

Corollary 6.1.3 Let L be a finite distributive lattice and suppose $X, Y \subseteq L$. Then

$$|X| \cdot |Y| \leq |X \vee Y| \cdot |X \wedge Y|.$$

We close this section by presenting a very simple consequence of the last corollary, first proved by Marica and Schonheim (1969).

Corollary 6.1.4 Let $\mathcal{A}$ be a family of subsets of a finite set N and define

$$\mathcal{A} \setminus \mathcal{A} = \{F \setminus F' : F, F' \in \mathcal{A}\}$$

Then $|\mathcal{A} \setminus \mathcal{A}| \geq |\mathcal{A}|$.

Proof. Let L be the distributive lattice of all subsets of N . By applying Corollary 6.1.3 to $\mathcal{A}$ and $\mathcal{B} = \{N \setminus F : F \in \mathcal{A}\}$ we obtain

$$|\mathcal{A}|^2 = |\mathcal{A}| \cdot |\mathcal{B}| \leq |\mathcal{A} \cup \mathcal{B}| \cdot |\mathcal{A} \cap \mathcal{B}| = |\mathcal{A} \setminus \mathcal{A}|^2.$$

The desired result follows.

17.6.2 6.2 The Fkg Inequality

A function $\mu : L \rightarrow \mathbb{R}^+$, where L is a finite distributive lattice, is called logsupermodular if

$$\mu(x)\mu(y) \leq \mu(x \vee y)\mu(x \wedge y)$$

for all $x, y \in L$. A function $f : L \rightarrow \mathbb{R}^+$ is increasing if $f(x) \leq f(y)$ whenever $x \leq y$ and is decreasing if $f(x) \geq f(y)$ whenever $x \leq y$.

Motivated by a problem from statistical mechanics, Fortuin et al. (1971) proved the following useful inequality, which has become known as the FKG Inequality.

Theorem 6.2.1 [The FKG Inequality] Let L be a finite distributive lattice and let $\mu : L \rightarrow \mathbb{R}^+$ be a log-supermodular function. Then, for any two increasing functions $f, g : L \rightarrow \mathbb{R}^+$ we have

$$\left(\sum_{x \in L} \mu(x) f(x) \right) \left(\sum_{x \in L} \mu(x) g(x) \right) \leq \left(\sum_{x \in L} \mu(x) f(x) g(x) \right) \left(\sum_{x \in L} \mu(x) \right). \quad (6.6)$$

Proof. Define four functions $\alpha, \beta, \gamma, \delta : L \rightarrow \mathbb{R}^+$ as follows. For each $x \in L$,

$$\begin{aligned} \alpha(x) &= \mu(x) f(x), & \beta(x) &= \mu(x) g(x), \\ \gamma(x) &= \mu(x) f(x) g(x), & \delta(x) &= \mu(x). \end{aligned}$$

We claim that these functions satisfy the hypothesis of the Ahlswede-Daykin Theorem, stated in Corollary 6.1.2. Indeed, if $x, y \in L$ then, by the supermodularity of μ and since f and g are increasing,

$$\begin{aligned} \alpha(x) \beta(y) &= \mu(x) f(x) \mu(y) g(y) \leq \mu(x \vee y) f(x) g(y) \mu(x \wedge y) \\ &\leq \mu(x \vee y) f(x \vee y) g(x \vee y) \mu(x \wedge y) = \gamma(x \vee y) \delta(x \wedge y) \end{aligned}$$

Therefore by Corollary 6.1.2 (with $X = Y = L$),

$$\alpha(L) \beta(L) \leq \gamma(L) \delta(L)$$

which is the desired result.

Note that the conclusion of Theorem 6.2.1 holds also if both f and g are decreasing (simply interchange γ and δ in the proof). In case f is increasing and g is decreasing (or vice versa) the opposite inequality holds:

$$\left(\sum_{x \in L} \mu(x) f(x) \right) \left(\sum_{x \in L} \mu(x) g(x) \right) \geq \left(\sum_{x \in L} \mu(x) f(x) g(x) \right) \left(\sum_{x \in L} \mu(x) \right).$$

To prove it, simply apply Theorem 6.2.1 to the two increasing functions $f(x)$ and $k - g(x)$, where k is the constant $\max_{x \in L} g(x)$. [This constant is needed to guarantee that $k - g(x) \geq 0$ for all $x \in L$.]

It is helpful to view μ as a measure on L . Assuming μ is not identically zero we can define, for any function $f : L \rightarrow \mathbb{R}^+$, its expectation,

$$\langle f \rangle = \frac{\sum_{x \in L} f(x) \mu(x)}{\sum_{x \in L} \mu(x)}.$$

With this notation, the FKG Inequality asserts that if μ is log-supermodular and $f, g : L \rightarrow \mathbb{R}^+$ are both increasing or both decreasing then

$$\langle fg \rangle \geq \langle f \rangle \langle g \rangle$$

Similarly, if f is increasing and g is decreasing (or vice versa), then

$$\langle fg \rangle \leq \langle f \rangle \langle g \rangle$$

This formulation demonstrates clearly the probabilistic nature of the inequality, some of whose many interesting consequences are presented in the rest of this chapter.

17.6.3 6.3 Monotone Properties

Recall that a family $\mathcal{A}$ of subsets of $N = \{1, 2, \dots, n\}$ is monotone decreasing if $A \in \mathcal{A}$ and $A' \subseteq A \Rightarrow A' \in \mathcal{A}$. Similarly, it is monotone increasing if $A \in \mathcal{A}$ and $A \subseteq A' \Rightarrow A' \in \mathcal{A}$. By considering the power set $P(N)$ as a symmetric probability space, one naturally defines the probability of $\mathcal{A}$ by

$$\Pr[\mathcal{A}] = \frac{|\mathcal{A}|}{2^n}.$$

Thus $\Pr[\mathcal{A}]$ is simply the probability that a randomly chosen subset of N lies in $\mathcal{A}$. Kleitman's Lemma, which was the starting point of all the correlation inequalities considered in this chapter, is the following.

Proposition 6.3.1 Let $\mathcal{A}$ and $\mathcal{B}$ be two monotone increasing families of subsets of $N = \{1, 2, \dots, n\}$ and let $\mathcal{C}$ and $\mathcal{D}$ be two monotone decreasing families of subsets of N . Then

$$\Pr[\mathcal{A} \cap \mathcal{B}] \geq \Pr[\mathcal{A}] \cdot \Pr[\mathcal{B}],$$

$$\Pr[\mathcal{C} \cap \mathcal{D}] \geq \Pr[\mathcal{C}] \cdot \Pr[\mathcal{D}],$$

$$\Pr[\mathcal{A} \cap \mathcal{C}] \leq \Pr[\mathcal{A}] \cdot \Pr[\mathcal{C}].$$

In terms of cardinalities, this can be read as follows:

$$2^n |\mathcal{A} \cap \mathcal{B}| \geq |\mathcal{A}| \cdot |\mathcal{B}|,$$

$$2^n |\mathcal{C} \cap \mathcal{D}| \geq |\mathcal{C}| \cdot |\mathcal{D}|,$$

$$2^n |\mathcal{A} \cap \mathcal{C}| \leq |\mathcal{A}| \cdot |\mathcal{C}|,$$

where here and in what follows, $\mathcal{A} \cap \mathcal{B}$, $\mathcal{C} \cap \mathcal{D}$ and $\mathcal{A} \cap \mathcal{C}$ denote usual intersections of families.

Proof. Let $f : P(N) \rightarrow \mathbb{R}^+$ be the characteristic function of $\mathcal{A}$; that is, $f(A) = 0$ if $A \notin \mathcal{A}$ and $f(A) = 1$ if $A \in \mathcal{A}$. Similarly, let g be the characteristic function of $\mathcal{B}$. By the assumptions, f and g are both increasing. Applying the FKG Inequality with the trivial measure $\mu \equiv 1$ we get

$$\Pr[\mathcal{A} \cap \mathcal{B}] = \langle fg \rangle \geq \langle f \rangle \langle g \rangle = \Pr[\mathcal{A}] \cdot \Pr[\mathcal{B}]$$

The other two inequalities follow similarly from Theorem 6.2.1 and the paragraph following it.

It is worth noting that this proposition can be also derived easily from the Ahlswede-Daykin Theorem or from Corollary 6.1.3.

The last proposition has several interesting combinatorial consequences, some of which appear already in Kleitman's original paper. Since those are direct combinatorial consequences and do not contain any additional probabilistic ideas, we omit their exact statement and turn to a version of Proposition 6.3.1 in a more general probability space.

For a real vector $p = (p_1, \dots, p_n)$, where $0 \leq p_i \leq 1$, consider the probability space whose elements are all members of the power set $P(N)$, where, for each $A \subseteq N$, $\Pr[A] = \prod_{i \in A} p_i \prod_{j \notin A} (1 - p_j)$. Clearly this probability distribution is obtained if we choose a random $A \subseteq N$ by choosing each element $i \in N$, independently, with probability p_i . Let us denote, for each $\mathcal{A} \subseteq P(N)$, its probability in this space by $\Pr_p[\mathcal{A}]$. In particular, if all the probabilities p_i are $1/2$ then $\Pr_p[\mathcal{A}]$ is the quantity denoted as $\Pr[\mathcal{A}]$ in Proposition 6.3.1. Define $\mu = \mu_p : P(N) \rightarrow \mathbb{R}^+$ by $\mu(A) = \prod_{i \in A} p_i \prod_{j \notin A} (1 - p_j)$.

It is easy to check that μ is log-supermodular. This is because for $A, B \subseteq N$, $\mu(A)\mu(B) = \mu(A \cup B)\mu(A \cap B)$, as can be checked by comparing the contribution arising from each $i \in N$ to the left-hand side and to the right-hand side of the last equality. Hence one can apply the FKG Inequality and obtain the following generalization of Proposition 6.3.1.

Theorem 6.3.2 Let $\mathcal{A}$ and $\mathcal{B}$ be two monotone increasing families of subsets of N and let $\mathcal{C}$ and $\mathcal{D}$ be two monotone decreasing families of subsets of N . Then, for any real vector $p = (p_1, \dots, p_n)$, $0 \leq p_i \leq 1$,

$$\begin{aligned}\Pr_p[\mathcal{A} \cap \mathcal{B}] &\geq \Pr_p[\mathcal{A}] \cdot \Pr_p[\mathcal{B}], \\ \Pr_p[\mathcal{C} \cap \mathcal{D}] &\geq \Pr_p[\mathcal{C}] \cdot \Pr_p[\mathcal{D}], \\ \Pr_p[\mathcal{A} \cap \mathcal{C}] &\leq \Pr_p[\mathcal{A}] \cdot \Pr_p[\mathcal{C}].\end{aligned}$$

This theorem can be applied in many cases and will be used in Chapter 8 to derive the Janson Inequalities. As a simple illustration suppose that $A_1, A_2, \dots, A_k$ are arbitrary subsets of N and one chooses a random subset A of N by choosing each $i \in N$, independently, with probability p . Then, Theorem 6.3.2 easily implies that

$$\Pr[A \text{ intersects each } A_i] \geq \prod_{i=1}^k \Pr[A \text{ intersects } A_i]$$

Note that this is false, in general, for other similar probabilistic models. For example, if A is a randomly chosen ℓ -element subset of N then the last inequality may fail.

By viewing the members of N as the $n = \binom{m}{2}$ edges of the complete graph on the set of vertices $V = \{1, 2, \dots, m\}$, we can derive a correlation inequality for random graphs. Let $G = (V, E)$ be a random graph on the set of vertices V generated by choosing, for each $i, j \in V, i \neq j$, independently, the pair $\{i, j\}$ to be an edge with probability p . (This model of random graphs is discussed in detail in Chapter 10.) A property of graphs is a subset of the set of all graphs on V , closed under isomorphism. Thus, for example, connectivity is a property (corresponding to all connected graphs on V) and planarity is another property. A property Q is monotone increasing if whenever G has Q and H is obtained from G by adding edges then H has Q too. A monotone decreasing property is defined in a similar manner. By interpreting the members of N in Theorem 6.3.2 as the $\binom{m}{2}$ pairs $\{i, j\}$ with $i, j \in V, i \neq j$ we obtain the following.

Theorem 6.3.3 Let Q_1, Q_2, Q_3 and Q_4 be graph properties, where Q_1, Q_2 are monotone increasing and Q_3, Q_4 are monotone decreasing. Let $G = (V, E)$ be a random graph on V obtained by picking every edge, independently, with probability p . Then

$$\begin{aligned}\Pr[G \in Q_1 \cap Q_2] &\geq \Pr[G \in Q_1] \cdot \Pr[G \in Q_2], \\ \Pr[G \in Q_3 \cap Q_4] &\geq \Pr[G \in Q_3] \cdot \Pr[G \in Q_4], \\ \Pr[G \in Q_1 \cap Q_3] &\leq \Pr[G \in Q_1] \cdot \Pr[G \in Q_3].\end{aligned}$$

Thus, for example, the probability that G is both Hamiltonian and planar does not exceed the product of the probability that it is Hamiltonian by that it is planar. It seems hopeless to try and prove such a statement directly, without using one of the correlation inequalities.

17.6.4 6.4 Linear Extensions of Partially Ordered Sets

Let $(P, \leq)$ be a partially ordered set with n elements. A linear extension of P is a one to one mapping $\sigma : P \rightarrow \{1, 2, \dots, n\}$, which is order preserving; that is, if $x, y \in P$ and $x \leq y$ then $\sigma(x) \leq \sigma(y)$. Intuitively, σ is a ranking of the elements of P that preserves the partial order of P . Consider the probability space of all linear extensions of P , where each possible extension is equally likely. In this space we can consider events of the form, for example, $x \leq y$ or $(x \leq y) \wedge (x \leq z)$ (for $x, y, z \in P$) and compute their probabilities. It turns out that the FKG Inequality is a very useful tool for studying the correlation between such events. The best known result of this

form was conjectured by Rival and Sands and proved by Shepp (1982). [See also Fishburn

(1992) for a strengthening.] It asserts that for any partially ordered set P and any three elements $x, y, z \in P$: $\Pr[x \leq y \wedge x \leq z] \geq \Pr[x \leq y] \Pr[x \leq z]$.

This result became known as the *XYZ Theorem*. Although it looks intuitively obvious, its proof is nontrivial and contains a clever application of the FKG Inequality. In this section we present this result and its elegant proof.

Theorem 6.4.1 Let P be a partially ordered set with n elements $a_1, a_2, \dots, a_n$. Then

$$\Pr[a_1 \leq a_2 \wedge a_1 \leq a_3] \geq \Pr[a_1 \leq a_2] \Pr[a_1 \leq a_3]$$

Proof. Let m be a large integer (which will later tend to infinity) and let L be the set of all ordered n -tuples $\mathbf{x} = (x_1, \dots, x_n)$, where $x_i \in M = \{1, 2, \dots, m\}$. (Note that we do not assume that the numbers x_i are distinct.) Define an order relation $\leq$ on L as follows. For $\mathbf{y} = (y_1, \dots, y_n) \in L$ and $\mathbf{x}$ as above $\mathbf{x} \leq \mathbf{y}$ iff $x_1 \geq y_1$ and $x_i - x_1 \leq y_i - y_1$ for all $2 \leq i \leq n$. It is not too difficult to check that $(L, \leq)$ is a lattice in which the i th component of the meet $\mathbf{x} \wedge \mathbf{y}$ is $(\mathbf{x} \wedge \mathbf{y})_i = \min\{x_i - x_1, y_i - y_1\} + \max\{x_1, y_1\}$ and the i th component of the join $\mathbf{x} \vee \mathbf{y}$ is $(\mathbf{x} \vee \mathbf{y})_i = \max\{x_i - x_1, y_i - y_1\} + \min\{x_1, y_1\}$.

Moreover, the lattice L is distributive. This follows by an easy computation from the fact that the trivial lattice of integers (with respect to the usual order) is distributive and hence for any three integers a, b and c ,

$$\min\{a, \max\{b, c\}\} = \max\{\min\{a, b\}, \min\{a, c\}\}, \quad (6.7)$$

and

$$\max\{a, \min\{b, c\}\} = \min\{\max\{a, b\}, \max\{a, c\}\}. \quad (6.8)$$

Let us show how this implies that L is distributive. Let $\mathbf{x} = (x_1, \dots, x_n)$, $\mathbf{y} = (y_1, \dots, y_n)$ and $\mathbf{z} = (z_1, \dots, z_n)$ be three elements of L . We must show that

$$\mathbf{x} \wedge (\mathbf{y} \vee \mathbf{z}) = (\mathbf{x} \wedge \mathbf{y}) \vee (\mathbf{x} \wedge \mathbf{z})$$

The i th component of $\mathbf{x} \wedge (\mathbf{y} \vee \mathbf{z})$ is

$$\begin{aligned} (\mathbf{x} \wedge (\mathbf{y} \vee \mathbf{z}))_i &= \min\{x_i - x_1, (\mathbf{y} \vee \mathbf{z})_i - (\mathbf{y} \vee \mathbf{z})_1\} \\ &\quad + \max\{x_1, (\mathbf{y} \vee \mathbf{z})_1\} \\ &= \min\{x_i - x_1, \max\{y_i - y_1, z_i - z_1\}\} \\ &\quad + \max\{x_1, \min\{y_1, z_1\}\} \end{aligned}$$

Similarly, the i th component of $(\mathbf{x} \wedge \mathbf{y}) \vee (\mathbf{x} \wedge \mathbf{z})$ is

$$\begin{aligned} &((\mathbf{x} \wedge \mathbf{y}) \vee (\mathbf{x} \wedge \mathbf{z}))_i \\ &= \max\{(\mathbf{x} \wedge \mathbf{y})_i - (\mathbf{x} \wedge \mathbf{y})_1, (\mathbf{x} \wedge \mathbf{z})_i - (\mathbf{x} \wedge \mathbf{z})_1\} \\ &\quad + \min\{(\mathbf{x} \wedge \mathbf{y})_1, (\mathbf{x} \wedge \mathbf{z})_1\} \\ &= \max\{\min\{x_i - x_1, y_i - y_1\}, \min\{x_i - x_1, z_i - z_1\}\} \\ &\quad + \min\{\max\{x_1, y_1\}, \max\{x_1, z_1\}\}. \end{aligned}$$

These two quantities are equal, as follows by applying (6.7) with $a = x_i - x_1$, $b = y_i - y_1$, $c = z_i - z_1$ and (6.8) with $a = x_1$, $b = y_1$, $c = z_1$.

Thus L is distributive. To apply the FKG Inequality we need the measure function μ and the two functions f and g . Let μ be the characteristic function of P ; that is, for $\mathbf{x} = (x_1, \dots, x_n) \in L$, $\mu(\mathbf{x}) = 1$ if $x_i \leq x_j$ whenever $a_i \leq a_j$ in P and $\mu(\mathbf{x}) = 0$ otherwise. To show that μ is log-supermodular it suffices to check that if $\mu(\mathbf{x}) = \mu(\mathbf{y}) = 1$ then $\mu(\mathbf{x} \vee \mathbf{y}) = \mu(\mathbf{x} \wedge \mathbf{y}) = 1$. However, if $\mu(\mathbf{x}) = \mu(\mathbf{y}) = 1$ and $a_i \leq a_j$ in P then $x_i \leq x_j$ and $y_i \leq y_j$ and hence

$$\begin{aligned}(x \vee y)_i &= \max \{x_i - x_1, y_i - y_1\} + \min \{x_1, y_1\} \\ &\leq \max \{x_j - x_1, y_j - y_1\} + \min \{x_1, y_1\} = (x \vee y)_j,\end{aligned}$$

that is, $\mu(\mathbf{x} \vee \mathbf{y}) = 1$. Similarly, $\mu(\mathbf{x}) = \mu(\mathbf{y}) = 1$ implies $\mu(\mathbf{x} \wedge \mathbf{y}) = 1$ too.

Not surprisingly, we define the functions f and g as the characteristic functions of the two events $x_1 \leq x_2$ and $x_1 \leq x_3$, respectively; that is, $f(\mathbf{x}) = 1$ if $x_1 \leq x_2$ and $f(\mathbf{x}) = 0$ otherwise, and $g(\mathbf{x}) = 1$ if $x_1 \leq x_3$ and $g(\mathbf{x}) = 0$ otherwise. Trivially, both f and g are increasing. Indeed, if $\mathbf{x} \leq \mathbf{y}$ and $f(\mathbf{x}) = 1$ then $0 \leq x_2 - x_1 \leq y_2 - y_1$ and hence $f(\mathbf{y}) = 1$, and similarly for g .

We therefore have all the necessary ingredients for applying the FKG Inequality (Theorem 6.2.1). This gives that in L the probability that an n -tuple $(x_1, \dots, x_n)$ that satisfies the inequalities in P satisfies both $x_1 \leq x_2$ and $x_1 \leq x_3$ is at least as big as the product of the probability that it satisfies $x_1 \leq x_2$ by that it satisfies $x_1 \leq x_3$. Note that this is not yet what we wanted to prove; the n -tuples in L are not n -tuples of distinct integers and thus do not correspond to linear extensions of P . However, as $m \rightarrow \infty$, the probability that $x_i = x_j$ for some $i \neq j$ in a member $\mathbf{x} = (x_1, \dots, x_n)$ of L tends to 0 and the assertion of the theorem follows.

17.6.5 6.5 Exercises

1. Let G be a graph and let P denote the probability that a random subgraph of G obtained by picking each edge of G with probability $1/2$, independently, is connected (and spanning). Let Q denote the probability that in a random two-coloring of G , where each edge is chosen, randomly and independently, to be either red or blue, the red graph and the blue graph are both connected (and spanning). Is $Q \leq P^2$?
2. A family of subsets $\mathcal{G}$ is called intersecting if $G_1 \cap G_2 \neq \emptyset$ for all $G_1, G_2 \in \mathcal{G}$. Let $\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k$ be k intersecting families of subsets of $\{1, 2, \dots, n\}$. Prove that

$$\left| \bigcup_{i=1}^k \mathcal{F}_i \right| \leq 2^n - 2^{n-k}$$

3. Show that the probability that in the random graph $G(2k, 1/2)$ the maximum degree is at most $k-1$ is at least $1/4^k$.

The Probabilistic Lens: Turán's Theorem

In a graph $G = (V, E)$ let d_v denote the degree of a vertex v and let $\alpha(G)$ be the maximal size of an independent set of vertices. The following result was proved by Caro and Wei.

Theorem $1\alpha(G) \geq \sum_{v \in V} \frac{1}{d_v + 1}$.

Proof. Let $<$ be a uniformly chosen total ordering of V . Define

$$I = \{v \in V : \{v, w\} \in E \Rightarrow v < w\}.$$

Let X_v be the indicator random variable for $v \in I$ and $X = \sum_{v \in V} X_v = |I|$. For each v ,

$$\mathbb{E}[X_v] = \Pr[v \in I] = \frac{1}{d_v + 1}$$

since $v \in I$ if and only if v is the least element among v and its neighbors. Hence

$$\mathbb{E}[X] = \sum_{v \in V} \frac{1}{d_v + 1}$$

and so there exists a specific ordering $<$ with

$$|I| \geq \sum_{v \in V} \frac{1}{d_v + 1}$$

But if $x, y \in I$ and $\{x, y\} \in E$ then $x < y$ and $y < x$, a contradiction. Thus I is independent and $\alpha(G) \geq |I|$.

For any $m \leq n$ let q, r satisfy $n = mq + r, 0 \leq r < m$, and let $e = r \binom{q+1}{2} + (m-r) \binom{q}{2}$. Define a graph $G = G_{n,e}$ on n vertices and e edges by splitting the vertex set into m classes as evenly as possible and joining two vertices if and only if they lie in the same class. Clearly $\alpha(G_{n,e}) = m$.

Theorem 2 [Turán (1941)] Let H have n vertices and e edges. Then $\alpha(H) \geq m$ and $\alpha(H) = m \Leftrightarrow H \cong G_{n,e}$.

Proof. $G_{n,e}$ has $\sum_{v \in V} (d_v + 1)^{-1} = m$ since each clique contributes 1 to the sum. Fixing $e = \sum_{v \in V} d_v/2$, $\sum_{v \in V} (d_v + 1)^{-1}$ is minimized with the d_v as close together as possible. Thus for any H ,

$$\alpha(H) \geq \sum_{v \in V} \frac{1}{d_v + 1} \geq m.$$

For $\alpha(H) = m$ we must have equality on both sides above. The second equality implies the d_v must be as close together as possible. Letting $X = |I|$ as in the previous theorem, assume $\alpha(H) = E[X]$. But $\alpha(H) \geq X$ for all values of $<$ so X must be a constant. Suppose H is not a union of cliques. Then, there exist $x, y, z \in V$ with $\{x, y\}, \{x, z\} \in E, \{y, z\} \notin E$. Let $<$ be an ordering that begins x, y, z and $<'$ the same ordering except that it begins y, z, x , and let I, I' be the corresponding sets of vertices all of whose neighbors are "greater." Then I, I' are identical except that $x \in I, y, z \notin I$, whereas $x \notin I', y, z \in I'$. Thus X is not constant. That is, $\alpha(H) = E[X]$ implies that H is the union of cliques and so $H \cong G_{n,e}$.

17.7 7 Martingales and Tight Concentration

Mathematics seems much more real to me than business - in the sense that, well, what's the reality in a McDonald's stand? It's here today and gone tomorrow. Now, the integers - that's reality. When you prove a theorem, you've really done something that has substance to it, to which no business venture can compare for reality.

- Jim Simons

17.7.1 7.1 Definitions

A martingale is a sequence $X_0, \dots, X_m$ of random variables so that for $0 \leq i < m$,

$$E[X_{i+1} \mid X_i, X_{i-1}, \dots, X_0] = X_i$$

Imagine a gambler walking into a casino with X_0 dollars. The casino contains a variety of games of chance. All games are "fair" in that their expectations are zero. The gambler may allow previous history to determine his choice of game and bet. He might employ the gambler's definition of martingale - double the bet until you win. He might play roulette until he wins three times and then switch to keno. Let X_i be the gambler's fortune at time i . Given that $X_i = a$ the conditional expectation of X_{i+1} must be a and so this is a martingale.

A simple but instructive martingale occurs when the gambler plays "flip a coin" for stakes of one dollar each time. Let $Y_1, \dots, Y_m$ be independent coin flips, each $+1$ or -1 with probability

1/2. Normalize so that $X_0 = 0$ is the gambler's initial stake, though he has unlimited credit. Then $X_i = Y_1 + \dots + Y_i$ has distribution S_i .

Our martingales will look quite different, at least from the outside.

The Edge Exposure Martingale. Let the random graph $G(n, p)$ be the underlying probability space. Label the potential edges $\{i, j\} \subseteq [n]$ by $e_1, \dots, e_m$, setting $m = \binom{n}{2}$ for convenience, in any specific manner. Let f be any graph theoretic function. We define a martingale $X_0, \dots, X_m$ by giving the values $X_i(H)$. $X_m(H)$ is simply $f(H)$. $X_0(H)$ is the expected value of $f(G)$ with $G \sim G(n, p)$. Note that X_0 is a constant. In general (including the cases $i = 0$ and $i = m$),

$$X_i(H) = E[f(G) \mid e_j \in G \Leftrightarrow e_j \in H, 1 \leq j \leq i].$$

In words, to find $X_i(H)$ we first expose the first i pairs $e_1, \dots, e_i$ and see if they are in H . The remaining edges are not seen and considered to be random. $X_i(H)$ is then the conditional expectation of $f(G)$ with this partial information. When $i = 0$ nothing is exposed and X_0 is a constant. When $i = m$ all is exposed and X_m is the function f . The martingale moves from no information to full information in small steps.




Fig. 7.1 The edge exposure martingale with $n = m = 3$, f is the chromatic number and the edges exposed in the order "bottom, left, right." The values $X_{\{i\}}$ are given by tracing from the central node to the leaf labeled H .

The figure shows why this is a martingale. The conditional expectation of $f(H)$ knowing the first $i - 1$ edges is the weighted average of the conditional expectations of $f(H)$, where the i th edge has been exposed. More generally - in what is sometimes referred to as a Doob martingale process - X_i may be the conditional expectation of $f(H)$ after certain information is revealed as long as the information known at time i includes the information known at time $i - 1$.

The Vertex Exposure Martingale. Again let $G(n, p)$ be the underlying probability space and f any graph theoretic function. Define $X_1, \dots, X_n$ by

$$X_i(H) = E[f(G) \mid \text{for } x, y \leq i, \{x, y\} \in G \Leftrightarrow \{x, y\} \in H].$$

In words, to find $X_i(H)$ we expose the first i vertices and all their internal edges and take the conditional expectation of $f(G)$ with that partial information. By ordering the edges appropriately the vertex exposure martingale may be considered a subsequence of the edge exposure martingale. Note that $X_1(H) = E[f(G)]$ is constant as no edges have been exposed and $X_n(H) = f(H)$ as all edges have been exposed.

17.7.2 7.2 Large Deviations

Maurey (1979) applied a large deviation inequality for martingales to prove an isoperimetric inequality for the symmetric group S_n . This inequality was useful in the study of normed spaces; see Milman and Schechtman (1986) for many related results. The applications of martingales

in graph theory also all involve the same underlying martingale result used by Maurey, which is the following.

Theorem 7.2.1 [Azuma's Inequality] Let $0 = X_0, \dots, X_m$ be a martingale with

$$|X_{i+1} - X_i| \leq 1$$

for all $0 \leq i < m$. Let $\lambda > 0$ be arbitrary. Then

$$\Pr [X_m > \lambda\sqrt{m}] < e^{-\lambda^2/2}$$

In the "flip a coin" martingale X_m has distribution S_m and this result is Theorem A.1.1. Indeed, the general proof is quite similar.

Proof. Set, with foresight, $\alpha = \lambda/\sqrt{m}$. Set $Y_i = X_i - X_{i-1}$ so that $|Y_i| \leq 1$ and $E[Y_i | X_{i-1}, X_{i-2}, \dots, X_0] = 0$. Then, as in Theorem A.1.16,

$$E[e^{\alpha Y_i} | X_{i-1}, X_{i-2}, \dots, X_0] \leq \cosh(\alpha) \leq e^{\alpha^2/2}$$

Hence

$$\begin{aligned} E[e^{\alpha X_m}] &= E\left[\prod_{i=1}^m e^{\alpha Y_i}\right] \\ &= E\left[\left(\prod_{i=1}^{m-1} e^{\alpha Y_i}\right) E[e^{\alpha Y_m} | X_{m-1}, X_{m-2}, \dots, X_0]\right] \\ &\leq E\left[\prod_{i=1}^{m-1} e^{\alpha^2/2}\right] e^{\alpha^2/2} \leq e^{\alpha^2 m/2} \end{aligned}$$

Therefore

$$\begin{aligned} \Pr [X_m > \lambda\sqrt{m}] &= \Pr [e^{\alpha X_m} > e^{\alpha\lambda\sqrt{m}}] \\ &< E[e^{\alpha X_m}] e^{-\alpha\lambda\sqrt{m}} \\ &\leq e^{\alpha^2 m/2 - \alpha\lambda\sqrt{m}} = e^{-\lambda^2/2} \end{aligned}$$

as needed.

Corollary 7.2.2 Let $c = X_0, \dots, X_m$ be a martingale with

$$|X_{i+1} - X_i| \leq 1$$

for all $0 \leq i < m$. Then

$$\Pr [|X_m - c| > \lambda\sqrt{m}] < 2e^{-\lambda^2/2}$$

A graph theoretic function f is said to satisfy the edge Lipschitz condition if whenever H and H' differ in only one edge then $|f(H) - f(H')| \leq 1$. It satisfies the vertex Lipschitz condition if whenever H and H' differ at only one vertex, $|f(H) - f(H')| \leq 1$.

Theorem 7.2.3 When f satisfies the edge Lipschitz condition the corresponding edge exposure martingale satisfies $|X_{i+1} - X_i| \leq 1$. When f satisfies the vertex Lipschitz condition the corresponding vertex exposure martingale satisfies $|X_{i+1} - X_i| \leq 1$.

We prove these results in a more general context later. They have the intuitive sense that if knowledge of a particular vertex or edge cannot change f by more than one then exposing a vertex or edge should not change the expectation of f by more than one. Now we give a simple application of these results.

Theorem 7.2.4 [Shamir and Spencer (1987)] Let n, p be arbitrary and let $c = E[\chi(G)]$, where $G \sim G(n, p)$. Then

$$\Pr[|\chi(G) - c| > \lambda\sqrt{n-1}] < 2e^{-\lambda^2/2}$$

Proof. Consider the vertex exposure martingale $X_1, \dots, X_n$ on $G(n, p)$ with $f(G) = \chi(G)$. A single vertex can always be given a new color so the vertex Lipschitz condition applies. Now apply Azuma's Inequality in the form of Corollary 7.2.2.

Letting $\lambda \rightarrow \infty$ arbitrarily slowly, this result shows that the distribution of $\chi(G)$ is "tightly concentrated" around its mean. The proof gives no clue as to where the mean is.

17.7.3 7.3 Chromatic Number

In Theorem 10.3.1 we prove that $\chi(G) \sim n/2 \log_2 n$ almost surely, where $G \sim G(n, \frac{1}{2})$. Here we give the original proof of Béla Bollobás using martingales. We follow the notations of Section 10.3, setting $f(k) = \binom{n}{k} 2^{-\binom{k}{2}}$, k_0 so that $f(k_0 - 1) > 1 > f(k_0)$, $k = k_0 - 4$ so that $k \sim 2 \log_2 n$ and $f(k) > n^{3+o(1)}$. Our goal is to show

$$\Pr[\omega(G) < k] = e^{-n^{2+o(1)}},$$

where $\omega(G)$ is the size of the maximum clique of G . We shall actually show in Theorem 7.3.2 a more precise bound. The remainder of the argument is given in Section 10.3.

Let $Y = Y(H)$ be the maximal size of a family of edge disjoint cliques of size k in H . This ingenious and unusual choice of function is key to the martingale proof.

Lemma 7.3.1 $E[Y] \geq (1 + o(1)) (n^2/2k^4)$.

Proof. Let $\mathcal{K}$ denote the family of k -cliques of G so that $f(k) = \mu = E[|\mathcal{K}|]$. Let W denote the number of unordered pairs $\{A, B\}$ of k -cliques of G with $2 \leq |A \cap B| < k$. Then $E[W] = \Delta/2$, with Δ as described in Section 10.3 (see also Section 4.5), $\Delta \sim \mu^2 k^4 n^{-2}$. Let $\mathcal{C}$ be a random subfamily of $\mathcal{K}$ defined by setting, for each $A \in \mathcal{K}$, $\Pr[A \in \mathcal{C}] = q$, q to be determined. Let W' be the number of unordered pairs $\{A, B\}$, $A, B \in \mathcal{C}$ with $2 \leq |A \cap B| < k$. Then

$$E[W'] = E[W]q^2 = \Delta q^2/2.$$

Delete from $\mathcal{C}$ one set from each such pair $\{A, B\}$. This yields a set $\mathcal{C}^*$ of edge disjoint k -cliques of G and

$$E[Y] \geq E[|\mathcal{C}^*|] \geq E[|\mathcal{C}|] - E[W'] = \mu q - \Delta q^2/2 = \mu^2/2\Delta \sim n^2/2k^4,$$

where we choose $q = \mu/\Delta (< 1)$ to minimize the quadratic.

We conjecture that Lemma 7.3.1 may be improved to $E[Y] > cn^2/k^2$. That is, with positive probability there is a family of k -cliques that are edge disjoint and cover a positive proportion of the edges.

Theorem 7.3.2

$$\Pr[\omega(G) < k] < e^{-(c+o(1))(n^2/\ln^8 n)}$$

with c a positive constant.

Proof. Let $Y_0, \dots, Y_m$, $m = \binom{n}{2}$, be the edge exposure martingale on $G(n, \frac{1}{2})$ with the function Y just defined. The function Y satisfies the edge Lipschitz condition as adding a single edge can only add at most one clique to a family of edge disjoint cliques. (Note that the Lipschitz condition would not be satisfied for the number

of k -cliques as a single edge might yield many new cliques.) G has no k -clique if and only if $Y = 0$. Apply Azuma's Inequality with $m = \binom{n}{2} \sim n^2/2$ and $E[Y] \geq (1 + o(1)) (n^2/2k^4)$. Then

$$\begin{aligned} \Pr[\omega(G) < k] &= \Pr[Y = 0] \\ &\leq \Pr[Y - E[Y] \leq -E[Y]] \\ &\leq e^{-E[Y]^2/2\binom{n}{2}} \\ &\leq e^{-(c'+o(1))n^2/k^8} \\ &= e^{-(c+o(1))n^2/\ln^8 n}, \end{aligned}$$

as desired.

Here is another example where the martingale approach requires an inventive choice of graph theoretic function.

Theorem 7.3.3 Let $p = n^{-\alpha}$, where $\alpha > \frac{5}{6}$ is fixed. Let $G = G(n, p)$. Then there exists $u = u(n, p)$ so that almost always

$$u \leq \chi(G) \leq u + 3$$

That is, $\chi(G)$ is concentrated in four values.

We first require a technical lemma that has been well known.

Lemma 7.3.4 Let α, c be fixed, $\alpha > \frac{5}{6}$. Let $p = n^{-\alpha}$. Then almost always every $c\sqrt{n}$ vertices of $G = G(n, p)$ may be three-colored.

Proof. If not, let T be a minimal set that is not three-colorable. As $T - \{x\}$ is three-colorable, x must have internal degree at least 3 in T for all $x \in T$. Thus if T has t vertices it must have at least $3t/2$ edges. The probability of this occurring for some T with at most $c\sqrt{n}$ vertices is bounded from above by

$$\sum_{t=4}^{c\sqrt{n}} \binom{n}{t} \binom{\binom{t}{2}}{3t/2} p^{3t/2}$$

We bound

$$\binom{n}{t} \leq \left(\frac{ne}{t}\right)^t \quad \text{and} \quad \binom{\binom{t}{2}}{3t/2} \leq \left(\frac{te}{3}\right)^{3t/2},$$

so each term is at most

$$\left(\frac{ne}{t} \frac{t^{3/2} e^{3/2}}{3^{3/2}} n^{-3\alpha/2}\right)^t \leq (c_1 n^{1-3\alpha/2} t^{1/2})^t \leq (c_2 n^{1-3\alpha/2} n^{1/4})^t = (c_2 n^{-\epsilon})^t$$

with $\epsilon = \frac{3}{2}\alpha - \frac{5}{4} > 0$ and the sum is therefore $o(1)$.

Proof [Theorem 7.3.3]. Let $\epsilon > 0$ be arbitrarily small and let $u = u(n, p, \epsilon)$ be the least integer so that

$$\Pr[\chi(G) \leq u] > \epsilon.$$

Now define $Y(G)$ to be the minimal size of a set of vertices S for which $G - S$ may be u -colored. This Y satisfies the vertex Lipschitz condition since at worst one could add a vertex to S . Apply the vertex exposure martingale on $G(n, p)$ to Y . Letting $\mu = E[Y]$,

$$\begin{aligned} \Pr[Y \leq \mu - \lambda\sqrt{n-1}] &< e^{-\lambda^2/2} \\ \Pr[Y \geq \mu + \lambda\sqrt{n-1}] &< e^{-\lambda^2/2} \end{aligned}$$

Let λ satisfy $e^{-\lambda^2/2} = \epsilon$ so that these tail events each have probability less than ϵ . We defined u so that, with probability at least ϵ , G would be u -colorable and hence $Y = 0$. That is, $\Pr[Y = 0] > \epsilon$. The first inequality therefore forces $\mu \leq \lambda\sqrt{n-1}$. Now employing the second inequality,

$$\Pr[Y \geq 2\lambda\sqrt{n-1}] \leq \Pr[Y \geq \mu + \lambda\sqrt{n-1}] \leq \epsilon.$$

With probability at least $1 - \epsilon$ there is a u -coloring of all but at most $c'\sqrt{n}$ vertices. By the lemma almost always, and so with probability at least $1 - \epsilon$, these points may be colored with three further colors, giving a $(u + 3)$ -coloring of G . The minimality of u guarantees that with probability at least $1 - \epsilon$ at least u colors are needed for G . Altogether

$$\Pr[u \leq \chi(G) \leq u + 3] \geq 1 - 3\epsilon,$$

and ϵ was arbitrarily small.

Using the same technique, similar results can be achieved for other values of α . Together with some related ideas it can be shown that for any fixed $\alpha > \frac{1}{2}$, $\chi(G)$ is concentrated on at most two values. See Łuczak (1991) and Alon and Krivelevich (1997) for the detailed proofs.

17.7.4 7.4 Two General Settings

The martingales useful in studying random graphs generally can be placed in the following general setting, which is essentially the one considered in Maurey (1979) and in Milman and Schechtman (1986). Let $\Omega = A^B$ denote the set of functions $g : B \rightarrow A$. (With B the set of pairs of vertices on n vertices and $A = \{0, 1\}$ we may identify $g \in A^B$ with a graph on n vertices.) We define a measure by giving values p_{ab} and setting

$$\Pr[g(b) = a] = p_{ab}$$

with the values $g(b)$ assumed mutually independent. [In $G(n, p)$ all $p_{1b} = p, p_{0b} = 1 - p$.] Now fix a gradation

$$\emptyset = B_0 \subset B_1 \subset \cdots \subset B_m = B.$$

Let $L : A^B \rightarrow R$ be a functional (e.g., clique number). We define a martingale $X_0, X_1, \dots, X_m$ by setting

$$X_i(h) = E[L(g) \mid g(b) = h(b) \text{ for all } b \in B_i].$$

X_0 is a constant, the expected value of L of the random g . X_m is L itself. The values $X_i(g)$ approach $L(g)$ as the values of $g(b)$ are "exposed." We say the functional L satisfies the Lipschitz condition relative to the gradation if for all $0 \leq i < m$,

$$h, h' \text{ differ only on } B_{i+1} - B_i \Rightarrow |L(h') - L(h)| \leq 1.$$

Theorem 7.4.1 Let L satisfy the Lipschitz condition. Then the corresponding martingale satisfies

$$|X_{i+1}(h) - X_i(h)| \leq 1$$

for all $0 \leq i < m, h \in A^B$.

Proof. Let H be the family of h' that agree with h on B_{i+1} . Then

$$X_{i+1}(h) = \sum_{h' \in H} L(h') w_{h'},$$

where $w_{h'}$ is the conditional probability that $g = h'$ given that $g = h$ on B_{i+1} . For each $h' \in H$ let $H[h']$ denote the family of h^* that agree with h' on all points except (possibly) $B_{i+1} - B_i$. The $H[h']$ partition the family of h^* agreeing with h on B_i . Thus we may express

$$X_i(h) = \sum_{h' \in H} \sum_{h^* \in H[h']} [L(h^*) q_{h^*}] w_{h'},$$

where q_{h^*} is the conditional probability that g agrees with h^* on B_{i+1} given that it agrees with h on B_i . (This is because for $h^* \in H[h']$, $w_{h'}$ is also the conditional probability that $g = h^*$ given that $g = h^*$ on B_{i+1} .) Thus

$$\begin{aligned} |X_{i+1}(h) - X_i(h)| &= \left| \sum_{h' \in H} w_{h'} \left[L(h') - \sum_{h^* \in H[h']} L(h^*) q_{h^*} \right] \right| \\ &\leq \sum_{h' \in H} w_{h'} \sum_{h^* \in H[h']} |q_{h^*} [L(h') - L(h^*)]| \end{aligned}$$

The Lipschitz condition gives $|L(h') - L(h^*)| \leq 1$ so

$$|X_{i+1}(h) - X_i(h)| \leq \sum_{h' \in H} w_{h'} \sum_{h^* \in H[h']} q_{h^*} = \sum_{h' \in H} w_{h'} = 1.$$

Now we can express Azuma's Inequality in a general form.

Theorem 7.4.2 Let L satisfy the Lipschitz condition relative to a gradation of length m and let $\mu = E[L(g)]$. Then for all $\lambda > 0$,

$$\begin{aligned} \Pr[L(g) \geq \mu + \lambda\sqrt{m}] &< e^{-\lambda^2/2} \\ \Pr[L(g) \leq \mu - \lambda\sqrt{m}] &< e^{-\lambda^2/2} \end{aligned}$$

The second general setting is taken from Alon, Kim and Spencer (1997). We assume our underlying probability space is generated by a finite set of mutually independent Yes/No choices, indexed by $i \in I$. We are given a random variable Y on this space. Let p_i denote the probability that choice i is Yes. Let c_i be such that changing choice i (keeping all else the same) can change Y by at most c_i . We call c_i the effect of i . Let C be an upper bound on all c_i . We call $p_i(1 - p_i)c_i^2$ the variance of choice i .

Now consider a solitaire game in which Paul finds the value of Y by making queries of an always truthful oracle Carole. The queries are always of a choice $i \in I$. Paul's choice of query can depend on Carole's previous responses. A strategy for Paul can then naturally be represented in a decision tree form. A "line of questioning" is a path from the root to a leaf of this tree, a sequence of questions and responses that determine Y . The total variance of a line of questioning is the sum of the variances of the queries in it.

Theorem 7.4.3 For all $\epsilon > 0$ there exists $\delta > 0$ so that the following holds. Suppose Paul has a strategy for finding Y such that every line of questioning has total variance at most σ^2 . Then

$$\Pr[|Y - E[Y]| > \alpha\sigma] \leq 2e^{-\alpha^2/2(1+\epsilon)} \quad (7.1)$$

for all positive α with $\alpha C < \sigma(1 + \epsilon)\delta$.

Applications. For a specific suboptimal bound we may take $\epsilon = \delta = 1$. If $C = O(1)$, $\alpha \rightarrow \infty$ and $\alpha = o(\sigma)$ the upper bound of (7.1) is $\exp[-\Omega(\alpha^2)]$. In many cases Paul queries all $i \in I$. Then we may take σ with $\sigma^2 = \sum_{i \in I} p_i(1 - p_i)c_i^2$. For example, consider an edge Lipschitz Y on $G(n, p)$ with $p = p(n) \rightarrow 0$. I is the set of $m = \binom{n}{2}$ potential edges, all $p_i = p$, $C = 1$ so that $\sigma = \Theta(\sqrt{n^2 p})$. If $\alpha \rightarrow \infty$ with $\alpha = o(\sqrt{n^2 p})$ the upper bound of (7.1) is again $\exp[-\Omega(\alpha^2)]$.

Proof. For simplicity we replace Y by $Y - E[Y]$ so that we shall henceforth assume $E[Y] = 0$. By symmetry we shall bound only the upper tail of Y . We set, with foresight, $\lambda = \alpha/[\sigma(1 + \epsilon)]$. Our side assumption gives that $C\lambda < \delta$. We will show

$$E[e^{\lambda Y}] \leq e^{(1+\epsilon)\lambda^2\sigma^2/2} \quad (7.2)$$

The Martingale Inequality then follows by the Markov bound

$$\Pr[Y > \alpha\sigma] < e^{-\lambda\alpha\sigma} E[e^{\lambda Y}] \leq e^{-\alpha^2/2(1+\epsilon)}$$

We first claim that for all $\epsilon > 0$ there exists $\delta > 0$ so that for $0 \leq p \leq 1$ and $|a| \leq \delta$

$$pe^{(1-p)a} + (1-p)e^{-pa} \leq e^{(1+\epsilon)p(1-p)a^2/2} \quad (7.3)$$

Take the Taylor series in a of the left-hand side. The constant term is 1, the linear term 0, the coefficient of a^2 is $\frac{1}{2}p(1-p)$ and for $j \geq 3$ the coefficient of a^j is at most

$$\frac{1}{j!}p(1-p)(p^{j-1} + (1-p)^{j-1}) \leq \frac{1}{j!}p(1-p).$$

Pick δ so that $|a| \leq \delta$ implies

$$\sum_{j=3}^{\infty} \frac{a^j}{j!} < \epsilon a^2/2.$$

(In particular, this holds for $\epsilon = \delta = 1$.) Then

$$pe^{(1-p)a} + (1-p)e^{-pa} \leq 1 + p(1-p)\frac{a^2}{2}(1+\epsilon)$$

and (7.3) follows from the inequality $1 + x \leq e^x$.

Using this δ we show (7.2) by induction on the depth M of the decision tree. For $M = 0$, Y is constant and (7.2) is immediate. Otherwise, let $p, c, v = p(1-p)c^2$ denote the probability, effect and variance, respectively, of Paul's first query. Let μ_y, μ_n denote the conditional expectations of Y if Carole's response is Yes or No, respectively. Then $0 = E[Y]$ can be split into

$$0 = p\mu_y + (1-p)\mu_n$$

The difference $\mu_y - \mu_n$ is the expected change in Y when all other choices are made independent with their respective probabilities and the root choice is changed from Yes to No. As this always changes Y by at most c ,

$$|\mu_y - \mu_n| \leq c$$

Thus we may parametrize

$$\mu_y = (1-p)b \quad \text{and} \quad \mu_n = -pb$$

with $|b| \leq c$. From (7.3)

$$pe^{\lambda\mu_y} + (1-p)e^{\lambda\mu_n} \leq e^{(1+\epsilon)p(1-p)b^2\lambda^2/2} \leq e^{(1+\epsilon)v\lambda^2/2}$$

Let A_y denote the expectation of $e^{\lambda(Y-\mu_y)}$ conditional on Carole's first response being Yes and let A_n denote the analogous quantity for No. Given Carole's first response Paul has a decision tree (one of the two main subtrees) that determines Y with total variation at most $\sigma^2 - v$ and the tree has depth at most $M - 1$. So by induction $A_y, A_n \leq A^-$, where we set

$$A^- = e^{(1+\epsilon)\lambda^2(\sigma^2-v)/2}$$

Now we split

$$\begin{aligned} E[e^{\lambda Y}] &= pe^{\lambda\mu_y}A_y + (1-p)e^{\lambda\mu_n}A_n \\ &\leq [pe^{\lambda\mu_y} + (1-p)e^{\lambda\mu_n}]A^- \\ &\leq e^{(1+\epsilon)\lambda^2(v+(\sigma^2-v))/2} \end{aligned}$$

completing the proof of (7.2) and hence of Theorem 7.4.3.

We remark that this formal inductive proof somewhat masks the martingale. A martingale $E[Y] = Y_0, \dots, Y_M = Y$ can be defined with Y_t the conditional expectation of Y after the first t queries and responses. Theorem 7.4.3 can be thought of as bounding the tail of Y by that of a normal distribution of greater or equal variance. For very large distances from the mean, large α , this bound fails.

17.7.5 7.5 Four Illustrations

Let g be the random function from $\{1, \dots, n\}$ to itself, all n^n possible functions equally likely. Let $L(g)$ be the number of values not hit, that is, the number of y for which $g(x) = y$ has no solution. By linearity of expectation,

$$E[L(g)] = n \left(1 - \frac{1}{n}\right)^n$$

and this quantity is at most n/e and at least $n(1 - 1/n)^{n-1} \cdot (1 - 1/n) > (n-1)/e$.

Set $B_i = \{1, \dots, i\}$. L satisfies the Lipschitz condition relative to this gradation since changing the value of $g(i)$ can change $L(g)$ by at most 1. Thus we have the following.

Theorem 7.5.1 $\Pr[|L(g) - \frac{n}{e}| > \lambda\sqrt{n} + 1] < 2e^{-\lambda^2/2}$.

Deriving these asymptotic bounds from first principles is quite cumbersome.

As a second illustration let B be any normed space and let $v_1, \dots, v_n \in B$ with all $|v_i| \leq 1$. Let $\epsilon_1, \dots, \epsilon_n$ be independent with $\Pr[\epsilon_i = +1] = \Pr[\epsilon_i = -1] = 1/2$ and set

$$X = |\epsilon_1 v_1 + \dots + \epsilon_n v_n|$$

Theorem 7.5.2

$$\begin{aligned} \Pr[X - E[X] > \lambda\sqrt{n}] &< e^{-\lambda^2/2} \\ \Pr[X - E[X] < -\lambda\sqrt{n}] &< e^{-\lambda^2/2} \end{aligned}$$

Proof. Consider $\{-1, +1\}^n$ as the underlying probability space with all $(\epsilon_1, \dots, \epsilon_n)$ equally likely. Then X is a random variable and we define a martingale $X_0, \dots, X_n = X$ by exposing one ϵ_i at a time. The value of ϵ_i can only change X by 2, so direct application of Theorem 7.4.1 gives $|X_{i+1} - X_i| \leq 2$. But let ϵ, ϵ' be two n -tuples differing only in the i th coordinate:

$$X_i(\epsilon) = (X_{i+1}(\epsilon) + X_{i+1}(\epsilon'))/2$$

so that

$$|X_i(\epsilon) - X_{i+1}(\epsilon)| = |X_{i+1}(\epsilon') - X_{i+1}(\epsilon)|/2 \leq 1.$$

Now apply Azuma's Inequality.

For a third illustration let ρ be the Hamming metric on $\{0, 1\}^n$. For $A \subseteq \{0, 1\}^n$ let $B(A, s)$ denote the set of $y \in \{0, 1\}^n$ so that $\rho(x, y) \leq s$ for some $x \in A$. [$A \subseteq B(A, s)$ as we may take $x = y$.]

Theorem 7.5.3 Let $\epsilon, \lambda > 0$ satisfy $e^{-\lambda^2/2} = \epsilon$. Then

$$|A| \geq \epsilon 2^n \quad \Rightarrow \quad |B(A, 2\lambda\sqrt{n})| \geq (1 - \epsilon)2^n$$

Proof. Consider $\{0, 1\}^n$ as the underlying probability space, all points equally likely. For $y \in \{0, 1\}^n$ set

$$X(y) = \min_{x \in A} \rho(x, y).$$

Let $X_0, X_1, \dots, X_n = X$ be the martingale given by exposing one coordinate of $\{0, 1\}^n$ at a time. The Lipschitz condition holds for X : If y, y' differ in just one coordinate then $|X(y) - X(y')| \leq 1$. Thus, with $\mu = E[X]$,

$$\Pr[X < \mu - \lambda\sqrt{n}] < e^{-\lambda^2/2} = \epsilon,$$

$$\Pr[X > \mu + \lambda\sqrt{n}] < e^{-\lambda^2/2} = \epsilon.$$

But

$$\Pr[X = 0] = |A|2^{-n} \geq \epsilon,$$

so $\mu \leq \lambda\sqrt{n}$. Thus

$$\Pr[X > 2\lambda\sqrt{n}] < \epsilon$$

and

$$|B(A, 2\lambda\sqrt{n})| = 2^n \Pr[X \leq 2\lambda\sqrt{n}] \geq 2^n(1 - \epsilon).$$

Actually, a much stronger result is known. Let $B(s)$ denote the ball of radius s about $(0, \dots, 0)$. The Isoperimetric Inequality proved by Harper (1966) states that

$$|A| \geq |B(r)| \quad \Rightarrow \quad |B(A, s)| \geq |B(r + s)|.$$

One may actually use this inequality as a beginning to give an alternate proof that $\chi(G) \sim n/2 \log_2 n$ and to prove a number of the other results we have shown using martingales.

We illustrate Theorem 7.4.3 with a key technical lemma (in simplified form) from Alon, Kim and Spencer (1997). Let $G = (V, E)$ be a graph on N vertices, each vertex having degree D . Asymptotics will be for $N, D \rightarrow \infty$. Set $p = 1/D$. Define a random subgraph $H \subseteq G$ by placing each edge $e \in E$ in H with independent probability p . Let M (for matching) be the set of isolated edges of H . Let V^* be those $v \in V$ not in any $\{v, w\} \in M$. For $v \in V$ set $\deg^*(v)$ equal to the number of $w \in V^*$ with $\{v, w\} \in E$. As

$$\Pr[w \notin V^*] = \sum_{\{v, w\} \in E} p(1-p)^{2D-1} = e^{-2} + O(D^{-1}),$$

linearity of expectation gives

$$E[\deg^*(v)] = D(1 - e^{-2}) + O(1)$$

We want $\deg^*(v)$ tightly concentrated about its mean.

In the notation of Theorem 7.4.3 the probability space is determined by the choices $e \in H$ for all $e \in E$. All $p_i = p$. Changing $e \in H$ to $e \notin H$ can change $\deg^*(v)$ by at most $C = 4$.

Paul needs to find $\deg^*(v)$ by queries of the form "Is $e \in H$?" For each w with $\{v, w\} \in E$ he determines if $w \in V^*$ by the following line of inquiry. First, for all u with $\{w, u\} \in E$ he queries if $\{w, u\} \in H$. If no $\{w, u\} \in H$ then $w \in V^*$. If two (or more) $\{w, u_1\}, \{w, u_2\} \in H$ then w cannot be in an isolated edge of H so $w \in V^*$. Now suppose $\{w, u\} \in H$ for precisely one u . Paul then asks (using his acquired knowledge!) for each $z \neq w$ with $\{u, z\} \in E$ if $\{u, z\} \in H$. The replies determine if $\{w, u\}$ is an isolated edge of H and hence if $w \in V^*$. Paul has made at most $D + (D - 1)$ queries for each w for a total of at most $D(2D - 1) = O(D^2)$ queries. We deduce

$$\Pr [|\deg^*(v) - D(1 - e^{-2})| > \lambda D^{1/2}] = \exp [-\Omega(\lambda^2)]$$

when $\lambda \rightarrow \infty$ and $\lambda = o(D^{1/2})$.

In application one wishes to iterate this procedure (now applying it to the restriction of G to V^*) in order to find a large matching. This is somewhat akin to the Rödl nibble of Section 4.7. There are numerous further complications but the tight concentration of $\deg^*(v)$ about its mean plays an indispensable role.

17.7.6 7.6 Talagrand's Inequality

Let $\Omega = \prod_{i=1}^n \Omega_i$, where each Ω_i is a probability space and Ω has the product measure. Let $A \subseteq \Omega$ and let $\vec{x} = (x_1, \dots, x_n) \in \Omega$. Talagrand (1996) gives an unusual, subtle and ultimately powerful notion of the distance - denoted $\rho(A, \vec{x})$ - from $\vec{x}$ to A . We imagine moving from $\vec{x}$ to some $\vec{y} = (y_1, \dots, y_n) \in A$ by changing coordinates. $\rho(A, \vec{x})$ will measure the minimal cost of such a move when a suitably restricted adversary sets the cost of each change.

Definition $2\rho(A, \vec{x})$ is the least value such that for any $\vec{\alpha} = (\alpha_1, \dots, \alpha_n) \in \mathbb{R}^n$ with $|\vec{\alpha}| = 1$ there exists $\vec{y} = (y_1, \dots, y_n) \in A$ with

$$\sum_{x_i \neq y_i} \alpha_i \leq \rho(A, \vec{x}).$$

Note that $\vec{y}$ can, and generally will, depend on $\vec{\alpha}$.

We define for any real $t \geq 0$,

$$A_t = \{\vec{x} \in \Omega : \rho(A, \vec{x}) \leq t\}.$$

Note that $A_0 = A$ as when $\vec{x} \in A$ one can select $\vec{y} = \vec{x}$.

Theorem 7.6.1 [talagrand's Inequality]

$$\Pr[A] (1 - \Pr[A_t]) \leq e^{-t^2/4}$$

In particular, if $\Pr[A] \geq \frac{1}{2}$ (or any fixed constant) and t is "very large" then all but a very small proportion of Ω is within "distance" t of A .

Example. Take $\Omega = \{0, 1\}^n$ with the uniform distribution and let τ be the Hamming (L^1) metric. Then $\rho(A, \vec{x}) \geq \min_{\vec{y} \in A} \tau(\vec{x}, \vec{y}) n^{-1/2}$ as the adversary can choose all $\alpha_i = n^{-1/2}$. Suppose to move from $\vec{x}$ to A the values $x_1, \dots, x_l$ (or any particular l coordinates) must be changed. Then $\rho(A, \vec{x}) \geq l^{1/2}$ as the adversary could choose $\alpha_i = l^{-1/2}$ for $1 \leq i \leq l$ and zero elsewhere.

Define $U(A, \vec{x})$ to be the set of $\vec{s} = (s_1, \dots, s_n) \in \{0, 1\}^n$ with the property that there exists $\vec{y} \in A$ such that

$$x_i \neq y_i \Rightarrow s_i = 1$$

We may think of $U(A, \vec{x})$ as representing the possible paths from $\vec{x}$ to A . Note that when $s_i = 1$ we, for somewhat technical reasons, do not require $x_i \neq y_i$. With this notation $\rho(A, \vec{x})$ is the least real so that for all $\vec{\alpha}$ with $|\vec{\alpha}| = 1$ there exists $\vec{s} \in U(A, \vec{x})$ with $\vec{\alpha} \cdot \vec{s} \leq \rho(A, \vec{x})$.

Now define $V(A, \vec{x})$ to be the convex hull of $U(A, \vec{x})$. The following result gives an alternate characterization of ρ that supplies the concept with much of its richness.

Theorem 7.6.2

$$\rho(A, \vec{x}) = \min_{\vec{v} \in V(A, \vec{x})} |\vec{v}|.$$

Proof. Let $\vec{v} \in V(A, \vec{x})$ achieve this minimum. The hyperplane through $\vec{v}$ perpendicular to the line from the origin to $\vec{v}$ then separates $V(A, \vec{x})$ from the origin so that all $\vec{s} \in V(A, \vec{x})$ have $\vec{s} \cdot \vec{v} \geq \vec{v} \cdot \vec{v}$. Set $\vec{\alpha} = \vec{v}/|\vec{v}|$. Then all $\vec{s} \in U(A, \vec{x}) \subseteq V(A, \vec{x})$ have $\vec{s} \cdot \vec{\alpha} \geq \vec{v} \cdot \vec{v}/|\vec{v}| = |\vec{v}|$. Conversely, take any $\vec{\alpha}$ with $|\vec{\alpha}| = 1$. Then $\vec{\alpha} \cdot \vec{v} \leq |\vec{v}|$.

As $\vec{v} \in V(A, \vec{x})$ we may write $\vec{v} = \sum \lambda_i \vec{s}_i$ for some $\vec{s}_i \in U(A, \vec{x})$, with all $\lambda_i \geq 0$ and $\sum \lambda_i = 1$. Then

$$|\vec{v}| \geq \sum \lambda_i (\vec{\alpha} \cdot \vec{s}_i)$$

and hence some $\vec{\alpha} \cdot \vec{s}_i \leq |\vec{v}|$.

The case $\Omega = \{0, 1\}^n$ is particularly important and instructive. There $\rho(A, \vec{x})$ is simply the Euclidean distance from $\vec{x}$ to the convex hull of A .

Theorem 7.6.3

$$\int_{\Omega} \exp \left[\frac{1}{4} \rho^2(A, \vec{x}) \right] d\vec{x} \leq \frac{1}{\Pr[A]}$$

Theorem 7.6.1 is an immediate corollary of the above result. Indeed, fix A and consider the random variable $X = \rho(A, \vec{x})$. Then

$$\Pr[A_t] = \Pr[X \geq t] = \Pr[e^{X^2/4} \geq e^{t^2/4}] \leq \mathbb{E}[e^{X^2/4}] e^{-t^2/4}$$

and the theorem states $\mathbb{E}[e^{X^2/4}] \leq 1/\Pr[A]$.

Proof [Theorem 7.6.3]. We use induction on the dimension n . For $n = 1$, $\rho(A, \vec{x}) = 1$ if $\vec{x} \notin A$, zero otherwise so that

$$\int \exp \left[\frac{1}{4} \rho^2(A, \vec{x}) \right] = \Pr[A] + (1 - \Pr[A])e^{1/4} \leq \frac{1}{\Pr[A]}$$

as the inequality $u + (1 - u)e^{1/4} \leq u^{-1}$ for $0 < u \leq 1$ is a simple calculus exercise.

Assume the result for n . Write $\text{OLD} = \prod_{i=1}^n \Omega_i$, $\text{NEW} = \Omega_{n+1}$ so that $\Omega = \text{OLD} \times \text{NEW}$ and any $z \in \Omega$ can be uniquely written $z = (x, \omega)$ with $x \in \text{OLD}, \omega \in \text{NEW}$. Set

$$B = \{x \in \text{OLD} : (x, \omega) \in A \text{ for some } \omega \in \text{NEW}\}$$

and for any $\omega \in \text{NEW}$ set

$$A_{\omega} = \{x \in \text{OLD} : (x, \omega) \in A\}$$

Given $z = (x, \omega) \in \Omega$ we can move to A in two basic ways - either by changing ω , which reduces the problem to moving from x to B , or by not changing ω , which reduces the problem to moving from x to A_{ω} . Thus

$$\vec{s} \in U(B, x) \Rightarrow (\vec{s}, 1) \in U(A, (x, \omega))$$

and

$$\vec{t} \in U(A_\omega, x) \Rightarrow (\vec{t}, 0) \in U(A, (x, \omega))$$

Taking the convex hulls, if $\vec{s} \in V(B, x)$ and $\vec{t} \in V(A_\omega, x)$ then $(\vec{s}, 1)$ and $(\vec{t}, 0)$ are in $V(A, (x, \omega))$ and hence for any $\lambda \in [0, 1]$,

$$((1 - \lambda)\vec{s} + \lambda\vec{t}, 1 - \lambda) \in V(A, (x, \omega)).$$

Then, by convexity,

$$\rho^2(A, (x, \omega)) \leq (1 - \lambda)^2 + |(1 - \lambda)\vec{s} + \lambda\vec{t}|^2 \leq (1 - \lambda)^2 + (1 - \lambda)|\vec{s}|^2 + \lambda|\vec{t}|^2.$$

Selecting $\vec{s}, \vec{t}$ with minimal norms yields the critical inequality

$$\rho^2(A, (x, \omega)) \leq (1 - \lambda)^2 + \lambda\rho^2(A_\omega, x) + (1 - \lambda)\rho^2(B, x).$$

Quoting from Talagrand, "The main trick of the proof is to resist the temptation to optimize now over λ ." Rather, we first fix ω and bound

$$\begin{aligned} & \int_x \exp \left[\frac{1}{4} \rho^2(A, (x, \omega)) \right] \\ & \leq e^{(1-\lambda)^2/4} \int_x \left(\exp \left[\frac{1}{4} \rho^2(A_\omega, x) \right] \right)^\lambda \left(\exp \left[\frac{1}{4} \rho^2(B, x) \right] \right)^{1-\lambda}. \end{aligned}$$

By Hölder's Inequality this is at most

$$e^{(1-\lambda)^2/4} \left[\int_x \exp \left[\frac{1}{4} \rho^2(A_\omega, x) \right] \right]^\lambda \left[\int_x \exp \left[\frac{1}{4} \rho^2(B, x) \right] \right]^{1-\lambda}$$

which by induction is at most

$$e^{(1-\lambda)^2/4} \left(\frac{1}{\Pr[A_\omega]} \right)^\lambda \left(\frac{1}{\Pr[B]} \right)^{1-\lambda} = \frac{1}{\Pr[B]} e^{(1-\lambda)^2/4} r^{-\lambda}$$

where $r = \Pr[A_\omega] / \Pr[B] \leq 1$. Now we use calculus and minimize $e^{(1-\lambda)^2/4} r^{-\lambda}$ by choosing $\lambda = 1 + 2 \ln r$ for $e^{-1/2} \leq r \leq 1$ and $\lambda = 0$ otherwise. Further (somewhat tedious but simple) calculation shows $e^{(1-\lambda)^2/4} r^{-\lambda} \leq 2 - r$ for this $\lambda = \lambda(r)$. Thus

$$\int_x \exp \left[\frac{1}{4} \rho^2(A, (x, \omega)) \right] \leq \frac{1}{\Pr[B]} \left(2 - \frac{\Pr[A_\omega]}{\Pr[B]} \right)$$

We integrate over ω giving

$$\int_\omega \int_x \exp \left[\frac{1}{4} \rho^2(A, (x, \omega)) \right] \leq \frac{1}{\Pr[B]} \left(2 - \frac{\Pr[A]}{\Pr[B]} \right) = \frac{1}{\Pr[A]} x(2 - x),$$

where $x = \Pr[A] / \Pr[B] \in [0, 1]$. But $x(2 - x) \leq 1$, completing the induction and hence the theorem.

17.7.7 7.7 Applications of Talagrand's Inequality

Let $\Omega = \prod_{i=1}^n \Omega_i$, where each Ω_i is a probability space and Ω has the product measure. Let $h : \Omega \rightarrow R$. Talagrand's Inequality enables us, under certain conditions, to show that the random variable $X = h(\cdot)$ is tightly concentrated. In this sense it can serve the same function Azuma's Inequality does for martingales and there are many cases in which it gives far stronger results.

We call $h : \Omega \rightarrow R$ Lipschitz if $|h(x) - h(y)| \leq 1$ whenever x, y differ in at most one coordinate. Talagrand's Inequality is most effective on those Lipschitz functions with the property that when $h(x) \geq s$ there are a relatively small number of coordinates that will certify that $h(x) \geq s$. We formalize this notion as follows.

Definition 3 Let $f : N \rightarrow N$. h is f -certifiable if whenever $h(x) \geq s$ there exists $I \subseteq \{1, \dots, n\}$ with $|I| \leq f(s)$ so that all $y \in \Omega$ that agree with x on the coordinates I have $h(y) \geq s$.

Example. Consider $G(n, p)$ as the product of $\binom{n}{2}$ coin flips and let $h(G)$ be the number of triangles in G . Then h is f -certifiable with $f(s) = 3s$. For if $h(G) \geq s$ there exist s triangles that together have at most $3s$ edges and any other G' with those $3s$ edges has $h(G') \geq s$. Note that I , here the indices for those $3s$ edges, very much depends on G . Also note that we need certify only lower bounds for h .

Theorem 7.7.1 Under the above assumptions and for all b, t ,

$$\Pr[X \leq b - t\sqrt{f(b)}] \Pr[X \geq b] \leq e^{-t^2/4}.$$

Proof. Set $A = \{x : h(x) < b - t\sqrt{f(b)}\}$. Now suppose $h(y) \geq b$. We claim $y \notin A_t$. Let I be a set of indices of size at most $f(b)$ that certifies $h(y) \geq b$ as given above. Define $\alpha_i = 0$ when $i \notin I$, $\alpha_i = |I|^{-1/2}$ when $i \in I$. If $y \in A_t$ there exists a $z \in A$ that differs from y in at most $t|I|^{1/2} \leq t\sqrt{f(b)}$ coordinates of I though at arbitrary coordinates outside I . Let y' agree with y on I and agree with z outside I . By the certification $h(y') \geq b$. Now y', z differ in at most $t\sqrt{f(b)}$ coordinates and so, by Lipschitz,

$$h(z) \geq h(y') - t\sqrt{f(b)} \geq b - t\sqrt{f(b)},$$

but then $z \notin A$, a contradiction. So $\Pr[X \geq b] \leq \Pr[A_t]$ and from Talagrand's Theorem,

$$\Pr[X < b - t\sqrt{f(b)}] \Pr[X \geq b] \leq e^{-t^2/4}.$$

As the right-hand side is continuous in t we may replace ' $<$ ' by ' $\leq$ ' giving the theorem.

□

A small generalization is sometimes useful. Call $h : \Omega \rightarrow R$ K -Lipschitz if $|h(x) - h(y)| \leq K$ whenever x, y differ in only one coordinate. Applying the above theorem to h/K , which is Lipschitz, we find

$$\Pr[X \leq b - tK\sqrt{f(b)}] \Pr[X \geq b] \leq e^{-t^2/4}$$

In applications one often takes b to be the median so that for t large the probability of being $t\sqrt{f(b)}$ under the median goes sharply to zero. But it works both ways, by parametrizing so that $m = b - t\sqrt{f(b)}$ is the median one usually gets $b \sim m + t\sqrt{f(m)}$ and that the probability of being $t\sqrt{f(b)}$ above the median goes sharply to zero. Martingales, via Azuma's Inequality, generally produce a concentration result around the mean μ of X while Talagrand's Inequality yields a concentration result about the median m . Means tend to be easy to compute, medians notoriously difficult, but a tight concentration result will generally allow us to show that the mean and median are not far away.

Let $x = (x_1, \dots, x_n)$, where the x_i are independently and uniformly chosen from $[0, 1]$. Set $X = h(x)$ to be the length of the longest increasing subsequence of x . Elementary methods

give that $c_1 n^{1/2} < X < c_2 n^{1/2}$ almost surely for some positive constants c_1, c_2 and that the mean μ and median m of X are both in that range. Also X is Lipschitz, as changing one x_i can only change X by at most one. How concentrated is X ? We can apply Azuma's Inequality to deduce that if $s \gg n^{1/2}$ then $|X - \mu| \leq s$ almost surely. This is not particularly good since X itself is only of order $n^{1/2}$. Now consider Talagrand's Inequality. X is f -certifiable with $f(s) = s$ since if x has an increasing subsequence of length s then those s coordinates certify that $X \geq s$. Then $\Pr[X < m - tm^{1/2}] \leq e^{-t^2/4} / \Pr[X \geq m] \leq 2e^{-t^2/4}$ as m is the median value. But $m = \Theta(n^{1/2})$. Thus when $s \gg n^{1/4}$ we have $X > m - s$ almost surely. For the other side suppose $t \rightarrow \infty$ slowly and let b be such that $b - tb^{1/2} = m$. Then $\Pr[X \geq b] \leq e^{-t^2/4} / \Pr[X \leq m] \leq 2e^{-t^2/4}$. Then $X \leq b$ almost surely. But $b = m + (1 + o(1))tm^{1/2}$ so that $X \leq m + tm^{1/2}$ almost surely. Combining, if $s \gg n^{1/4}$ then $|X - m| < s$ almost surely. A much stronger result, determining the precise asymptotic distribution of X , has been obtained by Baik, Deift and Johansson (1999), using deep analytic tools.

Let's reexamine the bound (Theorem 7.3.2) that $G(n, \frac{1}{2})$ has no clique of size k with k as defined there. We let, as there, Y be the maximal number of edge disjoint k -cliques. From the work there $E[Y] = \Omega(n^2 k^{-4})$ and Y is tightly concentrated about $E[Y]$ so that the median m of Y must also have $m = \Omega(n^2 k^{-4})$. As before Y is Lipschitz. Further Y is f -certifiable with $f(s) = \binom{k}{2}s$ as the edges of the s -cliques certify that $Y \geq s$. Hence

$$\Pr\left[Y \leq m - tm^{1/2} \binom{k}{2}^{1/2}\right] \Pr[Y \geq m] < e^{-t^2/4}$$

Set $t = \Theta(m^{1/2}/k)$ so that $m = tm^{1/2} \binom{k}{2}^{1/2}$. Then

$$\Pr[\omega(G) < k] = \Pr[Y \leq 0] < 2e^{-t^2/4} < \exp\left[-\Omega\left(\frac{n^2}{\ln^6 n}\right)\right],$$

which improves the bound of Theorem 7.3.2. Still, we should note that application of the Extended Janson Inequality in Section 10.3 does even better.

17.7.8 7.8 Kim-vu Polynomial Concentration

The approach of Kim and Vu(2000) is often useful. Let $H = (V(H), E(H))$ be a hypergraph and let each edge $e \in E(H)$ have a nonnegative weight w_e . Let $t_i, i \in V(H)$ be mutually independent indicator random variables with $E[t_i] = p_i$. Consider the random variable polynomial

$$Y = \sum_{e \in E(H)} w_e \prod_{i \in e} t_i.$$

We allow $e = \emptyset$ in which case $\prod_{i \in e} t_i$ is by convention 1. We want to show that Y is concentrated about its mean.

Let $S \subseteq V(H)$ be a random set given by $\Pr[i \in S] = p_i$, these events mutually independent over $i \in V(H)$. Then Y is the weighted number of hyperedges e in the restriction of H to S . In applications we generally have all weights equal one so that Y simply counts the hyperedges in the random S . But we may also think abstractly of Y as simply any polynomial over the indicators t_i having all nonnegative coefficients.

We set $n = |V(H)|$, the number of vertices of H (number of variables t_i). Let k be an upper bound on the size of all hyperedges (upper bound on the degree of the polynomial Y).

Let $A \subseteq V(H)$ with $|A| \leq k$. We truncate Y to Y_A as follows: For those terms $\prod_{i \in e} t_i$ with $A \subseteq e$ we set $t_i = 1$ for all $i \in A$, replacing the term by $\prod_{i \in e-A} t_i$. All other terms (where e does not contain A) are deleted. For example, with $A = \{1\}$, $2t_1t_2 + 5t_1t_3t_4 + 7t_2t_4$ becomes

$2t_2 + 5t_3t_4$. Intriguingly, as polynomials in the t_i , Y_A is the partial derivative of Y with respect to the t_i , $i \in A$. Set $E_A = \mathbb{E}[Y_A]$. That is, E_A is the expected number of hyperedges in S that contain A , conditional on all vertices of A being in S . Set E_i equal to the maximal E_A over all $A \subseteq V(H)$ of size i . Set $\mu = \mathbb{E}[Y]$ for convenience and set

$$E' = \max_{1 \leq i \leq k} E_i \quad \text{and} \quad E = \max\{\mu, E'\}$$

Theorem 7.8.1 [Kim-Vu Polynomial Concentration] With the above hypotheses

$$\Pr\left[|Y - \mu| > a_k (EE')^{1/2} \lambda^k\right] < d_k e^{-\lambda} n^{k-1}$$

for any $\lambda > 1$.

Here, for definiteness, we may take $a_k = 8^k k!^{1/2}$ and $d_k = 2e^2$.

We omit the proof, which combines martingale inequalities similar to those of Theorem 7.4.3 with a subtle induction on the degree k . There may well be room for improvement in the a_k, d_k and n^{k-1} terms. In applications one generally has k fixed and $\lambda \gg \ln n$ so that the $e^{-\lambda}$ term dominates the probability bound.

Applications of Kim-Vu polynomial concentration tend to be straightforward. Let $G \sim G(n, p)$ with $p = n^{-\alpha}$ and assume $0 < \alpha < 2/3$. Fix a vertex x of G and let $Y = Y(x)$ be the number of triangles containing x . Set $\mu = \mathbb{E}[Y] = \binom{n-1}{2} p^3 \sim \frac{1}{2} n^{2-3\alpha}$. Let $\delta > 0$ be fixed. We want to bound $\Pr[|Y - \mu| > \delta\mu]$.

The random graph G is defined by the random variables t_{ij} , one for each unordered pair of vertices, which are indicators of the adjacency of the two vertices. In that context

$$Y = \sum_{i,j \neq x} t_{xi} t_{xj} t_{ij}$$

This is a polynomial of degree $k = 3$. When A consists of a single edge $\{x, i\}$ we find $E_A = (n-2)p^2$; when it consists of three edges forming a triangle containing x we find $E_A = 1$. When $A = \emptyset$, $E_A = \mu$. Other cases give smaller E_A . Basically $E' \sim \max\{np^2, 1\}$. Calculation gives $E' \sim c\mu n^{-\epsilon}$ for some positive ϵ (dependent on α) throughout our range. We apply Kim-Vu polynomial concentration with $\lambda = c'n^{\epsilon/6}$, c' a small positive constant, to bound $\Pr[|Y - \mu| > \delta\mu]$ by $\exp[-\Omega(n^{\epsilon/6})]$. Note that the n^{k-1} factor is absorbed by the exponential.

In particular, as this probability is $o(n^{-1})$, we have that almost surely every vertex x is in $\sim \mu$ triangles. This result generalizes. Fix $\alpha \in (0, 1)$ and suppose (R, H) is a rooted graph, safe, in the sense of Section 10.4, with respect to α . Let $G \sim G(n, p)$ with $p = n^{-\alpha}$. For distinct vertices $x_1, \dots, x_r$ let $Y = Y(x_1, \dots, x_r)$ denote the number of extensions in G to H . Set $\mu = \mathbb{E}[Y]$. Kim-Vu polynomial concentration gives an exponentially small upper bound on the probability that Y is not near μ . In particular, this probability is $o(n^{-r})$. Hence almost surely every r vertices have $\sim \mu$ extensions to H .

17.7.9 7.9 Exercises

1. Let $G = (V, E)$ be the graph whose vertices are all 7^n vectors of length n over $\mathbb{Z}_7$, in which two vertices are adjacent iff they differ in precisely one coordinate. Let $U \subset V$ be a set of 7^{n-1} vertices of G , and let W be the set of all vertices of G whose distance from U exceeds $(c+2)\sqrt{n}$, where $c > 0$ is a constant. Prove that $|W| \leq 7^n \cdot e^{-c^2/2}$.
2. (*) Let $G = (V, E)$ be a graph with chromatic number $\chi(G) = 1000$. Let $U \subset V$ be a random subset of V chosen uniformly among all $2^{|V|}$ subsets of V . Let $H = G[U]$ be the induced subgraph of G on U . Prove that

$$\Pr[\chi(H) \leq 400] < 1/100$$

3. Prove that there is an absolute constant c such that for every $n > 1$ there is an interval I_n of at most $c\sqrt{n}/\log n$ consecutive integers such that the probability that the chromatic number of $G(n, 0.5)$ lies in I_n is at least 0.99 .

The Probabilistic Lens: Weierstrass Approximation Theorem

The well-known Weierstrass Approximation Theorem asserts that the set of real polynomials over $[0, 1]$ is dense in the space of all continuous real functions over $[0, 1]$. This is stated in the following theorem.

Theorem 1 [Weierstrass Approximation Theorem] For every continuous real function $f : [0, 1] \rightarrow \mathbb{R}$ and every $\epsilon > 0$, there is a polynomial $p(x)$ such that $|p(x) - f(x)| \leq \epsilon$ for all $x \in [0, 1]$.

Bernstein (1912) gave a charming probabilistic proof of this theorem, based on the properties of the binomial distribution. His proof is the following.

Proof. Since a continuous $f : [0, 1] \rightarrow \mathbb{R}$ is uniformly continuous there is a $\delta > 0$ such that if $x, x' \in [0, 1]$ and $|x - x'| \leq \delta$ then $|f(x) - f(x')| \leq \epsilon/2$. In addition, since f must be bounded there is an $M > 0$ such that $|f(x)| \leq M$ in $[0, 1]$.

Let $B(n, x)$ denote the binomial random variable with n independent trials and probability of success x for each of them. Thus the probability that $B(n, x) = j$ is precisely $\binom{n}{j} x^j (1-x)^{n-j}$. The expectation of $B(n, x)$ is nx and its standard deviation is $\sqrt{nx(1-x)} \leq \sqrt{n}$. Therefore, by Chebyshev's Inequality discussed in Chapter 4, for every integer n , $\Pr[|B(n, x) - nx| > n^{2/3}] \leq 1/n^{1/3}$. It follows that there is an integer n such that

$$\Pr[|B(n, x) - nx| > n^{2/3}] < \frac{\epsilon}{4M}$$

and

$$\frac{1}{n^{1/3}} < \delta.$$

Define

$$P_n(x) = \sum_{i=0}^n \binom{n}{i} x^i (1-x)^{n-i} f(i/n).$$

We claim that for every $x \in [0, 1]$, $|P_n(x) - f(x)| \leq \epsilon$. Indeed, since

$$\sum_{i=0}^n \binom{n}{i} x^i (1-x)^{n-i} = 1$$

we have

$$\begin{aligned}
|P_n(x) - f(x)| &\leq \sum_{i: |i-nx| \leq n^{2/3}} \binom{n}{i} x^i (1-x)^{n-i} |f(i/n) - f(x)| \\
&\quad + \sum_{i: |i-nx| > n^{2/3}} \binom{n}{i} x^i (1-x)^{n-i} [|f(i/n)| + |f(x)|] \\
&\leq \sum_{i: |i/n-x| \leq n^{-1/3} < \delta} \binom{n}{i} x^i (1-x)^{n-i} |f(i/n) - f(x)| \\
&\quad + 2M \Pr[|B(n, x) - nx| > n^{2/3}] \\
&\leq \frac{\epsilon}{2} + 2M \frac{\epsilon}{4M} = \epsilon
\end{aligned}$$

This completes the proof.

17.8 8 The Poisson Paradigm

One of the things that attracts us most when we apply ourselves to a mathematical problem is precisely that within us we always hear the call: here is the problem, search for the solution, you can find it by pure thought, for in mathematics there is no ignorabimus.

- David Hilbert

When X is the sum of many rare indicator "mostly independent" random variables and $\mu = E[X]$, we would like to say that X is close to a Poisson distribution with mean μ and, in particular, that $\Pr[X = 0]$ is nearly $e^{-\mu}$. We call this rough statement the Poisson Paradigm. In this chapter we give a number of situations in which this paradigm may be rigorously proved.

17.8.1 8.1 The Janson Inequalities

In many instances we would like to bound the probability that none of a set of bad events $\{B_i\}_{i \in I}$ occur. If the events are mutually independent then

$$\Pr \left[\bigwedge_{i \in I} \overline{B_i} \right] = \prod_{i \in I} \Pr [\overline{B_i}].$$

When the B_i are "mostly" independent the Janson Inequalities allow us, sometimes, to say that these two quantities are "nearly" equal.

Let Ω be a finite universal set and let R be a random subset of Ω given by

$$\Pr[r \in R] = p_r,$$

these events mutually independent over $r \in \Omega$. Let $\{A_i\}_{i \in I}$ be subsets of Ω , I a finite index set. Let B_i be the event $A_i \subseteq R$. (That is, each point $r \in \Omega$ "flips a coin" to determine if it is in R . B_i is the event that the coins for all $r \in A_i$ came up "heads.") Let X_i be the indicator random variable for B_i and $X = \sum_{i \in I} X_i$ the number of $A_i \subseteq R$. The event $\bigwedge_{i \in I} \overline{B_i}$ and $X = 0$ are then identical. For $i, j \in I$ we write $i \sim j$ if $i \neq j$ and $A_i \cap A_j \neq \emptyset$. Note that when $i \neq j$ and not $i \sim j$ then B_i, B_j are independent events since they involve separate coin flips. Furthermore, and this plays a crucial role in the proofs, if $i \notin J \subset I$ and not $i \sim j$ for all $j \in J$ then B_i is mutually independent of $\{B_j\}_{j \in J}$, that is, independent of any Boolean function of those B_j . This is because the coin flips on A_i and on $\cup_{j \in J} A_j$ are independent. We define

$$\Delta = \sum_{i \sim j} \Pr[B_i \wedge B_j]$$

Here the sum is over ordered pairs so that $\Delta/2$ gives the same sum over unordered pairs. We set

$$M = \prod_{i \in I} \Pr[\overline{B_i}]$$

the value of $\Pr[\bigwedge_{i \in I} \overline{B_i}]$ if the B_i were independent. Finally, we set

$$\mu = \mathbb{E}[X] = \sum_{i \in I} \Pr[B_i]$$

The following results were given in Janson, Łuczak and Ruciński (1990).

Theorem 8.1.1 [The Janson Inequality] Let $\{B_i\}_{i \in I}, \Delta, M, \mu$ be as above and assume all $\Pr[B_i] \leq \epsilon$. Then

$$M \leq \Pr\left[\bigwedge_{i \in I} \overline{B_i}\right] \leq M e^{[1/(1-\epsilon)]\Delta/2}$$

and, further,

$$\Pr\left[\bigwedge_{i \in I} \overline{B_i}\right] \leq e^{-\mu + \Delta/2}$$

For each $i \in I$

$$\Pr[\overline{B_i}] = 1 - \Pr[B_i] \leq e^{-\Pr[B_i]}$$

so, multiplying over $i \in I$,

$$M \leq e^{-\mu}$$

The two upper bounds for Theorem 8.1.1 are generally quite similar; we tend to use the second for convenience. In many asymptotic instances a simple calculation gives $M \sim e^{-\mu}$. In particular, this is always the case when $\epsilon = o(1)$ and $\epsilon\mu = o(1)$.

Perhaps the simplest example of Theorem 8.1.1 is the asymptotic probability that $G(n, c/n)$ is triangle-free, given in Section 10.1. There, as is often the case, $\epsilon = o(1)$, $\Delta = o(1)$ and μ approaches a constant k . In those instances $\Pr[\bigwedge_{i \in I} \overline{B_i}] \rightarrow e^{-k}$. This is no longer the case when Δ becomes large. Indeed, when $\Delta \geq 2\mu$ the upper bound of Theorem 8.1.1 becomes useless. Even for Δ slightly less it is improved by the following result.

Theorem 8.1.2 [The Extended Janson Inequality] Under the assumptions of Theorem 8.1.1 and the further assumption that $\Delta \geq \mu$,

$$\Pr\left[\bigwedge_{i \in I} \overline{B_i}\right] \leq e^{-\mu^2/2\Delta}.$$

Theorem 8.1.2 (when it applies) often gives a much stronger result than Chebyshev's Inequality as used in Chapter 4. In Section 4.3 we saw $\text{Var}[X] \leq \mu + \Delta$ so that

$$\Pr\left[\bigwedge_{i \in I} \overline{B_i}\right] = \Pr[X = 0] \leq \frac{\text{Var}[X]}{\mathbb{E}[X]^2} \leq \frac{\mu + \Delta}{\mu^2}.$$

Suppose $\mu \rightarrow \infty, \mu \ll \Delta$ and $\gamma = \mu^2/\Delta \rightarrow \infty$. Chebyshev's upper bound on $\Pr[X = 0]$ is then roughly γ^{-1} while Janson's upper bound is roughly $e^{-\gamma}$.

17.8.2 8.2 The Proofs

The original proofs of Janson are based on estimates of the Laplace transform of an appropriate random variable. The proof we present here follows that of Boppana and Spencer (1989). We shall use the inequalities

$$\Pr \left[B_i \mid \bigwedge_{j \in J} \overline{B_j} \right] \leq \Pr [B_i],$$

valid for all index sets $J \subset I, i \notin J$ and

$$\Pr \left[B_i \mid B_k \wedge \bigwedge_{j \in J} \overline{B_j} \right] \leq \Pr [B_i \mid B_k],$$

valid for all index sets $J \subset I, i, k \notin J$. The first follows from Theorem 6.3.2. The second is equivalent to the first since conditioning on B_k is the same as assuming $p_r = \Pr[r \in R] = 1$ for all $r \in A_k$.

Proof [Theorem 8.1.1]. The lower bound follows immediately. Order the index set $I = \{1, \dots, m\}$ for convenience. For $1 \leq i \leq m$,

$$\Pr \left[B_i \mid \bigwedge_{1 \leq j < i} \overline{B_j} \right] \leq \Pr [B_i]$$

so

$$\Pr \left[\overline{B_i} \mid \bigwedge_{1 \leq j < i} \overline{B_j} \right] \geq \Pr [\overline{B_i}]$$

and

$$\Pr \left[\bigwedge_{i \in I} \overline{B_i} \right] = \prod_{i=1}^m \Pr \left[\overline{B_i} \mid \bigwedge_{1 \leq j < i} \overline{B_j} \right] \geq \prod_{i=1}^m \Pr [\overline{B_i}].$$

Now the first upper bound. For a given i renumber, for convenience, so that $i \sim j$ for $1 \leq j \leq d$ and not for $d+1 \leq j < i$. We use the inequality $\Pr[A \mid B \wedge C] \geq \Pr[A \wedge B \mid C]$, valid for any A, B, C . With $A = B_i, B = \overline{B_1} \wedge \dots \wedge \overline{B_d}$ and $C = \overline{B_{d+1}} \wedge \dots \wedge \overline{B_{i-1}}$,

$$\begin{aligned} \Pr \left[B_i \mid \bigwedge_{1 \leq j < i} \overline{B_j} \right] &= \Pr[A \mid B \wedge C] \geq \Pr[A \wedge B \mid C] \\ &= \Pr[A \mid C] \Pr[B \mid A \wedge C] \end{aligned}$$

From the mutual independence $\Pr[A \mid C] = \Pr[A]$. We bound

$$\Pr[B \mid A \wedge C] \geq 1 - \sum_{j=1}^d \Pr[B_j \mid B_i \wedge C] \geq 1 - \sum_{j=1}^d \Pr[B_j \mid B_i]$$

from the Correlation Inequality. Thus

$$\Pr \left[B_i \mid \bigwedge_{1 \leq j < i} \overline{B_j} \right] \geq \Pr[B_i] - \sum_{j=1}^d \Pr[B_j \wedge B_i].$$

Reversing

$$\begin{aligned} \Pr \left[\overline{B_i} \mid \bigwedge_{1 \leq j < i} \overline{B_j} \right] &\leq \Pr[\overline{B_i}] + \sum_{j=1}^d \Pr[B_j \wedge B_i] \\ &\leq \Pr[\overline{B_i}] \left(1 + \frac{1}{1-\epsilon} \sum_{j=1}^d \Pr[B_j \wedge B_i] \right) \end{aligned}$$

since $\Pr[\overline{B_i}] \geq 1 - \epsilon$. Employing the inequality $1 + x \leq e^x$,

$$\Pr \left[\overline{B_i} \mid \bigwedge_{1 \leq j < i} \overline{B_j} \right] \leq \Pr[\overline{B_i}] \exp \left(\frac{1}{1-\epsilon} \sum_{j=1}^d \Pr[B_j \wedge B_i] \right).$$

For each $1 \leq i \leq m$ we plug this inequality into

$$\Pr \left[\bigwedge_{i \in I} \overline{B_i} \right] = \prod_{i=1}^m \Pr \left[\overline{B_i} \mid \bigwedge_{1 \leq j < i} \overline{B_j} \right]$$

The terms $\Pr[\overline{B_i}]$ multiply to M . The exponents add: for each $i, j \in I$ with $j < i$ and $j \sim i$ the term $\Pr[B_j \wedge B_i]$ appears once so they add to $\Delta/2$.

For the second upper bound we instead bound

$$\begin{aligned} \Pr \left[\overline{B_i} \mid \bigwedge_{1 \leq j < i} \overline{B_j} \right] &\leq 1 - \Pr[B_i] + \sum_{j=1}^d \Pr[B_j \wedge B_i] \\ &\leq \exp \left(-\Pr[B_i] + \sum_{j=1}^d \Pr[B_j \wedge B_i] \right) \end{aligned}$$

Again, the $\Pr[B_j \wedge B_i]$ terms add to $\Delta/2$ while the $-\Pr[B_i]$ terms add to $-\mu$.

Proof [Theorem 8.1.2]. The second upper bound of Theorem 8.1.1 may be rewritten

$$-\ln \left(\Pr \left[\bigwedge_{i \in I} \overline{B_i} \right] \right) \geq \sum_{i \in I} \Pr[B_i] - \frac{1}{2} \sum_{i \sim j} \Pr[B_i \wedge B_j].$$

For any set of indices $S \subset I$ the same inequality applied only to $\{B_i\}_{i \in S}$ gives

$$-\ln \left(\Pr \left[\bigwedge_{i \in S} \overline{B_i} \right] \right) \geq \sum_{i \in S} \Pr[B_i] - \frac{1}{2} \sum_{i, j \in S, i \sim j} \Pr[B_i \wedge B_j].$$

Let now S be a random subset of I given by

$$\Pr[i \in S] = p$$

with p a constant to be determined, the events mutually independent. (Here we are using probabilistic methods to prove a probability theorem!) Each term $\Pr[B_i]$ then appears with probability p and each term $\Pr[B_i \wedge B_j]$ with probability p^2 so that

$$\begin{aligned} \mathbb{E} \left[-\ln \left(\Pr \left[\bigwedge_{i \in S} \overline{B_i} \right] \right) \right] &\geq \mathbb{E} \left[\sum_{i \in S} \Pr[B_i] \right] - \frac{1}{2} \mathbb{E} \left[\sum_{i, j \in S, i \sim j} \Pr[B_i \wedge B_j] \right] \\ &= p\mu - p^2 \frac{\Delta}{2} \end{aligned}$$

We set

$$p = \frac{\mu}{\Delta}$$

so as to maximize this quantity. The added assumption of Theorem 8.1.2 assures us that the probability p is at most 1. Then

$$\mathbb{E} \left[-\ln \left(\Pr \left[\bigwedge_{i \in S} \overline{B_i} \right] \right) \right] \geq \frac{\mu^2}{2\Delta}$$

Thus there is a specific $S \subset I$ for which

$$-\ln \left(\Pr \left[\bigwedge_{i \in S} \overline{B_i} \right] \right) \geq \frac{\mu^2}{2\Delta}$$

That is,

$$\Pr \left[\bigwedge_{i \in S} \overline{B_i} \right] \leq e^{-\mu^2/2\Delta}$$

But

$$\Pr \left[\bigwedge_{i \in I} \overline{B_i} \right] \leq \Pr \left[\bigwedge_{i \in S} \overline{B_i} \right],$$

completing the proof.

17.8.3 8.3 Brun's Sieve

The more traditional approach to the Poisson Paradigm is called Brun's sieve, for its use by the number theorist T. Brun. Let $B_1, \dots, B_m$ be events, X_i the indicator random variable for B_i and $X = X_1 + \dots + X_m$ the number of B_i that hold. Let there be a hidden parameter n (so that actually $m = m(n)$, $B_i = B_i(n)$, $X = X(n)$), which will define our o, O notation. Define

$$S^{(r)} = \sum \Pr[B_{i_1} \wedge \dots \wedge B_{i_r}]$$

the sum over all sets $\{i_1, \dots, i_r\} \subseteq \{1, \dots, m\}$, and put

$$X^{(r)} = X(X-1) \dots (X-r+1).$$

The inclusion-exclusion principle gives that

$$\Pr[X=0] = \Pr[\overline{B_1} \wedge \dots \wedge \overline{B_m}] = 1 - S^{(1)} + S^{(2)} - \dots + (-1)^r S^{(r)} \dots$$

Theorem 8.3.1 Suppose there is a constant μ so that

$$\mathbb{E}[X] = S^{(1)} \rightarrow \mu$$

and such that for every fixed r ,

$$\mathbb{E} \left[\binom{X}{r} \right] = S^{(r)} \rightarrow \frac{\mu^r}{r!}$$

Then

$$\Pr[X = 0] \rightarrow e^{-\mu}$$

and indeed for every t

$$\Pr[X = t] \rightarrow \frac{\mu^t}{t!} e^{-\mu}$$

Proof. We do only the case $t = 0$. Fix $\epsilon > 0$. Choose s so that

$$\left| \sum_{r=0}^{2s} (-1)^r \frac{\mu^r}{r!} - e^{-\mu} \right| \leq \frac{\epsilon}{2}$$

The Bonferroni Inequalities state that, in general, the inclusion-exclusion formula alternately over- and underestimates $\Pr[X = 0]$. In particular,

$$\Pr[X = 0] \leq \sum_{r=0}^{2s} (-1)^r S^{(r)}$$

Select n_0 (the hidden variable) so that for $n \geq n_0$,

$$\left| S^{(r)} - \frac{\mu^r}{r!} \right| \leq \frac{\epsilon}{2(2s+1)}$$

for $0 \leq r \leq 2s$. For such n

$$\Pr[X = 0] \leq e^{-\mu} + \epsilon$$

Similarly, taking the sum to $2s+1$ we find n_0 so that for $n \geq n_0$,

$$\Pr[X = 0] \geq e^{-\mu} - \epsilon$$

As ϵ was arbitrary $\Pr[X = 0] \rightarrow e^{-\mu}$.

The threshold functions for $G \sim G(n, p)$ to contain a copy of a given graph H , derived in Section 10.1 via the Janson Inequality, were originally found using Brun's sieve. Here is an example where both methods are used. Let $G \sim G(n, p)$, the random graph of Chapter 10. Let EPIT represent the statement that every vertex lies in a triangle.

Theorem 8.3.2 Let $c > 0$ be fixed and let $p = p(n), \mu = \mu(n)$ satisfy

$$\binom{n-1}{2} p^3 = \mu$$

$$e^{-\mu} = \frac{c}{n}$$

Then

$$\lim_{n \rightarrow \infty} \Pr[G(n, p) \text{ satisfies EPIT}] = e^{-c}$$

In Spencer (1990b) threshold functions are found for a very wide class of "extension statements" that every r vertices lie in a copy of some fixed H .

Proof. First fix $x \in V(G)$. For each unordered $y, z \in V(G) - \{x\}$ let B_{xyz} be the event that $\{x, y, z\}$ is a triangle of G . Let C_x be the event $\bigwedge_{y,z} \bar{B}_{xyz}$ and X_x the corresponding indicator random variable. We use Janson's Inequality to bound $E[X_x] = \Pr[C_x]$. Here $p = o(1)$ so $\epsilon = o(1)$. $\sum \Pr[B_{xyz}] = \mu$ as defined above. Dependency $xyz \sim xuv$ occurs if and only if the sets overlap (other than in x). Hence

$$\Delta = \sum_{y,z,z'} \Pr[B_{xyz} \wedge B_{xyz'}] = O(n^3 p^5) = o(1)$$

since $p = n^{-2/3+o(1)}$. Thus

$$E[X_x] \sim e^{-\mu} = \frac{c}{n}$$

Now define

$$X = \sum_{x \in V(G)} X_x$$

the number of vertices x not lying in a triangle. Then from linearity of expectation,

$$E[X] = \sum_{x \in V(G)} E[X_x] \rightarrow c$$

We need to show that the Poisson Paradigm applies to X . Fix r . Then

$$E\left[\binom{X}{r}\right] = S^{(r)} = \sum \Pr[C_{x_1} \wedge \cdots \wedge C_{x_r}],$$

the sum over all sets of vertices $\{x_1, \dots, x_r\}$. All r -sets look alike so

$$E\left[\binom{X}{r}\right] = \binom{n}{r} \Pr[C_{x_1} \wedge \cdots \wedge C_{x_r}] \sim \frac{n^r}{r!} \Pr[C_{x_1} \wedge \cdots \wedge C_{x_r}],$$

where $x_1, \dots, x_r$ are some particular vertices. But

$$C_{x_1} \wedge \cdots \wedge C_{x_r} = \bigwedge \overline{B_{x_i y z}}$$

the conjunction over $1 \leq i \leq r$ and all y, z . We apply Janson's Inequality to this conjunction. Again $\epsilon = p^3 = o(1)$. The number of $\{x_i, y, z\}$ is $r \binom{n-1}{2} - O(n)$, the overcount coming from those triangles containing two (or three) of the x_i . (Here it is crucial that r is fixed.) Thus

$$\sum \Pr[B_{x_i y z}] = p^3 \left(r \binom{n-1}{2} - O(n) \right) = r\mu + O(n^{-1+o(1)})$$

As before Δ is p^5 times the number of pairs $x_i y z \sim x_j y' z'$. There are $O(rn^3) = O(n^3)$ terms with $i = j$ and $O(r^2 n^2) = O(n^2)$ terms with $i \neq j$ so again $\Delta = o(1)$. Therefore

$$\Pr[C_{x_1} \wedge \cdots \wedge C_{x_r}] \sim e^{-r\mu}$$

and

$$E\left[\binom{X}{r}\right] \sim \frac{(ne^{-\mu})^r}{r!} = \frac{c^r}{r!}.$$

Hence the conditions of Theorem 8.3.1 are met for X .

17.8.4 8.4 Large Deviations

We return to the formulation of Section 8.1. Our object is to derive large deviation results on X similar to those in Appendix A. Given a point in the probability space (i.e., a selection of R) we call an index set $J \subseteq I$ a disjoint family (abbreviated disfam) if

- B_j for every $j \in J$.
- For no $j, j' \in J$ is $j \sim j'$.

If, in addition,

- If $j' \notin J$ and $B_{j'}$ then $j \sim j'$ for some $j \in J$,
then we call J a maximal disjoint family (maxdisfam). We give some general results on the possible sizes of maxdisfams. The connection to X must then be done on an ad hoc basis.

Lemma 8.4.1 With the above notation and for any integer s ,

$$\Pr[\text{there exists a disfam } J, |J| = s] \leq \frac{\mu^s}{s!}.$$

Proof. Let $\sum^*$ denote the sum over all s -sets $J \subseteq I$ with no $j \sim j'$. Let $\sum^o$ denote the sum over ordered s -tuples $(j_1, \dots, j_s)$ with $\{j_1, \dots, j_s\}$ forming such a J . Let $\sum^a$ denote the sum over all ordered s -tuples $(j_1, \dots, j_s)$. Then

$$\begin{aligned} \Pr[\text{there exists a disfam } J, |J| = s] &\leq \sum^* \Pr \left[\bigwedge_{j \in J} B_j \right] \\ &= \sum^* \prod_{j \in J} \Pr[B_j] = \frac{1}{s!} \sum^o \Pr[B_{j_1}] \cdots \Pr[B_{j_s}] \\ &\leq \frac{1}{s!} \sum^a \Pr[B_{j_1}] \cdots \Pr[B_{j_s}] \leq \frac{1}{s!} \left(\sum_{i \in I} \Pr[B_i] \right)^s = \frac{\mu^s}{s!}. \end{aligned}$$

Lemma 8.4.1 gives an effective upper bound when $\mu^s \ll s!$ - basically if $s > \mu\alpha$ for $\alpha > e$. For smaller s we look at the further condition of J being a maxdisfam. To that end we let μ_s denote the minimum, over all $j_1, \dots, j_s \in I$ of $\sum \Pr[B_i]$, the sum taken over all $i \in I$ except those i with $i \sim j_l$ for some $1 \leq l \leq s$. In application s will be small (otherwise we use Lemma 8.4.1) and μ_s will be close to μ . For some applications it is convenient to set

$$\nu = \max_{j \in I} \sum_{i \sim j} \Pr[B_i]$$

and note that $\mu_s \geq \mu - s\nu$.

Lemma 8.4.2 With the above notation and for any integer s ,

$$\begin{aligned} \Pr[\text{there exists a maxdisfam } J, |J| = s] &\leq \frac{\mu^s}{s!} e^{-\mu_s} e^{\Delta/2} \\ &\leq \frac{\mu^s}{s!} e^{-\mu} e^{s\nu} e^{\Delta/2} \end{aligned}$$

Proof. As in Lemma 8.4.1 we bound this probability by $\sum^*$ of $J = \{j_1, \dots, j_s\}$ being a maxdisfam. For this to occur J must first be a disfam and then $\bigwedge^* \overline{B_i}$, where $\bigwedge^*$ is the conjunction over all $i \in I$ except those with $i \sim j_l$ for some $1 \leq l \leq s$. We apply Janson's Inequality to give an upper bound to $\Pr[\bigwedge^* \overline{B_i}]$. The associated values μ^*, Δ^* satisfy

$$\begin{aligned}\mu^* &\geq \mu_s \\ \Delta^* &\leq \Delta\end{aligned}$$

the latter since Δ^* has simply fewer addends. Thus

$$\Pr \left[\bigwedge_{i=1}^* \overline{B_i} \right] \leq e^{-\mu_s} e^{\Delta/2}$$

and

$$\begin{aligned}\sum_{j=1}^* \Pr[J \text{ maxdisfam}] &\leq e^{-\mu_s} e^{\Delta/2} \sum_{j=1}^* \Pr \left[\bigwedge_{j \in J} B_j \right] \\ &\leq e^{-\mu_s} e^{\Delta/2} \frac{\mu^s}{s!}\end{aligned}$$

When $\Delta = o(1)$ and $\nu\mu = o(1)$ or, more generally, $\mu_{3\mu} = \mu + o(1)$, then Lemma 8.4.2 gives a close approximation to the Poisson distribution since

$$\Pr[\text{there exists a maxdisfam } J, |J| = s] \leq (1 + o(1)) \frac{\mu^s}{s!} e^{-\mu}$$

for $s \leq 3\mu$ and the probability is quite small for larger s by Lemma 8.4.1.

17.8.5 8.5 Counting Extensions

We begin with a case that uses the basic large deviation results of Appendix A.

Theorem 8.5.1 Set $p = [(\ln n)/n]\omega(n)$, where $\omega(n) \rightarrow \infty$ arbitrarily slowly. Then in $G(n, p)$ almost always

$$\deg(x) \sim (n-1)p$$

for all vertices x .

This is actually a large deviation result. It suffices to show the following.

Theorem 8.5.2 Set $p = [(\ln n)/n]\omega(n)$, where $\omega(n) \rightarrow \infty$ arbitrarily slowly. Let $x \in G$ be fixed. Fix $\epsilon > 0$. Then

$$\Pr[|\deg(x) - (n-1)p| > \epsilon(n-1)p] = o(n^{-1})$$

Proof. As $\deg(x) \sim B(n-1, p)$, that is, it is a binomial random variable with the above parameters, we have from Corollary A.1.14 that

$$\Pr[|\deg(x) - (n-1)p| > \epsilon(n-1)p] < 2e^{-c_\epsilon(n-1)p} = o(n^{-1}),$$

as c_ϵ is fixed and $(n-1)p \gg \ln n$.

This result illustrates why logarithmic terms appear so often in the study of random graphs. We want every x to have a property, hence we try to get the failure probability down to $o(n^{-1})$. When the Poisson Paradigm applies the failure probability is roughly an exponential, and hence we want the exponent to be logarithmic. This often leads to a logarithmic term for the edge probability p .

In Section 8.3 we found the threshold function for every vertex to lie on a triangle. It basically occurred when the expected number of extensions of a given vertex to a triangle reached $\ln n$. Now set $N(x)$ to be the number of triangles containing x . Set $\mu = \binom{n-1}{2}p^3 = E[N(x)]$.

Theorem 8.5.3 Let p be such that $\mu \gg \ln n$. Then almost always

$$N(x) \sim \mu$$

for all $x \in G(n, p)$.

As above, this is actually a large deviation result. We actually show the following.

Theorem 8.5.4 Let p be such that $\mu \gg \ln n$. Let $x \in G$ be fixed. Fix $\epsilon > 0$. Then

$$\Pr[|N(x) - \mu| > \epsilon\mu] = o(n^{-1})$$

Proof. We shall prove this under the further assumption $p = n^{-2/3+o(1)}$ (or equivalently, $\mu = n^{o(1)}$), which could be removed by technical methods. We now have, in the notation of Lemmas 8.4.1 and 8.4.2, $\nu\mu, \Delta = o(1)$. Let P denote the Poisson distribution with mean μ . Then

$$\begin{aligned} & \Pr[\text{there exists a maxdisfam } J, |J| \leq \mu(1 - \epsilon)] \\ & \leq (1 + o(1)) \Pr[P \leq \mu(1 - \epsilon)], \\ & \Pr[\text{there exists a maxdisfam } J, \mu(1 + \epsilon) \leq |J| \leq 3\mu] \\ & \leq (1 + o(1)) \Pr[\mu(1 + \epsilon) \leq P \leq 3\mu], \\ & \Pr[\text{there exists a maxdisfam } J, |J| \geq 3\mu] \\ & \leq \Pr[\text{there exists a disfam } J, |J| \geq 3\mu] \leq \sum_{s=3\mu}^{\infty} \frac{\mu^s}{s!} = O((1 - c)^\mu), \end{aligned}$$

where $c > 0$ is an absolute constant. Since $\mu \gg \ln n$ the third term is $o(n^{-1})$. The first and second terms are $o(n^{-1})$ by Theorem A.1.15. With probability $1 - o(n^{-1})$ every maxdisfam J has size between $(1 - \epsilon)\mu$ and $(1 + \epsilon)\mu$.

Fix one such J . (There always is some maximal disfam - even if no B_i held we could take $J = \emptyset$.) The elements of J are triples xyz that form triangles, hence $N(x) \geq |J| \geq (1 - \epsilon)\mu$. The upper bound is ad hoc. The probability that there exist five triangles of the form $xyz_1, xyz_2, xyz_3, xyz_4, xyz_5$ is at most $n^6 p^{11} = o(n^{-1})$. The probability that there exist triangles $xy_i z_i, xy_i z'_i, 1 \leq i \leq 4$, all vertices distinct, is at most $n^{12} p^{20} = o(n^{-1})$. Consider the graph whose vertices are the triangles xyz , with $\sim$ giving the edge relation. There are $N(x)$ vertices, the maxdisfam J are the maximal independent sets. In this graph, with probability $1 - o(n^{-1})$, each vertex xyz has degree at most nine and there is no set of four disjoint edges. This implies that for any $J, |J| \geq N(x) - 27$ and

$$N(x) \leq (1 + \epsilon)\mu + 27 \leq (1 + \epsilon')\mu.$$

For any graph H with "roots" $x_1, \dots, x_r$ we can examine in $G(n, p)$ the number of extensions $N(x_1, \dots, x_r)$ of a given set of r vertices to a copy of H . In Spencer (1990a) some general results are given that generalize Theorems 8.5.2 and 8.5.4. Under fairly wide assumptions (see Exercise 5, Chapter 10), when the expected number μ of extensions satisfies $\mu \gg \ln n$ then almost always all $N(x_1, \dots, x_r) \sim \mu$.

17.8.6 8.6 Counting Representations

The results of this section shall use the following very basic and very useful result.

Lemma 8.6.1 [The Borel-Cantelli Lemma] Let $\{A_n\}_{n \in \mathbb{N}}$ be events with

$$\sum_{n=1}^{\infty} \Pr[A_n] < \infty$$

Then

$$\Pr \left[\bigwedge_{i=1}^{\infty} \bigvee_{j=i}^{\infty} A_j \right] = 0$$

That is, almost always A_n is false for all sufficiently large n . In application we shall aim for $\Pr[A_n] < n^{-c}$ with $c > 1$ in order to apply this lemma.

Again we begin with a case that involves only the large deviation results of Appendix A. For a given set S of natural numbers let (for every $n \in \mathbb{N}$) $f(n) = f_S(n)$ denote the number of representations $n = x + y, x, y \in S, x < y$.

Theorem 8.6.2 [Erdős (1956)] There is a set S for which $f(n) = \Theta(\ln n)$. That is, there is a set S and constants c_1, c_2 so that for all sufficiently large n

$$c_1 \ln n \leq f(n) \leq c_2 \ln n.$$

Proof. Define S randomly by

$$\Pr[x \in S] = p_x = \min \left\{ 10 \sqrt{\frac{\ln x}{x}}, 1 \right\}$$

Fix n . Now $f(n)$ is a random variable with mean

$$\mu = E[f(n)] = \frac{1}{2} \sum_{\substack{x+y=n \\ x \neq y}} p_x p_y$$

Roughly there are n addends with $p_x p_y > p_n^2 = (100 \ln n)/n$. We have $p_x p_x = \Theta((\ln n)/n)$ except in the regions $x = o(n), y = o(n)$ and care must be taken that those terms don't contribute significantly to μ . Careful asymptotics (and first year calculus!) yield

$$\mu \sim 50 \ln n \int_0^1 \frac{dx}{\sqrt{x(1-x)}} = 50\pi \ln n$$

The negligible effect of the $x = o(n), y = o(n)$ terms reflects the finiteness of the indefinite integral at poles $x = 0$ and $x = 1$. The possible representations $x + y = n$ are mutually independent events so that from Corollary A.1.14,

$$\Pr[|f(n) - \mu| > \epsilon \mu] < 2e^{-\delta \mu}$$

for constants $\epsilon, \delta = \delta(\epsilon)$. To be specific we can take $\epsilon = 0.9, \delta = 0.1$ and

$$\Pr[|f(n) - \mu| > 0.9\mu] < 2e^{-5\pi \ln n} < n^{-1.1}$$

for n sufficiently large. Take $c_1 < 0.1(50\pi)$ and $c_2 > 1.9(50\pi)$.

Let A_n be the event that $c_1 \ln n \leq f(n) \leq c_2 \ln n$ does not hold. We have $\Pr[A_n] < n^{-1.1}$ for n sufficiently large. The Borel-Cantelli Lemma applies, almost always all A_n fail for n sufficiently large. Thus there exists a specific point in the probability space, that is, a specific set S , for which $c_1 \ln n \leq f(n) \leq c_2 \ln n$ for all sufficiently large n .

The development of the infinite probability space used here, and below, has been carefully done in the book Sequences by H. Halberstam and K. F. Roth (1983).

The use of the infinite probability space leaves a number of questions about the existential nature of the proof that go beyond the algorithmic. For example, does there exist a recursive set S having the property of Theorem 8.6.2? An affirmative answer is given in Kolountzakis (1999).

Now for a given set S of natural numbers let $g(n) = g_S(n)$ denote the number of representations $n = x + y + z, x, y, z \in S, x < y < z$. The following result was actually proved for representations of n as the sum of k terms for any fixed k . For simplicity we present here only the proof for $k = 3$.

Theorem 8.6.3 [Erdős and Tetali (1990)] There is a set S for which $g(n) = \Theta(\ln n)$. That is, there is a set S and constants c_1, c_2 so that for all sufficiently large n ,

$$c_1 \ln n \leq g(n) \leq c_2 \ln n.$$

Proof. Define S randomly by

$$\Pr[x \in S] = p_x = \min \left\{ 10 \left(\frac{\ln x}{x^2} \right)^{1/3}, \frac{1}{2} \right\}$$

Fix n . Now $g(n)$ is a random variable and

$$\mu = E[g(n)] = \sum_{x+y+z=n} p_x p_y p_z$$

Careful asymptotics give

$$\mu \sim \frac{10^3}{6} \ln n \int_{x=0}^1 \int_{y=0}^{1-x} \frac{dx dy}{[xy(1-x-y)]^{2/3}} = K \ln n$$

where K is large. (We may make K arbitrarily large by increasing "10.") We apply Lemma 8.4.2. Here

$$\Delta = \sum p_x p_y p_z p_{y'} p_{z'}$$

the sum over all five-tuples with $x+y+z = x+y'+z' = n$. Roughly there are n^3 terms, each $\sim p_n^5 = n^{-10/3+o(1)}$ so that the sum is $o(1)$. Again, care must be taken that those terms with one (or more) small variables don't contribute much to the sum. We bound $s \leq 3\mu = \Theta(\ln n)$ and consider μ_s . This is the minimal possible $\sum p_x p_y p_z$ over all those x, y, z with $x+y+z = n$ that do not intersect a given set of s representations; let us weaken that and say a given set of $3s$ elements. Again one needs that the weight of $\sum_{x+y+z=n} p_x p_y p_z$ is not on the edges but "spread" in

the center and one shows $\mu_s \sim \mu$. Now, as in Section 8.5, let P denote the Poisson distribution with mean μ . The probability that there exists a maxdisfam J of size less than $\mu(1-\epsilon)$ or between $\mu(1+\epsilon)$ and 3μ is asymptotically the probability that P lies in that range. For moderate ϵ , as K is large, these - as well as the probability of having a disfam of size bigger than 3μ - will be $o(n^{-c})$ with $c > 1$. By the Borel-Cantelli Lemma almost always all sufficiently large n will have all maxdisfam J of size between $c_1 \ln n$ and $c_2 \ln n$. Then $g(n) \geq c_1 \ln n$ immediately.

The upper bound is again ad hoc. With this p let $f(n)$ be, as before, the number of representations of n as the sum of two elements of S . We use only that $p_x = x^{-2/3+o(1)}$. We calculate

$$E[f(n)] = \sum_{x+y=n} (xy)^{-2/3+o(1)} = n^{-1/3+o(1)},$$

again watching the "pole" at 0. Here the possible representations are mutually independent so

$$\Pr[f(n) \geq 4] \leq E[f(n)]^4 / 4! = n^{-4/3+o(1)},$$

and by the Borel-Cantelli Lemma almost always $f(n) \leq 3$ for all sufficiently large n . But then almost always there is a C so that $f(n) \leq C$ for all n . For all sufficiently large n there is a maxdisfam (with representations as the sum of three terms) of size less than $c_2 \ln n$. Every triple $x, y, z \in S$ with $x + y + z = n$ must contain at least one of these at most $3c_2 \ln n$ points. The number of triples $x, y, z \in S$ with $x + y + z = n$ for a particular x is simply $f(n - x)$, the number of representations $n - x = y + z$ (possibly one less since $y, z \neq x$), and so is at most C . But then there are at most $C(3c_2 \ln n)$ total representations $n = x + y + z$.

17.8.7 8.7 Further Inequalities

Here we discuss some further results that allow one, sometimes, to apply the Poisson Paradigm. Let $B_i, i \in I$ be events in an arbitrary probability space. As in the Lovász Local Lemma of Chapter 5 we say that a symmetric binary relation ' $\sim$ ' on I is a dependency digraph if for each $i \in I$ the event B_i is mutually independent of $\{B_j : i \not\sim j\}$. [The digraph of Section 5.1 has $E = \{(i, j) : i \sim j\}$.] Suppose the events B_i satisfy the inequalities of Section 8.2:

$$\Pr \left[B_i \mid \bigwedge_{j \in J} \overline{B_j} \right] \leq \Pr [B_i]$$

valid for all index sets $J \subset I, i \notin J$ and

$$\Pr \left[B_i \mid B_k \wedge \bigwedge_{j \in J} \overline{B_j} \right] \leq \Pr [B_i \mid B_k]$$

valid for all index sets $J \subset I, i, k \notin J$. Then the Janson Inequalities in Theorems 8.1.1 and 8.1.2 and also Lemmas 8.4.1 and 8.4.2 hold as stated. The proofs are identical, the above are the only properties of the events B_i that were used.

Suen (1990) [see also Janson (1998) for significant variations] has given a very general result that allows the approximation of $\Pr [\bigwedge_{i \in I} \overline{B_i}]$ by $M = \prod_{i \in I} \Pr [\overline{B_i}]$. Again let $\{B_i\}_{i \in I}$ be events in an arbitrary probability space. We say that a binary relation $\sim$ on I is a superdependency digraph if the following holds: Let $J_1, J_2 \subset I$ be disjoint subsets so that $j_1 \sim j_2$ for no $j_1 \in J_1, j_2 \in J_2$. Let B^1 be any Boolean combination of the events $\{B_j\}_{j \in J_1}$ and let B^2 be any Boolean combination of the events $\{B_j\}_{j \in J_2}$. Then B^1, B^2 are independent. Note that the ' $\sim$ ' of Section 8.1 is indeed a superdependency digraph.

Theorem 8.7.1 [Suen] Under the above conditions,

$$\left| \Pr \left[\bigwedge_{i \in I} \overline{B_i} \right] - M \right| \leq M [e^{\sum_{i \sim j} y(i, j)} - 1]$$

where

$$y(i, j) = (\Pr [B_i \wedge B_j] + \Pr [B_i] \Pr [B_j]) \prod_{l \sim i \text{ or } l \sim j} (1 - \Pr [B_l])^{-1}.$$

We shall not prove Theorem 8.7.1. In many instances the above product is not large. Suppose it is less than two for all $i \sim j$. In that instance

$$\sum_{i \sim j} y(i, j) \leq 2 \left(\Delta + \sum_{i \sim j} \Pr [B_i] \Pr [B_j] \right).$$

In many instances $\sum_{i \sim j} \Pr[B_i] \Pr[B_j]$ is small relative to Δ (as in many instances when $i \sim j$ the events B_i, B_j are positively correlated). When, furthermore, $\Delta = o(1)$, Suen's Theorem gives the approximation of $\Pr[\bigwedge_{i \in I} \overline{B_i}]$ by M . Suen has applied this result to examinations of the number of induced copies of a fixed graph H in the random $G(n, p)$.

Janson (1990) has given a one-way large deviation result on the X of Section 8.1 that is somewhat simpler to apply than Lemmas 8.4.1 and 8.4.2.

Theorem 8.7.2 [Janson] With $\mu = E[X]$ and $\gamma > 0$ arbitrary,

$$\Pr[X \leq (1 - \gamma)\mu] < e^{-\gamma^2 \mu / [2 + (\Delta/\mu)]}$$

When $\Delta = o(\mu)$ this bound on the tail approximates that of the normal curve with mean and standard deviation μ . We shall not prove Theorem 8.7.2 here. The proofs of Theorems 8.7.1 and 8.7.2 as well as the original proofs by Janson of Theorems 8.1.1 and 8.1.2 are based on estimations of the Laplace transform of X , bounding $E[e^{-tX}]$.

17.8.8 8.8 Exercises

1. Prove that for every $\epsilon > 0$ there is some $n_0 = n_0(\epsilon)$ so that for every $n > n_0$ there is a graph on n vertices containing every graph on $k \leq (2 - \epsilon) \log_2 n$ vertices as an induced subgraph.
2. Find a threshold function for the property: $G(n, p)$ contains at least $n/6$ pairwise vertex disjoint triangles.

The Probabilistic Lens: Local Coloring

This result of Erdős (1962) gives further probabilistic evidence that the chromatic number $\chi(G)$ cannot be deduced from local considerations.

Theorem 1 For all k there exists $\epsilon > 0$ so that for all sufficiently large n there exist graphs G on n vertices with $\chi(G) > k$ and yet $\chi(G|_S) \leq 3$ for every set S of vertices of size at most ϵn .

Proof. For a given k let $c, \epsilon > 0$ satisfy (with foresight)

$$\begin{aligned} c &> 2k^2 H(1/k) \ln 2 \\ \epsilon &< e^{-5} 3^3 c^{-3} \end{aligned}$$

where $H(x) = -x \log_2 x - (1 - x) \log_2 (1 - x)$ is the entropy function. Set $p = c/n$ and let $G \sim G(n, p)$. We show that G almost surely satisfies the two conditions of the theorem.

If $\chi(G) \leq k$ there would be an independent set of size n/k . The expected number of such sets is

$$\binom{n}{n/k} (1 - p)^{\binom{n/k}{2}} < 2^{n(H(1/k) + o(1))} e^{-cn/2k^2(1+o(1))},$$

which is $o(1)$ by our condition on c . Hence almost surely $\chi(G) > k$.

Suppose some set S with $t \leq \epsilon n$ vertices required at least four colors. Then as in the proof of Lemma 7.3.4 there would be a minimal such set S . For any $v \in S$ there would be a three-coloring of $S - \{v\}$. If v had two or fewer neighbors in S then this could be extended to a three-coloring of S . Hence every $v \in S$ would have degree

at least three in $G|_S$ and so $G|_S$ would have at least $3t/2$ edges. The probability that some $t \leq \epsilon n$ vertices have at least $3t/2$ edges is less than

$$\sum_{t \leq \epsilon n} \binom{n}{t} \binom{\binom{t}{2}}{3t/2} \left(\frac{c}{n}\right)^{3t/2}$$

We outline the analysis. When $t = O(1)$ the terms are negligible. Otherwise we bound each term from above by

$$\left[\frac{ne}{t} \left(\frac{te}{3}\right)^{3/2} \left(\frac{c}{n}\right)^{3/2} \right]^t \leq \left[e^{5/2} 3^{-3/2} c^{3/2} \sqrt{t/n} \right]^t$$

Now since $t \leq \epsilon n$ the bracketed term is at most $e^{5/2} 3^{-3/2} c^{3/2} \epsilon^{1/2}$, which is less than one by our condition on ϵ . The full sum is $o(1)$; that is, almost surely no such S exists.

Many tempting conjectures are easily disproved by the probabilistic method. If every $n/\ln n$ vertices may be three-colored then can a graph G on n vertices be four-colored? This result shows that the answer is no.

17.9 9 Pseudorandomness

'A knot!', said Alice, always ready to make herself useful, and looking anxiously about her. 'Oh, do let me help to undo it!'

- from Alice in Wonderland, by Lewis Carroll

As shown in the various chapters of this book, the probabilistic method is a powerful tool for establishing the existence of combinatorial structures with certain properties. It is often the case that such an existence proof is not sufficient; we actually prefer an explicit construction. This is not only because an explicit construction may shed more light on the corresponding problem, but also because it often happens that a random-looking structure is useful for a certain algorithmic procedure; in this case we would like to have an algorithm and not merely to prove that it exists.

The problem of finding explicit constructions may look trivial; after all, since we are mainly dealing with finite cases, once we have a probabilistic proof of existence we can find an explicit example by exhaustive search. Moreover, many of the probabilistic proofs of existence actually show that most members of a properly chosen random space have the desired properties. We may thus expect that it would not be too difficult to find one such member. Although this is true in principle, it is certainly not practical to check all possibilities; it is thus common to define an explicit construction of a combinatorial object as one that can be performed efficiently; say, in time that is polynomial in the parameters of the object.

Let us illustrate this notion by one of the best known open problems in the area of explicit constructions, the problem of constructing explicit Ramsey graphs. The first example given in Chapter 1 is the proof of Erdős that for every n there are graphs on n vertices containing neither a clique nor an independent set on $2 \log_2 n$ vertices. This proof is an existence proof; can we actually describe such graphs explicitly? Erdős offered a prize of \$500 for the explicit construction of an infinite family of graphs, in which there is neither a clique nor an independent set of size more than a constant times the logarithm of the number of vertices, for some absolute constant. Of course, we can, in principle, for every fixed n , check all graphs on n vertices until we find a good one, but this does not give an efficient way of producing the desired graphs and hence is not explicit. Although the problem mentioned above received a considerable amount of attention, it is still open. The best known explicit construction is due to Frankl and Wilson (1981), who describe explicit graphs on n vertices which contain neither a clique nor an independent set on more than $2^{c\sqrt{\log n \log \log n}}$ vertices, for some absolute positive constant c .

Although the problem of constructing explicit Ramsey graphs is still open, there are several other problems for which explicit constructions are known. In this chapter we present a few examples and discuss briefly some of their algorithmic applications. We also describe several seemingly unrelated properties of a graph, which all turn out to be equivalent. All these are properties of the random graph and it is thus common to call a graph that satisfies these properties quasirandom. The equivalence of all these properties enables one to show, in several cases, that certain explicit graphs have many pseudorandom properties by merely showing that they possess one of them.

17.9.1 9.1 The Quadratic Residue Tournaments

Recall that a tournament on a set V of n players is an orientation $T = (V, E)$ of the set of edges of the complete graph on the set of vertices V . If (x, y) is a directed edge we say that x beats y . Given a permutation π of the set of players, a (directed) edge (x, y) of the tournament is consistent with π if x precedes y in π . If π is viewed as a ranking of the players, then it is reasonable to try and find rankings with as many consistent arcs as possible. Let $c(\pi, T)$ denote the number of arcs of T which are consistent with π , and define $c(T) = \max c(\pi, T)$, where the maximum is taken over all permutations π of the set of vertices of T . For every tournament T on n players, if $\pi = 1, 2, \dots, n$ and $\pi' = n, n-1, \dots, 1$ then $c(\pi, T) + c(\pi', T) = \binom{n}{2}$. Therefore $c(T) \geq \frac{1}{2} \binom{n}{2}$. In fact, it can be shown that for every such T , $c(T) \geq \frac{1}{2} \binom{n}{2} + \Omega(n^{3/2})$. On the other hand, a simple probabilistic argument shows that there are tournaments T on n players for which $c(T) \leq (1 + o(1)) \frac{1}{2} \binom{n}{2}$. [The best known estimate, which gives the right order of magnitude for the largest possible value of the difference of $c(T) - \frac{1}{2} \binom{n}{2}$ is more complicated and was given by de la Vega (1983), where he showed that there are tournaments T on n players for which $c(T) \leq \frac{1}{2} \binom{n}{2} + O(n^{3/2})$.]

Can we describe explicitly tournaments T on n vertices in which $c(T) \leq (1 + o(1)) \frac{1}{2} \binom{n}{2}$? This problem was mentioned by Erdős and Moon (1965) and by Spencer (1985a). It turns out that several such constructions can be given. Let us describe one.

Let $p \equiv 3 \pmod{4}$ be a prime and let $T = T_p$ be the tournament whose vertices are all elements of the finite field $GF(p)$ in which (i, j) is a directed edge iff $i - j$ is a quadratic residue. [Since $p \equiv 3 \pmod{4}$, -1 is a quadratic nonresidue modulo p and hence T_p is a well-defined tournament.]

Theorem 9.1.1 For the tournaments T_p described above,

$$c(T_p) \leq \frac{1}{2} \binom{p}{2} + O(p^{3/2} \log p)$$

In order to prove this theorem we need some preparations. Let χ be the quadratic residue character defined on the elements of the finite field $GF(p)$ by $\chi(y) = y^{(p-1)/2}$. Equivalently, $\chi(y)$ is 1 if y is a nonzero square, 0 if y is 0 and -1 otherwise. Let $D = (d_{ij})_{i,j=0}^{p-1}$ be the p by p matrix defined by $d_{ij} = \chi(i - j)$.

Fact 1 For every two distinct j and l , $\sum_{i \in GF(p)} d_{ij} d_{il} = -1$.

Proof.

$$\begin{aligned} \sum_i d_{ij} d_{il} &= \sum_i \chi(i - j) \chi(i - l) = \sum_{i \neq j, l} \chi(i - j) \chi(i - l) \\ &= \sum_{i \neq j, l} \chi\left(\frac{i - j}{i - l}\right) = \sum_{i \neq j, l} \chi\left(1 + \frac{l - j}{i - l}\right). \end{aligned}$$

As i ranges over all elements of $GF(p)$ besides j and l the quantity $(1 + (l - j)/(i - l))$ ranges over all elements of $GF(p)$ besides 0 and 1. Since the sum of $\chi(r)$ over all r in $GF(p)$ is 0 this implies that the right-hand side of the last equation is $0 - \chi(0) - \chi(1) = -1$, completing the proof of the fact.

For two subsets A and B of $GF(p)$, let $e(A, B)$ denote the number of directed edges of T_p that start in a vertex of A and end in a vertex of B . By the definition of the matrix D it follows that

$$\sum_{i \in A} \sum_{j \in B} d_{ij} = e(A, B) - e(B, A)$$

The following lemma is proved in Alon (1986b).

Lemma 9.1.2 For any two subsets A and B of $GF(p)$,

$$\left| \sum_{i \in A} \sum_{j \in B} d_{ij} \right| \leq |A|^{1/2} |B|^{1/2} p^{1/2}$$

Proof. By the Cauchy-Schwarz Inequality and by the fact above,

$$\begin{aligned} \left(\sum_{i \in A} \sum_{j \in B} d_{ij} \right)^2 &\leq |A| \sum_{i \in A} \left(\sum_{j \in B} d_{ij} \right)^2 \\ &\leq |A| \sum_{i \in GF(p)} \left(\sum_{j \in B} d_{ij} \right)^2 \\ &= |A| \sum_{i \in GF(p)} \left(|B| + 2 \sum_{j < l \in B} d_{ij} d_{il} \right) \\ &= |A| |B| p + 2 |A| \sum_{j < l \in B} \sum_{i \in GF(p)} d_{ij} d_{il} \\ &\leq |A| |B| p, \end{aligned}$$

completing the proof of the lemma.

Proof [Theorem 9.1.1]. Let r be the smallest integer satisfying $2^r \geq p$. Let $\pi = \pi_1, \dots, \pi_p$ be an arbitrary permutation of the vertices of T_p , and define $\pi' = \pi_p, \dots, \pi_1$. We must show that $c(\pi, T_p) \leq \frac{1}{2} \binom{p}{2} + O(p^{3/2} \log p)$ or, equivalently, that $c(\pi, T_p) - c(\pi', T_p) \leq O(p^{3/2} \log p)$. Let a_1 and a_2 be two integers satisfying $p = a_1 + a_2$ and $a_1 \leq 2^{r-1}, a_2 \leq 2^{r-1}$. Let A_1 be the set of the first a_1 vertices in the permutation π and let A_2 be the set of the last a_2 vertices in π . By Lemma 9.1.2,

$$e(A_1, A_2) - e(A_2, A_1) \leq (a_1 a_2 p)^{1/2} \leq 2^{r-1} p^{1/2}$$

Next, let $a_{11}, a_{12}, a_{21}, a_{22}$ be integers each of which does not exceed 2^{r-2} such that $a_1 = a_{11} + a_{12}$ and $a_2 = a_{21} + a_{22}$. Let A_{11} be the subset of A_1 consisting of those a_{11} elements of A_1 that appear first in π , and let A_{12} be the set of the a_{12} remaining elements of A_1 . The partition of A_2 into the two sets A_{21} and A_{22} is defined similarly. By applying Lemma 9.1.2 we obtain

$$\begin{aligned} &e(A_{11}, A_{12}) - e(A_{12}, A_{11}) + e(A_{21}, A_{22}) - e(A_{22}, A_{21}) \\ &\leq (a_{11} a_{12} p)^{1/2} + (a_{21} a_{22} p)^{1/2} \\ &\leq 2 \cdot 2^{r-2} p^{1/2}. \end{aligned}$$

Continuing in the same manner we obtain, in the i th step, a partition of the set of vertices into 2^i blocks, each consisting of at most 2^{r-i} consecutive elements in the permutation π . This

partition is obtained by splitting each block in the partition corresponding to the previous step into two parts. By applying Lemma 9.1.2 to each such pair $A_{\epsilon_1}, A_{\epsilon_2}$ (where here ϵ is a vector of length $i - 1$ with $\{1, 2\}$ -entries), and by summing we conclude that the sum over all these 2^{i-1} vectors ϵ of the differences $e(A_{\epsilon_1}, A_{\epsilon_2}) - e(A_{\epsilon_2}, A_{\epsilon_1})$ does not exceed

$$2^{i-1}2^{r-i}p^{1/2} \leq 2^{r-1}p^{1/2}$$

Observe that the sum of the left-hand sides of all these inequalities as i ranges from 1 to r is precisely the difference $c(\pi, T_p) - c(\pi', T_p)$. Therefore by summing we obtain

$$c(\pi, T_p) - c(\pi', T_p) \leq 2^{r-1}p^{1/2}r = O(p^{3/2} \log p),$$

completing the proof. $\square$

We note that any antisymmetric matrix with $\{1, -1\}$ -entries in which each two columns are roughly orthogonal can be used to give a construction of a tournament as above. Some related results appear in Frankl, Rödl and Wilson (1988). The tournaments T_p , however, have stronger pseudorandom properties than do some of these other tournaments. For example, for every $k \leq \frac{1}{4} \log p$, and for every set S of k vertices of T_p , the number of vertices of T_p that beat all the members of S is $(1 + o(1))p/2^k$. This was proved by Graham and Spencer (1971) by applying Weil's famous theorem known as the Riemann hypotheses for curves over finite fields [Weil (1948)]. Taking a sufficiently large p this supplies an explicit construction for the Schütte problem mentioned in Chapter 1.

17.9.2 9.2 Eigenvalues and Expanders

A graph $G = (V, E)$ is called an (n, d, c) -expander if it has n vertices, the maximum degree of a vertex is d , and for every set of vertices $W \subset V$ of cardinality $|W| \leq n/2$, the inequality $|N(W)| \geq c|W|$ holds, where $N(W)$ denotes the set of all vertices in $V \setminus W$ adjacent to some vertex in W . We note that sometimes a slightly different definition is used, but the difference is not essential. Expanders share many of the properties of sparse random graphs and are the subject of an extensive literature. A family of linear expanders of density d and expansion c is a sequence $\{G_i\}_{i=1}^\infty$, where G_i is an (n_i, d, c) -expander and n_i tends to infinity as i tends to infinity.

Such a family is the main component of the parallel sorting network of Ajtai, Komlós and Szemerédi (1983) and can be used for constructing certain fault tolerant linear arrays. It also forms the basic building block used in the construction of graphs with special connectivity properties and small number of edges. Some other examples of the numerous applications of these graphs to various problems in theoretical computer science can be found, for example, in Alon (1986b) and its references.

It is not too difficult to prove the existence of a family of linear expanders using probabilistic arguments. This was first done by Pinsker (1973). An explicit construction is much more difficult to find and was first given by Margulis (1973). This construction was later improved by various authors; most known constructions are Cayley graphs of certain groups of matrices, and their expansion properties are proved by estimating the eigenvalues of the adjacency matrices of the graphs and by relying on the close correspondence between the expansion properties of a graph and its spectral properties. This correspondence was first studied, independently, by Tanner (1984) and by Alon and Milman (1984). Since it is somewhat simpler for the case of regular graphs we restrict our attention here to this case.

Let $G = (V, E)$ be a d -regular graph and let $A = A_G = (a_{uv})_{u,v \in V}$ be its adjacency matrix given by $a_{uv} = 1$ if $uv \in E$ and $a_{uv} = 0$ otherwise. Since G is d -regular the largest eigenvalue of A is d , corresponding to the all 1 eigenvector. Let $\lambda = \lambda(G)$ denote the second largest

eigenvalue of G . For two (not necessarily disjoint) subsets B and C of V let $e(B, C)$ denote the number of ordered pairs (u, v) , where $u \in B, v \in C$ and uv is an edge of G . (Note that if B and C are disjoint this is simply the number of edges of G that connect a vertex of B with a vertex of C .)

Theorem 9.2.1 For every partition of the set of vertices V into two disjoint subsets B and C ,

$$e(B, C) \geq \frac{(d - \lambda)|B||C|}{n}$$

Proof. Put $|V| = n, b = |B|$ and $c = |C| = n - b$. Let $D = dI$ be the n by n scalar matrix with the degree of regularity of G on its diagonal. Observe that for any real vector x of length n (considered as a function $x : V \rightarrow \mathbb{R}$) we have

$$\begin{aligned} \langle (D - A)x, x \rangle &= \sum_{u \in V} \left(d(x(u))^2 - \sum_{v: uv \in E} x(v)x(u) \right) \\ &= d \sum_{u \in V} (x(u))^2 - 2 \sum_{uv \in E} x(v)x(u) = \sum_{uv \in E} (x(v) - x(u))^2. \end{aligned}$$

Define, now, a vector x by $x(v) = -c$ if $v \in B$ and $x(v) = b$ if $v \in C$. Note that A and $D - A$ have the same eigenvectors and that the eigenvalues of $D - A$ are precisely $d - \mu$, as μ ranges over all eigenvalues of A . Note, also, that $\sum_{v \in V} x(v) = 0$; that is, x is orthogonal to the constant vector, which is the eigenvector of the smallest eigenvalue of $D - A$. Since $D - A$ is a symmetric matrix, its eigenvectors are orthogonal to each other and form a basis of the n -dimensional space. It follows that x is a linear combination of the other eigenvectors of $D - A$ and hence, by the definition of λ and the fact that $d - \lambda$ is the second smallest eigenvalue of $D - A$, we conclude that

$$\langle (D - A)x, x \rangle \geq (d - \lambda)\langle x, x \rangle = (d - \lambda)(bc^2 + cb^2) = (d - \lambda)bcn.$$

By the second paragraph of the proof the left-hand side of the last inequality is $\sum_{uv \in E} (x(u) - x(v))^2 = e(B, C) \cdot (b + c)^2 = e(B, C) \cdot n^2$. Thus

$$e(B, C) \geq \frac{(d - \lambda)bc}{n}$$

completing the proof.

Corollary 9.2.2 If λ is the second largest eigenvalue of a d -regular graph G with n vertices, then G is an (n, d, c) -expander for $c = (d - \lambda)/2d$.

Proof. Let W be a set of $w \leq n/2$ vertices of G . By Theorem 9.2.1 there are at least $(d - \lambda)w(n - w)/n \geq (d - \lambda)w/2$ edges from W to its complement. Since no vertex in the complement is adjacent to more than d of these edges it follows that $|N(W)| \geq (d - \lambda)w/2d$.

The estimate for c in the last corollary can in fact be improved to $2(d - \lambda)/(3d - 2\lambda)$, as shown by Alon and Milman (1984). Each of these estimates shows that if the second largest eigenvalue of G is far from the first, then G is a good expander. The converse of this is also true, although more complicated. This is given in the following result, proved in Alon (1986a), which we state without its proof.

Theorem 9.2.3 If G is a d -regular graph which is an (n, d, c) -expander then

$$\lambda(G) \leq d - \frac{c^2}{4 + 2c^2}$$

The last two results supply an efficient algorithm for approximating the expanding properties of a d -regular graph; we simply compute (or estimate) its second largest eigenvalue. The larger the difference between this eigenvalue and d is, the better expanding properties of G follow. It is thus natural to ask how far from d this second eigenvalue can be. It is known [see Nilli (1991)] that the second largest eigenvalue of any d -regular graph with diameter k is at least $2\sqrt{d-1}(1 - O(1/k))$. Therefore, in any infinite family of d -regular graphs, the limsup of the second largest eigenvalue is at least $2\sqrt{d-1}$. Lubotzky, Phillips and Sarnak (1986), and independently, Margulis (1988), gave, for every $d = p + 1$ where p is a prime congruent to 1 modulo 4, explicit constructions of infinite families of d -regular graphs G_i with second largest eigenvalues $\lambda(G_i) \leq 2\sqrt{d-1}$. These graphs are Cayley graphs of factor groups of the group of all two by two invertible matrices over a finite field, and their eigenvalues are estimated by applying results of Eichler and Igusa concerning the Ramanujan Conjecture. Eichler's proof relies on Weil's Theorem mentioned in the previous section. The nonbipartite graphs G constructed in this manner satisfy a somewhat stronger assertion than $\lambda(G) \leq 2\sqrt{d-1}$. In fact, besides their largest eigenvalue d , they do not have eigenvalues whose absolute value exceed $2\sqrt{d-1}$. This fact implies some strong pseudorandom properties, as shown in the next results.

Theorem 9.2.4 Let $G = (V, E)$ be a d -regular graph on n vertices, and suppose the absolute value of each of its eigenvalues but the first one is at most λ . For a vertex $v \in V$ and a subset B of V denote by $N(v)$ the set of all neighbors of v in G , and let $N_B(v) = N(v) \cap B$ denote the set of all neighbors of v in B . Then, for every subset B of cardinality bn of V ,

$$\sum_{v \in V} (|N_B(v)| - bd)^2 \leq \lambda^2 b(1-b)n.$$

Observe that in a random d -regular graph each vertex v would tend to have about bd neighbors in each set of size bn . The above theorem shows that if λ is much smaller than d then for most vertices v , $N_B(v)$ is not too far from bd .

Proof. Let A be the adjacency matrix of G and define a vector $f : V \rightarrow \mathbb{R}$ by $f(v) = 1 - b$ for $v \in B$ and $f(v) = -b$ for $v \notin B$. Clearly $\sum_{v \in V} f(v) = 0$; that is, f is orthogonal to the eigenvector of the largest eigenvalue of A . Therefore

$$\langle Af, Af \rangle \leq \lambda^2 \langle f, f \rangle$$

The right-hand side of the last inequality is $\lambda^2 (bn(1-b)^2 + (1-b)nb^2) = \lambda^2 b(1-b)n$. The left-hand side is

$$\sum_{v \in V} ((1-b)|N_B(v)| - b(d - |N_B(v)|))^2 = \sum_{v \in V} (|N_B(v)| - bd)^2.$$

The desired result follows.

Corollary 9.2.5 Let $G = (V, E)$, d, n and λ be as in Theorem 9.2.4. Then for every two sets of vertices B and C of G , where $|B| = bn$ and $|C| = cn$, we have

$$|e(B, C) - cbdn| \leq \lambda \sqrt{bcn}$$

Proof. By Theorem 9.2.4,

$$\sum_{v \in C} (|N_B(v)| - bd)^2 \leq \sum_{v \in V} (|N_B(v)| - bd)^2 \leq \lambda^2 b(1-b)n.$$

Thus, by the Cauchy-Schwarz Inequality,

$$\begin{aligned}
|e(B, C) - cbdn| &\leq \sum_{v \in C} |N_B(v) - bd| \leq \sqrt{cn} \left(\sum_{v \in C} (|N_B(v)| - bd)^2 \right)^{1/2} \\
&\leq \sqrt{cn} \lambda \sqrt{b(1-b)n} \leq \lambda \sqrt{bcn}
\end{aligned}$$

The special case $B = C$ gives the following result. A slightly stronger estimate is proved in a similar way in Alon and Chung (1988).

Corollary 9.2.6 Let $G = (V, E)$, d, n and λ be as in Theorem 9.2.4. Let B be an arbitrary set of vertices of G and let $e(B) = \frac{1}{2}e(B, B)$ be the number of edges in the induced subgraph of G on B . Then

$$\left| e(B) - \frac{1}{2}b^2dn \right| \leq \frac{1}{2}\lambda bn.$$

A walk of length l in a graph G is a sequence $v_0, \dots, v_l$ of vertices of G , where for each $1 \leq i \leq l$, $v_{i-1}v_i$ is an edge of G . Obviously, the total number of walks of length l in a d -regular graph on n vertices is precisely $n \cdot d^l$. Suppose, now, that C is a subset of, say, $n/2$ vertices of G . How many of these walks do not contain any vertex of C ? If G is disconnected it may happen that half of these walks avoid C .

However, as shown by Ajtai, Komlós and Szemerédi (1987), there are many fewer such walks if all the eigenvalues of G but the largest are small. This result and some of its extensions have several applications in theoretical computer science, as shown in the above-mentioned paper [see also Cohen and Wigderson (1989)]. We conclude this section by stating and proving the result and one of its applications.

Theorem 9.2.7 Let $G = (V, E)$ be a d -regular graph on n vertices, and suppose that each of its eigenvalues but the first one is at most λ . Let C be a set of cn vertices of G . Then, for every l , the number of walks of length l in G that avoid C does not exceed $(1-c)n((1-c)d + c\lambda)^l$.

Proof. Let A be the adjacency matrix of G and let A' be the adjacency matrix of its induced subgraph on the complement of C . We claim that the maximum eigenvalue of A' is at most $(1-c)d + c\lambda$. To prove this claim we must show that for every vector $f : V \rightarrow \mathbb{R}$ satisfying $f(v) = 0$ for each $v \in C$ and $\sum_{v \in V} f(v)^2 = 1$, the inequality $\langle Af, f \rangle \leq (1-c)d + c\lambda$ holds. Let $f_1, f_2, \dots, f_n$ be an orthonormal basis of eigenvectors of A , where f_i is the eigenvector of λ_i , $\lambda_1 = d$ and each entry of f_1 is $1/\sqrt{n}$. Then $f = \sum_{i=1}^n c_i f_i$, where $\sum_{i=1}^n c_i^2 = 1$ and

$$\begin{aligned}
c_1 &= \sum_{v \in V} \frac{f(v)}{\sqrt{n}} = \sum_{v \in V \setminus C} \frac{f(v)}{\sqrt{n}} \\
&\leq \left(\sum_{v \in V \setminus C} f(v)^2 \right)^{1/2} \left((1-c)n \frac{1}{n} \right)^{1/2} = \sqrt{1-c},
\end{aligned}$$

where here we used the Cauchy-Schwarz Inequality. Therefore $\sum_{i=2}^n c_i^2 \geq c$ and

$$\langle Af, f \rangle = \sum_{i=1}^n c_i^2 \lambda_i \leq (1-c)d + c\lambda,$$

supplying the desired estimate for the largest eigenvalue of A' .

Let $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_m$ be the eigenvalues of A' , where $m = (1-c)n$. By the Perron-Frobenius Theorem it follows that the absolute value of each of them is at most $\gamma_1 \leq (1-c)d + c\lambda$. The total number of walks of length l that avoid C is precisely $(A'^l g, g)$, where g is the all 1-vector indexed by the vertices in $V - C$. By expressing g as a linear combination of the eigenvectors of A' , $g = \sum_{i=1}^m b_i g_i$, where g_i is the eigenvector of γ_i , we conclude that this number is precisely

$$\sum_{i=1}^m b_i^2 \gamma_i^l \leq \gamma_1^l \sum_{i=1}^m b_i^2 = m \gamma_1^l \leq m((1-c)d + c\lambda)^l$$

Substituting $m = (1-c)n$ the desired result follows.

A randomly chosen walk of length l in a graph G is a walk of length l in G chosen according to a uniform distribution among all walks of that length. Note that if G is d -regular such a walk can be chosen by choosing randomly its starting point v_0 , and then by choosing, for each $1 \leq i \leq l$, v_i randomly among the d neighbors of v_{i-1} .

Corollary 9.2.8 Let $G = (V, E)$, d, n, λ, C and c be as in Theorem 9.2.7 and suppose

$$(1-c)d + c\lambda \leq \frac{d}{\sqrt{2}}$$

Then, for every l , the probability that a randomly chosen walk of length l in G avoids C is at most $2^{-l/2}$.

Proof. The number of walks of length l in G that avoid C is at most $(1-c)n((1-c)d + c\lambda)^l \leq nd^l 2^{-l/2}$, by Theorem 9.2.7. Since the total number of walks is nd^l , the desired result follows.

The results above are useful for amplification of probabilities in randomized algorithms. Although such an amplification can be achieved for any Monte Carlo algorithm, we prefer, for simplicity, to consider one representative example: the primality testing algorithm of Rabin (1980).

For an odd integer q , define two integers a and b by $q-1 = 2^a b$, where b is odd. An integer x , $1 \leq x \leq q-1$, is called a witness (for the nonprimality of q) if for the sequence $x_0, \dots, x_a$ defined by $x_0 = x^b \pmod{q}$ and $x_i = x_{i-1}^2 \pmod{q}$ for $1 \leq i \leq a$ either $x_a \neq 1$ or there is an i such that $x_i \neq -1, 1$ and $x_{i+1} = 1$. One can show that if q is a prime then there are no such witnesses for q , whereas if q is an odd nonprime then at least half of the numbers between 1 and $q-1$ are witnesses for q . (In fact, at least $3/4$ are witnesses, as shown by Rabin.) This suggests the following randomized algorithm for testing if an odd integer q is a prime (for even integers there is a simpler algorithm!).

Choose, randomly, an integer x between 1 and $q-1$ and check if it is a witness. If it is, report that q is not a prime. Otherwise, report that q is a prime.

Observe that if q is a prime, the algorithm certainly reports it is a prime, whereas if q is not a prime, the probability that the algorithm makes a mistake and reports it as a prime is at most $1/2$. What if we wish to reduce the probability of making such a mistake? Clearly we can simply repeat the algorithm. If we repeat it l independent times, then the probability of making an error (i.e., reporting a nonprime as a prime) decreases to $1/2^l$. However, the number of random bits required for this procedure is $l \cdot \log(q-1)$.

Suppose we wish to use fewer random bits. By applying the properties of a randomly chosen walk on an appropriate graph, proved in the last two results, we can obtain the same estimate for the error probability by using only $\log(q-1) + O(l)$ random bits. This is done as follows.

Let G be a d -regular graph with $q-1$ vertices, labeled by all integers between 1 and $q-1$. Suppose G has no eigenvalue, but the first one, that exceeds λ and suppose that

$$\frac{d + \lambda}{2} \leq \frac{d}{\sqrt{2}} \tag{9.1}$$

Now choose randomly a walk of length $2l$ in the graph G , and check, for each of the numbers labeling its vertices, if it is a witness. If q is a nonprime, then at least half of the vertices of G are labeled by witnesses. Hence, by Corollary 9.2.8 and by (9.1), the probability that no witness is on the walk is at most $2^{-2l/2} = 2^{-l}$. Thus we obtain the same reduction in the error-probability as the one obtained by choosing l independent witnesses. Let us estimate the number of random bits required for choosing such a random walk.

The known constructions of expanders given by Lubotzky et al. (1986) or by Margulis (1988) give explicit families of graphs with degree d and with $\lambda \leq 2\sqrt{d-1}$, for each $d = p+1$, where p is a prime congruent to 1 modulo 4. [We note that these graphs will not have exactly $q-1$ vertices but this does not cause any real problem as we can take a graph with n vertices, where $q-1 \leq n \leq (1+o(1))(q-1)$, and label its i th vertex by $i \pmod{q-1}$. In this case the number of vertices labeled by witnesses would still be at least $(\frac{1}{2} + o(1))n$.] One can easily check that, for example, $d = 30$ and $\lambda = 2\sqrt{29}$ satisfy (9.1) and thus we can use a 30-regular graph. The number of random bits required for choosing a random walk of length $2l$ in it is less than $\log(q-1) + 10l + 1$, much less than the $l \log(q-1)$ bits that are needed in the repetition procedure.

17.9.3 9.3 Quasirandom Graphs

In this section we describe several pseudorandom properties of graphs, which, somewhat surprisingly, turn out to be all equivalent. All the properties are ones satisfied, almost surely, by a random graph in which every edge is chosen, independently, with probability $1/2$. The equivalence between some of these properties was first proved by several authors; see Thomason (1987), Frankl et al. (1988) and Alon and Chung (1988), but the first paper in which all of them (and some others) appear is the one by Chung, Graham and Wilson (1989). Our presentation here follows that paper, although, in order to simplify the presentation, we consider only the case of regular graphs.

We first need some notation. For two graphs G and H , let $N_G^*(H)$ be the number of labeled occurrences of H as an induced subgraph of G ; that is, the number of adjacency preserving injections $f : V(H) \rightarrow V(G)$ whose image is the set of vertices of an induced copy of H in G . Similarly, $N_G(H)$ denotes the number of labeled copies of H as a (not necessarily induced) subgraph of G . Note that $N_G(H) = \sum_L N_G^*(L)$, where L ranges over all graphs on the set of vertices of H obtained from H by adding to it a (possibly empty) set of edges.

Throughout this section G always denotes a graph with n vertices. We denote the eigenvalues of its adjacency matrix (taken with multiplicities) by $\lambda_1, \dots, \lambda_n$, where $|\lambda_1| \geq \dots \geq |\lambda_n|$. [Since we consider in this section only the eigenvalues of G we simply write λ_1 and not $\lambda_1(G)$.] Recall also the following notation, used in the previous section: for a vertex v of G , $N(v)$ denotes the set of its neighbors in G . If S is a set of vertices of G , $e(S)$ denotes the number of edges in the induced subgraph of G on S . If B and C are two (not necessarily disjoint) subsets of vertices of G ,

$e(B, C)$ denotes the number of ordered pairs (b, c) , where $b \in B, c \in C$ and bc is an edge of G . Thus $e(S) = \frac{1}{2}e(S, S)$.

We can now state the pseudorandom properties considered here. All the properties refer to a graph $G = (V, E)$ with n vertices. Throughout the section, we use the $o(\cdot)$ notation, without mentioning the precise behavior of each $o(\cdot)$. Thus occurrences of two $o(1)$, say, need not mean that both are identical and only mean that if we consider a family of graphs G and let their number of vertices n tend to infinity then each $o(1)$ tends to 0.

Property $P_1(s)$: For every graph $H(s)$ on s vertices

$$N_G^*(H(s)) = (1 + o(1))n^s 2^{-\binom{s}{2}}.$$

Property P_2 : For the cycle $C(4)$ with 4 vertices $N_G(C(4)) \leq (1 + o(1))(n/2)^4$.

Property P_3 : $|\lambda_2| = o(n)$.

Property P_4 : For every set S of vertices of G $e(S) = \frac{1}{4}|S|^2 + o(n^2)$.

Property P_5 : For every two sets of vertices B and C $e(B, C) = \frac{1}{2}|B||C| + o(n^2)$.

Property P_6 : $\sum_{u, v \in V} ||N(u) \cap N(v)| - n/4| = o(n^3)$.

It is easy to check that all the properties above are satisfied, almost surely, by a random graph

on n vertices. In this section we show that all these properties are equivalent for a regular graph with n vertices and degree of regularity about $n/2$. The fact that the innocent-looking property P_2 is strong enough to imply for such graphs $P_1(s)$ for every $s \geq 1$ is one of the interesting special cases of this result.

Graphs that satisfy any (and thus all) of the properties above are called quasirandom. As noted above the assumption that G is regular can be dropped (at the expense of slightly modifying property P_2 and slightly complicating the proofs).

Theorem 9.3.1 Let G be a d -regular graph on n vertices, where $d = (\frac{1}{2} + o(1))n$. If G satisfies any one of the seven properties $P_1(4), P_1(s)$ for all $s \geq 1, P_2, P_3, P_4, P_5, P_6$ then it satisfies all seven.

Proof. We show that

$$P_1(4) \Rightarrow P_2 \Rightarrow P_3 \Rightarrow P_4 \Rightarrow P_5 \Rightarrow P_6 \Rightarrow P_1(s) \text{ for all } s \geq 1 \quad (\Rightarrow P_1(4)).$$

1. $P_1(4) \Rightarrow P_2$.

Suppose G satisfies $P_1(4)$. Then $N_G(C(4)) = \sum_L N_G^*(L)$, as L ranges over the four labeled graphs obtained from a labeled $C(4)$ by adding to it a (possibly empty) set of edges. Since G satisfies $P_1(4)$, $N_G^*(L) = (1 + o(1))n^4 2^{-16}$ for each of these graphs L and hence $N_G(C(4)) = (1 + o(1))n^4 2^{-4}$, showing that G satisfies P_2 .

2. $P_2 \Rightarrow P_3$.

Suppose G satisfies P_2 and let A be its adjacency matrix. The trace of A^4 is precisely $\sum_{i=1}^n \lambda_i^4$. On the other hand it is easy to see that this trace is precisely the number of (labeled) closed walks of length 4 in G , that is, the number of sequences $v_0, v_1, v_2, v_3, v_4 = v_0$ of vertices of G such that $v_i v_{i+1}$ is an edge for each $0 \leq i \leq 3$. This number is $N_G(C(4))$ plus the number of such sequences in which $v_2 = v_0$, which is nd^2 , plus the number of such sequences in which $v_2 \neq v_0$ and $v_3 = v_1$, which is $nd(d-1)$. Thus

$$\begin{aligned} \sum_{i=1}^n \lambda_i^4 &= d^4 + \sum_{i=2}^n \lambda_i^4 = (1 + o(1))(n/2)^4 + \sum_{i=2}^n \lambda_i^4 \\ &= N_G(C(4)) + O(n^3) = (1 + o(1))(n/2)^4 \end{aligned}$$

It follows that $\sum_{i=2}^n \lambda_i^4 = o(n^4)$ and hence that $|\lambda_2| = o(n)$, as needed.

3. $P_3 \Rightarrow P_4$.

This is an immediate consequence of Corollary 9.2.6.

4. $P_4 \Rightarrow P_5$.

Suppose G satisfies P_4 . We first claim that it satisfies property P_5 for disjoint sets of vertices B and C . Indeed, if B and C are disjoint then

$$\begin{aligned} e(B, C) &= e(B \cup C) - e(B) - e(C) \\ &= \frac{1}{4}(|B| + |C|)^2 - \frac{1}{4}|B|^2 - \frac{1}{4}|C|^2 + o(n^2) \\ &= \frac{1}{2}|B||C| + o(n^2), \end{aligned}$$

proving the claim.

In case B and C are not disjoint we have

$$e(B, C) = e(B \setminus C, C \setminus B) + e(B \cap C, C \setminus B) + e(B \cap C, B \setminus C) + 2e(B \cap C).$$

Put $|B| = b, |C| = c$ and $|B \cap C| = x$. By the above expression for $e(B, C)$ and by the fact that G satisfies P_4 and P_5 for disjoint B and C we get

$$\begin{aligned}
e(B, C) &= \frac{1}{2}(b-x)(c-x) + \frac{1}{2}x(c-x) + \frac{1}{2}x(b-x) + \frac{2}{4}x^2 + o(n^2) \\
&= \frac{1}{2}bc + o(n^2) = \frac{1}{2}|B||C| + o(n^2)
\end{aligned}$$

showing that G satisfies P_5 .

5. $P_5 \Rightarrow P_6$.

Suppose that G satisfies P_5 and recall that G is d -regular, where $d = (\frac{1}{2} + o(1))n$. Let v be a fixed vertex of G , and let us estimate the sum

$$\sum_{u \in V, u \neq v} \left| |N(u) \cap N(v)| - \frac{n}{4} \right|.$$

Define

$$B_1 = \left\{ u \in V, u \neq v : |N(u) \cap N(v)| \geq \frac{n}{4} \right\}$$

and similarly

$$B_2 = \left\{ u \in V, u \neq v : |N(u) \cap N(v)| < \frac{n}{4} \right\}.$$

Let C be the set of all neighbors of v in G . Observe that

$$\begin{aligned}
\sum_{u \in B_1} \left| |N(u) \cap N(v)| - \frac{n}{4} \right| &= \sum_{u \in B_1} |N(u) \cap N(v)| - |B_1| \frac{n}{4} \\
&= e(B_1, C) - |B_1| \frac{n}{4}.
\end{aligned}$$

Since G satisfies P_5 and since $d = (\frac{1}{2} + o(1))n$ the last difference is $\frac{1}{2}|B_1|d + o(n^2) - |B_1|n/4 = o(n^2)$.

A similar argument implies that

$$\sum_{u \in B_2} \left| |N(u) \cap N(v)| - \frac{n}{4} \right| = o(n^2).$$

It follows that for every vertex v of G ,

$$\sum_{u \in V, u \neq v} \left| |N(u) \cap N(v)| - \frac{n}{4} \right| = o(n^2),$$

and by summing over all vertices v we conclude that G satisfies property P_6 .

6. $P_6 \Rightarrow P_1(s)$ for all $s \geq 1$.

Suppose $G = (V, E)$ satisfies P_6 . For any two distinct vertices u and v of G let $a(u, v)$ be 1 if $uv \in E$ and 0 otherwise. Also, define $s(u, v) = |\{w \in V : a(u, w) = a(v, w)\}|$. Since G is $d = (\frac{1}{2} + o(1))n$ -regular,

$$s(u, v) = 2|N(u) \cap N(v)| + n - 2d = 2|N(u) \cap N(v)| + o(n).$$

Therefore the fact that G satisfies P_6 implies that

$$\sum_{u, v \in V} \left| s(u, v) - \frac{n}{2} \right| = o(n^3) \quad (9.2)$$

Let $H = H(s)$ be an arbitrary fixed graph on s vertices, and put $N_s = N_G^*(H(s))$. We must show that

$$N_s = (1 + o(1))n^s 2^{-\binom{s}{2}}$$

Denote the vertex set of $H(s)$ by $\{v_1, \dots, v_s\}$. For each $1 \leq r \leq s$, put $V_r = \{v_1, \dots, v_r\}$, and let $H(r)$ be the induced subgraph of H on V_r . We prove, by induction on r , that for $N_r = N_G^*(H(r))$,

$$N_r = (1 + o(1))n_{(r)} 2^{-\binom{r}{2}} \quad (9.3)$$

where $n_{(r)} = n(n-1) \cdots (n-r+1)$.

This is trivial for $r = 1$. Assuming it holds for r , where $1 \leq r < s$, we prove it for $r+1$. For a vector $\alpha = (\alpha_1, \dots, \alpha_r)$ of distinct vertices of G , and for a vector $\epsilon = (\epsilon_1, \dots, \epsilon_r)$ of $(0, 1)$ -entries, define

$$f_r(\alpha, \epsilon) = |\{v \in V : v \neq \alpha_1, \dots, \alpha_r \text{ and } a(v, \alpha_j) = \epsilon_j \text{ for all } 1 \leq j \leq r\}|.$$

Clearly N_{r+1} is the sum of the N_r quantities $f_r(\alpha, \epsilon)$ in which $\epsilon_j = a(v_{r+1}, v_j)$ and α ranges over all N_r induced copies of $H(r)$ in G .

Observe that altogether there are precisely $n_{(r)} 2^r$ quantities $f_r(\alpha, \epsilon)$. It is convenient to view $f_r(\alpha, \epsilon)$ as a random variable defined on a sample space of $n_{(r)} 2^r$ points, each having an equal probability. To complete the proof we compute the expectation and the variance of this random variable. We show that the variance is so small that most of the quantities $f_r(\alpha, \epsilon)$ are very close to the expectation, and thus obtain a sufficiently accurate estimate for N_{r+1} which is the sum of N_r such quantities.

We start with the simple computation of the expectation $E[f_r]$ of $f_r(\alpha, \epsilon)$. We have

$$\begin{aligned} E[f_r] &= \frac{1}{n_{(r)} 2^r} \sum_{\alpha, \epsilon} f_r(\alpha, \epsilon) = \frac{1}{n_{(r)} 2^r} \sum_{\alpha} \sum_{\epsilon} f_r(\alpha, \epsilon) \\ &= \frac{1}{n_{(r)} 2^r} \sum_{\alpha} (n-r) = \frac{n-r}{2^r} \end{aligned}$$

where we used the fact that every vertex $v \neq \alpha_1, \dots, \alpha_r$ defines ϵ uniquely. Next, we estimate the quantity S_r defined by

$$S_r = \sum_{\alpha, \epsilon} f_r(\alpha, \epsilon) (f_r(\alpha, \epsilon) - 1)$$

We claim that

$$S_r = \sum_{u \neq v} s(u, v)_{(r)} \quad (9.4)$$

To prove this claim, observe that S_r can be interpreted as the number of ordered triples $(\alpha, \epsilon, (u, v))$, where $\alpha = (\alpha_1, \dots, \alpha_r)$ is an ordered set of r distinct vertices of G , $\epsilon = (\epsilon_1, \dots, \epsilon_r)$ is a binary vector of length r , and u, v is an ordered pair of additional vertices of G so that

$$a(u, \alpha_k) = a(v, \alpha_k) = \epsilon_k \text{ for all } 1 \leq k \leq r.$$

For each fixed α and ϵ , there are precisely $f_r(\alpha, \epsilon) (f_r(\alpha, \epsilon) - 1)$ choices for the pair (u, v) and hence S_r counts the number of these triples.

Now, let us compute this number by first choosing u and v . Once u, v are chosen, the additional vertices $\alpha_1, \dots, \alpha_r$ must all belong to the set $\{w \in V : a(u, w) = a(v, w)\}$. Since the cardinality of this set is $s(u, v)$ it follows that there are $s(u, v)_{(r)}$ choices for $\alpha_1, \dots, \alpha_r$. Once these are chosen the vector ϵ is determined and thus (9.4) follows.

We next claim that (9.2) implies

$$\sum_{u \neq v} s(u, v)_{(r)} = (1 + o(1))n^{r+2}2^{-r} \quad (9.5)$$

To prove this claim define $\epsilon_{uv} = s(u, v) - n/2$. Observe that, by (9.2), $\sum_{u \neq v} |\epsilon_{uv}| = o(n^3)$ and $|\epsilon_{uv}| \leq n/2 \leq n$ for each u, v . Hence, for every fixed $a \geq 1$,

$$\sum_{u \neq v} |\epsilon_{uv}|^a \leq n^{a-1} \sum_{u \neq v} |\epsilon_{uv}| = o(n^{a+2}).$$

This implies that

$$\begin{aligned} & \sum_{u \neq v} s(u, v)_{(r)} \\ &= \sum_{u \neq v} \left(\frac{n}{2} + \epsilon_{uv} \right)_{(r)} \\ &= \sum_{k=0}^r \sum_{u \neq v} c_k \left(\frac{n}{2} \right)^k \epsilon_{uv}^{r-k} \quad (\text{for appropriate constants } c_k) \\ &= \left(\frac{n}{2} \right)^r n_{(2)} + \sum_{k=0}^{r-1} \sum_{u \neq v} c_k \left(\frac{n}{2} \right)^k \epsilon_{uv}^{r-k} \\ &\leq \left(\frac{n}{2} \right)^r n_{(2)} + \sum_{k=0}^{r-1} \sum_{u \neq v} |c_k| n^k |\epsilon_{uv}|^{r-k} \\ &\leq n^{r+2}2^{-r} + c \sum_{k=0}^{r-1} n^k \sum_{u \neq v} |\epsilon_{uv}|^{r-k} \quad (\text{for an appropriate constant } c) \\ &\leq n^{r+2}2^{-r} + c \sum_{k=0}^{r-1} n^k \cdot o(n^{r-k+2}) \\ &= n^{r+2}2^{-r}(1 + o(1)), \end{aligned}$$

implying (9.5).

By (9.4) and (9.5), $S_r = (1 + o(1))n^{r+2}2^{-r}$. Therefore

$$\begin{aligned} & \sum_{\alpha, \epsilon} (f_r(\alpha, \epsilon) - \mathbb{E}[f_r])^2 \\ &= \sum_{\alpha, \epsilon} f_r^2(\alpha, \epsilon) - \sum_{\alpha, \epsilon} \mathbb{E}[f_r]^2, \\ & \sum_{\alpha, \epsilon} (f_r^2(\alpha, \epsilon) - f_r(\alpha, \epsilon)) + \sum_{\alpha, \epsilon} f_r(\alpha, \epsilon) - n_{(r)}2^r(n-r)^22^{-2r} \\ &= S_r + n_{(r)}2^r \mathbb{E}[f_r] - n_{(r)}2^r(n-r)^22^{-2r} \\ &= S_r + n_{(r+1)} - n_{(r)}2^r(n-r)^22^{-2r} = o(n^{r+2}) \end{aligned}$$

Recall that N_{r+1} is the summation of N_r quantities of the form $f_r(\alpha, \epsilon)$. Thus

$$|N_{r+1} - N_r \mathbb{E}[f_r]|^2 = \left| \sum_{N_r \text{ terms}} (f_r(\alpha, \epsilon) - \mathbb{E}[f_r]) \right|^2.$$

By Cauchy-Schwarz, the last expression is at most

$$\begin{aligned}
N_r \sum_{N_r \text{ terms}} (f_r(\alpha, \epsilon) - E[f_r])^2 &\leq N_r \sum_{\alpha, \epsilon} (f_r(\alpha, \epsilon) - E[f_r])^2 \\
&= N_r \cdot o(n^{r+2}) = o(n^{2r+2})
\end{aligned}$$

It follows that

$$|N_{r+1} - N_r E[f_r]| = o(n^{r+1})$$

and hence, by the induction hypothesis,

$$\begin{aligned}
N_{r+1} &= N_r E[f_r] + o(n^{r+1}) \\
&= (1 + o(1))n_{(r)}2^{-\binom{r}{2}} \cdot (n - r)2^{-r} + o(n^{r+1}) \\
&= (1 + o(1))n_{(r+1)}2^{-\binom{r+1}{2}}
\end{aligned}$$

This completes the proof of the induction step and establishes Theorem 9.3.1.

There are many examples of families of quasirandom graphs. The most widely used is probably the family of Paley graphs G_p defined as follows. For a prime p congruent to 1 modulo 4, let G_p be the graph whose vertices are the integers $0, 1, 2, \dots, p-1$ in which i and j are adjacent if and only if $i - j$ is a quadratic residue modulo p . The graphs G_p , which are the undirected analogues of the quadratic residue tournaments discussed in Section 9.1, are $(p-1)/2$ -regular. For any two distinct vertices i and j of G_p , the number of vertices k that are either adjacent to both i and j or nonadjacent to both is precisely the number of times the quotient $(k-i)/(k-j)$ is a quadratic residue. As k ranges over all numbers between 0 and $p-1$ but i and j , this quotient ranges over all numbers but 1 and 0 and hence it is a quadratic residue precisely $\frac{1}{2}(p-1) - 1$ times. (This is essentially the same assertion as that of the first fact given in the proof of Theorem 9.1.1.) We have thus shown that for every two vertices i and j of G_p , $s(i, j) = (p-3)/2$, and this, together with the fact that G_p is $(p-1)/2$ -regular, easily implies that it satisfies Property P_6 . Therefore it is quasirandom. As is the case with the quadratic residue tournaments, G_p satisfies, in fact, some stronger pseudorandom properties that are not satisfied by every quasirandom graph and that can be proved by applying Weil's Theorem.

17.9.4 9.4 Exercises

1. By considering a random bipartite three-regular graph on $2n$ vertices obtained by picking three random permutations between the two color classes, prove that there is a $c > 0$ such that for every n there exists a $(2n, 3, c)$ -expander.
2. Let $G = (V, E)$ be an (n, d, λ) -graph, suppose n is divisible by k , and let $C : V \rightarrow \{1, 2, \dots, k\}$ be a coloring of V by k colors, so that each color appears precisely n/k times. Prove that there is a vertex of G which has a neighbor of each of the k colors, provided $k\lambda \leq d$.
3. Let $G = (V, E)$ be a graph in which there is at least one edge between any two disjoint sets of size $a+1$. Prove that for every set Y of $5a$ vertices, there is a set X of at most a vertices, such that for every set Z satisfying $Z \cap (X \cup Y) = \emptyset$ and $|Z| \leq a$, the inequality $|N(Z) \cap Y| \geq 2|Z|$ holds.
4. Prove that for every $\epsilon > 0$ there exists an $n_0 = n_0(\epsilon)$ so that for every $(n, n/2, 2\sqrt{n})$ -graph $G = (V, E)$ with $n > n_0$, the number of triangles M in G satisfies $|M - n^3/48| \leq \epsilon n^3$.

The Probabilistic Lens: Random Walks

A vertex-transitive graph is a graph $G = (V, E)$ such that for any two vertices $u, v \in V$ there is an automorphism of G that maps u into v . A random walk of length l in G starting at a vertex v is a randomly chosen sequence $v = v_0, v_1, \dots, v_l$, where each v_{i+1} is chosen, randomly and independently, among the neighbors of v_i ($0 \leq i < l$).

The following theorem states that for every vertex-transitive graph G , the probability that a random walk of even length in G ends at its starting point is at least as big as the probability that it ends at any other vertex. Note that the proof requires almost no computation. We note also that the result does not hold for general regular graphs, and the vertex transitivity assumption is necessary.

Theorem 1 Let $G = (V, E)$ be a vertex-transitive graph. For an integer k and for two (not necessarily distinct) vertices u, v of G , let $P^k(u, v)$ denote the probability that a random walk of length k starting at u ends at v . Then, for every integer k and for every two vertices $u, v \in V$,

$$P^{2k}(u, u) \geq P^{2k}(u, v)$$

Proof. We need the following simple inequality, sometimes attributed to Chebyshev.

Claim 9.4.1 For every sequence $(a_1, \dots, a_n)$ of n reals and for any permutation π of $\{1, \dots, n\}$,

$$\sum_{i=1}^n a_i a_{\pi(i)} \leq \sum_{i=1}^n a_i^2$$

Proof. The inequality follows immediately from the fact that

$$\sum_{i=1}^n a_i^2 - \sum_{i=1}^n a_i a_{\pi(i)} = \frac{1}{2} \sum_{i=1}^n (a_i - a_{\pi(i)})^2 \geq 0.$$

Consider, now, a random walk of length $2k$ starting at u . By summing over all the possibilities of the vertex the walk reaches after k steps we conclude that for every vertex v ,

$$P^{2k}(u, v) = \sum_{w \in V} P^k(u, w) P^k(w, v) = \sum_{w \in V} P^k(u, w) P^k(v, w) \quad (1)$$

where the last equality follows from the fact that G is an undirected regular graph. Since G is vertex-transitive, the two vectors $(P^k(u, w))_{w \in V}$ and $(P^k(v, w))_{w \in V}$ can be obtained from each other by permuting the coordinates. Therefore, by the claim above, the maximum possible value of the sum in the right-hand side of (1) is when $u = v$, completing the proof of the theorem.

18 Part II Topics by Alon, Spencer

18.1 10 Random Graphs

It is six in the morning. The house is asleep. Nice music is playing. I prove and conjecture.

- Paul Erdős, in a letter to Vera Sós

Let n be a positive integer, $0 \leq p \leq 1$. The random graph $G(n, p)$ is a probability space over the set of graphs on the vertex set $\{1, \dots, n\}$ determined by

$$\Pr[\{i, j\} \in G] = p$$

with these events mutually independent. This model is often used in the probabilistic method for proving the existence of certain graphs. In this chapter we study the properties of $G(n, p)$ for their own sake.

Random graphs is an active area of research that combines probability theory and graph theory. The subject began in 1960 with the monumental paper On the Evolution of Random Graphs by Paul Erdős and Alfred Rényi. The book Random Graphs by Bollobás (2001) is the standard source for the field. Another book, also entitled Random Graphs by Janson, Łuczak and Ruciński (2000) is also excellent. In this chapter we explore only a few of the many topics in this fascinating area.

There is a compelling dynamic model for random graphs. For all pairs i, j let $x_{i,j}$ be selected uniformly from $[0, 1]$, the choices mutually independent. Imagine p going from 0 to 1. Originally, all potential edges are "off." The edge from i to j (which we may imagine as a neon light) is turned on when p reaches $x_{i,j}$ and then stays on. At $p = 1$ all edges are "on." At time p the graph of all "on" edges has distribution $G(n, p)$. As p increases $G(n, p)$ evolves from empty to full.

In their original paper, Erdős and Rényi let $G(n, e)$ be the random graph with n vertices and precisely e edges. Again there is a dynamic model: Begin with no edges and add edges randomly one by one until the graph becomes full. Generally $G(n, e)$ will have very similar properties as $G(n, p)$ with $p \sim e/\binom{n}{2}$. We will work on the probability model exclusively.

18.1.1 10.1 Subgraphs

The term "the random graph" is, strictly speaking, a misnomer. $G(n, p)$ is a probability space over graphs. Given any graph theoretic property A there will be a probability that $G(n, p)$ satisfies A , which we write $\Pr[G(n, p) \models A]$. When A is monotone $\Pr[G(n, p) \models A]$ is a monotone function of p . As an instructive example, let A be the event " G is triangle free." Let X be the number of triangles contained in $G(n, p)$. Linearity of expectation gives

$$\mathbb{E}[X] = \binom{n}{3} p^3$$

This suggests the parametrization $p = c/n$. Then

$$\lim_{n \rightarrow \infty} \mathbb{E}[X] = \lim_{n \rightarrow \infty} \binom{n}{3} p^3 = \frac{c^3}{6}$$

It turns out that the distribution of X is asymptotically Poisson. In particular,

$$\lim_{n \rightarrow \infty} \Pr[G(n, p) \models A] = \lim_{n \rightarrow \infty} \Pr[X = 0] = e^{-c^3/6}$$

Note that

$$\begin{aligned} \lim_{c \rightarrow 0} e^{-c^3/6} &= 1 \\ \lim_{c \rightarrow \infty} e^{-c^3/6} &= 0 \end{aligned}$$

When $p = 10^{-6}/n$, $G(n, p)$ is very unlikely to have triangles and when $p = 10^6/n$, $G(n, p)$ is very likely to have triangles. In the dynamic view the first triangles almost always appear at $p = \Theta(1/n)$. If we take a function such as $p(n) = n^{-0.9}$ with $p(n) \gg n^{-1}$ then $G(n, p)$ will almost always have triangles. Occasionally we will abuse notation and say, for example, that $G(n, n^{-0.9})$ contains a triangle - this meaning that the probability that it contains a triangle approaches 1 as n approaches infinity. Similarly, when $p(n) \ll n^{-1}$, for example, $p(n) = 1/(n \ln n)$, then $G(n, p)$ will almost always not contain a triangle and we abuse notation

and say that $G(n, 1/(n \ln n))$ is triangle-free. It was a central observation of Erdős and Rényi that many natural graph theoretic properties become true in a very narrow range of p . They made the following key definition.

Definition 4r(n) is called a threshold function for a graph theoretic property A if

1. When $p(n) \ll r(n)$, $\lim_{n \rightarrow \infty} \Pr[G(n, p) \models A] = 0$,
2. When $p(n) \gg r(n)$, $\lim_{n \rightarrow \infty} \Pr[G(n, p) \models A] = 1$, or vice versa.

In our example, $1/n$ is a threshold function for A . Note that the threshold function, when one exists, is not unique. We could equally have said that $10/n$ is a threshold function for A .

Let's approach the problem of $G(n, c/n)$ being triangle-free once more. For every set S of three vertices let B_S be the event that S is a triangle. Then $\Pr[B_S] = p^3$. Then "triangle-freeness" is precisely the conjunction $\bigwedge \overline{B_S}$ over all S . If the B_S were mutually independent then we would have

$$\Pr \left[\bigwedge \overline{B_S} \right] = \prod \Pr [\overline{B_S}] = (1 - p^3)^{\binom{n}{3}} \sim e^{-\binom{n}{3} p^3} \rightarrow e^{-c^3/6}.$$

The reality is that the B_S are not mutually independent though when $|S \cap T| \leq 1$, B_S and B_T are mutually independent.

We apply Janson's Inequality, Theorem 8.1.1. In the notation of Section 8.1, $I = \{S \subset V(G) : |S| = 3\}$ and $S \sim T$ if and only if $|S \cap T| = 2$. Here $\epsilon = p^3 = o(1)$, $\mu = \binom{n}{3} p^3 \sim c^3/6$ and $M = e^{-\mu(1+o(1))} = e^{-c^3/6+o(1)}$. There are $6\binom{n}{4} = O(n^4)$ pairs S, T of triples with $S \sim T$. For each $\Pr[B_S \wedge B_T] = p^5$. Thus

$$\Delta = O(n^4) p^5 = n^{-1+o(1)} = o(1).$$

When $\Delta = o(1)$ Janson's Inequality sandwiches an asymptotic bound:

$$\lim_{n \rightarrow \infty} \Pr \left[\bigwedge \overline{B_S} \right] = \lim_{n \rightarrow \infty} M = e^{-c^3/6}.$$

Can we duplicate this success with the property A that G contains no (not necessarily induced) copy of a general given graph H ? We use the definitions of balanced and strictly balanced of Section 4.4.

Theorem 10.1.1 Let H be a strictly balanced graph with v vertices, e edges and a automorphisms. Let $c > 0$ be arbitrary. Let A be the property that G contains no copy of H . Then with $p = cn^{-v/e}$,

$$\lim_{n \rightarrow \infty} \Pr[G(n, p) \models A] = \exp[-c^e/a].$$

Proof. Let $A_\alpha, 1 \leq \alpha \leq \binom{n}{v} v!/a$, range over the edge sets of possible copies of H and let B_α be the event $G(n, p) \supseteq A_\alpha$. We apply Janson's Inequality. As

$$\lim_{n \rightarrow \infty} \mu = \lim_{n \rightarrow \infty} \binom{n}{v} v! p^e / a = c^e / a,$$

we find

$$\lim_{n \rightarrow \infty} M = \exp[-c^e/a]$$

Now we examine (as in Theorem 4.4.2)

$$\Delta = \sum_{\alpha \sim \beta} \Pr[B_\alpha \wedge B_\beta]$$

We split the sum according to the number of vertices in the intersection of copies α and β . Suppose they intersect in j vertices. If $j = 0$ or $j = 1$ then $A_\alpha \cap A_\beta = \emptyset$ so that $\alpha \sim \beta$ cannot occur. For $2 \leq j \leq v$ let f_j be the maximal $|A_\alpha \cap A_\beta|$, where $\alpha \sim \beta$ and α, β intersect in j vertices. As $\alpha \neq \beta$, $f_v < e$. When $2 \leq j \leq v - 1$ the critical observation is that $A_\alpha \cap A_\beta$ is a subgraph of H and hence, as H is strictly balanced,

$$\frac{f_j}{j} < \frac{e}{v}$$

There are $O(n^{2v-j})$ choices of α, β intersecting in j points, since α, β are determined, except for order, by $2v - j$ points. For each such α, β ,

$$\Pr[B_\alpha \wedge B_\beta] = p^{|A_\alpha \cup A_\beta|} = p^{2e - |A_\alpha \cap A_\beta|} \leq p^{2e - f_j}.$$

Thus

$$\Delta = \sum_{j=2}^v O(n^{2v-j}) O(n^{-(v/e)(2e-f_j)})$$

But

$$2v - j - \frac{v}{e}(2e - f_j) = \frac{vf_j}{e} - j < 0$$

so each term is $o(1)$ and hence $\Delta = o(1)$. By Janson's Inequality,

$$\lim_{n \rightarrow \infty} \Pr \left[\bigwedge \overline{B_\alpha} \right] = \lim_{n \rightarrow \infty} M = \exp[-c^e/a]$$

completing the proof.

18.1.2 10.2 Clique Number

In this section we fix $p = 1/2$ (other values yield similar results) and consider the clique number $\omega(G(n, p))$. For a fixed $c > 0$ let $n, k \rightarrow \infty$ so that

$$\binom{n}{k} 2^{-\binom{k}{2}} \rightarrow c$$

As a first approximation,

$$n \sim \frac{k}{e\sqrt{2}} \sqrt{2}^k$$

and

$$k \sim \frac{2 \ln n}{\ln 2}.$$

Here $\mu \rightarrow c$ so $M \rightarrow e^{-c}$. The Δ term was examined in Section 4.5. For this k , $\Delta = o(E[X]^2)$ and so $\Delta = o(1)$. Therefore

$$\lim_{n, k \rightarrow \infty} \Pr[\omega(G(n, p)) < k] = \exp[-c].$$

Being more careful, let $n_0(k)$ be the minimum n for which

$$\binom{n}{k} 2^{-\binom{k}{2}} \geq 1.$$

Observe that for this n the left-hand side is $1 + o(1)$. Note that $\binom{n}{k}$ grows, in n , like n^k . For any $\lambda \in (-\infty, +\infty)$ if

$$n = n_0(k) \left[1 + \frac{\lambda + o(1)}{k} \right]$$

then

$$\binom{n}{k} 2^{-\binom{k}{2}} = \left[1 + \frac{\lambda + o(1)}{k} \right]^k = e^\lambda + o(1),$$

and so

$$\Pr[\omega(G(n, p)) < k] = e^{-e^\lambda} + o(1).$$

As λ ranges from $-\infty$ to $+\infty$, e^{-e^λ} ranges from 1 to 0. As $n_0(k+1) \sim \sqrt{2}n_0(k)$ the ranges will not "overlap" for different k . More precisely, let K be arbitrarily large and set

$$I_k = \left[n_0(k) \left[1 - \frac{K}{k} \right], n_0(k) \left[1 + \frac{K}{k} \right] \right].$$

For $k \geq k_0(K)$, $I_{k-1} \cap I_k = \emptyset$. Suppose $n \geq n_0(k_0(K))$. If n lies between the intervals (which occurs for "most" n), which we denote by $I_k < n < I_{k+1}$, then

$$\Pr[\omega(G(n, p)) < k] \leq e^{-e^K} + o(1),$$

nearly zero, and

$$\Pr[\omega(G(n, p)) < k+1] \geq e^{-e^{-K}} + o(1),$$

nearly one, so that

$$\Pr[\omega(G(n, p)) = k] \geq e^{-e^{-K}} - e^{-e^K} + o(1),$$

nearly one. When $n \in I_k$ we still have $I_{k-1} < n < I_{k+1}$ so that

$$\Pr[\omega(G(n, p)) = k \text{ or } k-1] \geq e^{-e^{-k}} - e^{-e^K} + o(1),$$

nearly one. As K may be made arbitrarily large this yields the celebrated twopoint concentration theorem on clique number, Corollary 4.5.2 in Section 4.5. Note, however, that for most n the concentration of $\omega(G(n, \frac{1}{2}))$ is actually on a single value!

18.1.3 10.3 Chromatic Number

In this section we fix $p = 1/2$ (there are similar results for other p) and let G be the random graph $G(n, \frac{1}{2})$. We shall find bounds on the chromatic number $\chi(G)$. A different derivation of the main result of this section is presented in Section 7.3. Set

$$f(k) = \binom{n}{k} 2^{-\binom{k}{2}}$$

Let $k_0 = k_0(n)$ be that value for which

$$f(k_0 - 1) > 1 > f(k_0)$$

Then $n = \sqrt{2}^{k(1+o(1))}$ so for $k \sim k_0$,

$$f(k+1)/f(k) = \frac{n}{k} 2^{-k} (1 + o(1)) = n^{-1+o(1)}$$

Set

$$k = k(n) = k_0(n) - 4$$

so that

$$f(k) > n^{3+o(1)}$$

Now we estimate $\Pr[\omega(G) < k]$ by the Extended Janson Inequality (Theorem 8.1.2). Here $\mu = f(k)$. (Note that Janson's Inequality gives a lower bound of $2^{-f(k)} = 2^{-n^{3+o(1)}}$ to this probability but this is way off the mark since with probability $2^{-\binom{n}{2}}$ the random G is empty!) The value Δ was examined in Section 4.5 where

$$\frac{\Delta}{\mu^2} = \frac{\Delta^*}{\mu} = \sum_{i=2}^{k-1} g(i)$$

There $g(2) \sim k^4/n^2$ and $g(k-1) \sim 2kn2^{-k}/\mu$ were the dominating terms. In our instance $\mu > n^{3+o(1)}$ and $2^{-k} = n^{-2+o(1)}$ so $g(2)$ dominates and

$$\Delta \sim \frac{\mu^2 k^4}{n^2}$$

Hence we bound the clique number probability

$$\Pr[\omega(G) < k] < e^{-\mu^2/2\Delta} = e^{-\Theta(n^2/(\ln n)^4)},$$

as $k = \Theta(\ln n)$. [The possibility that G is empty gives a lower bound so that we may say the probability is $e^{-n^{2+o(1)}}$, though a $o(1)$ in the hyperexponent leaves lots of room.]

Theorem 10.3.1 [Bollobás (1988)] Almost always

$$\chi(G) \sim \frac{n}{2 \log_2 n}$$

Proof. Let $\alpha(G) = \omega(\bar{G})$ denote, as usual, the independence number of G . The complement of G has the same distribution $G(n, \frac{1}{2})$. Hence $\alpha(G) \leq (2 + o(1)) \log_2 n$ almost always. Thus

$$\chi(G) \geq \frac{n}{\alpha(G)} \geq \frac{n}{2 \log_2 n} (1 + o(1))$$

almost always.

The reverse inequality was an open question for a full quarter century! Set $m = \lfloor n/\ln^2 n \rfloor$. For any set S of m vertices the restriction $G|_S$ has the distribution of $G(m, \frac{1}{2})$. Let $k = k(m) = k_0(m) - 4$ as above. Note

$$k \sim 2 \log_2 m \sim 2 \log_2 n$$

Then

$$\Pr[\alpha[G|_S] < k] < e^{-m^{2+o(1)}}$$

There are $\binom{n}{m} < 2^n = 2^{m^{1+o(1)}}$ such sets S . Hence

$$\Pr[\alpha[G] < k \text{ for some } m\text{-set } S] < 2^{m^{1+o(1)}} e^{-m^{2+o(1)}} = o(1)$$

That is, almost always every m vertices contain a k -element independent set. Now suppose G has this property. We pull out k -element independent sets and give each a distinct color until there are less than m vertices left. Then we give each point a distinct color. By this procedure

$$\begin{aligned}\chi(G) &\leq \left\lceil \frac{n-m}{k} \right\rceil + m \leq \frac{n}{k} + m \\ &= \frac{n}{2 \log_2 n} (1 + o(1)) + o\left(\frac{n}{\log_2 n}\right) \\ &= \frac{n}{2 \log_2 n} (1 + o(1))\end{aligned}$$

and this occurs for almost all G .

18.1.4 10.4 Zero-one Laws

In this section we restrict our attention to graph theoretic properties expressible in the first-order theory of graphs. The language of this theory consists of variables $(x, y, z, \dots)$, which always represent vertices of a graph, equality and adjacency $(x = y, x \sim y)$, the usual Boolean connectives $(\wedge, \neg, \dots)$ and universal and existential quantification $(\forall_x, \exists_y)$. Sentences must be finite. As examples, one can express the property of containing a triangle

$$\exists_x \exists_y \exists_z [x \sim y \wedge x \sim z \wedge y \sim z]$$

having no isolated point

$$\forall_x \exists_y [x \sim y]$$

and having radius at most two

$$\exists_x \forall_y [\neg(y = x) \wedge \neg(y \sim x) \longrightarrow \exists_z [z \sim y \wedge y \sim x]].$$

For any property A and any n, p we consider the probability that the random graph $G(n, p)$ satisfies A , denoted

$$\Pr[G(n, p) \models A].$$

Our objects in this section will be the theorem of Glebskii, Kogan, Liagonkii and Talanov (1969) and independently Fagin (1976) (Theorem 10.4.1), and that of Shelah and Spencer (1988) (Theorem 10.4.2).

Theorem 10.4.1 For any fixed $p, 0 < p < 1$ and any first-order A ,

$$\lim_{n \rightarrow \infty} \Pr[G(n, p) \models A] = 0 \text{ or } 1$$

Theorem 10.4.2 For any irrational $\alpha, 0 < \alpha < 1$, setting $p = p(n) = n^{-\alpha}$ and for any first-order A ,

$$\lim_{n \rightarrow \infty} \Pr[G(n, p) \models A] = 0 \text{ or } 1.$$

Both proofs are only outlined.

We shall say that a function $p = p(n)$ satisfies the Zero-One Law if the above equality holds for every first-order A .

The Glebskii-Fagin Theorem has a natural interpretation when $p = 0.5$ as then $G(n, p)$ gives equal weight to every (labeled) graph. It then says that any first-order property A holds

for either almost all graphs or for almost no graphs. The ShelahSpencer Theorem may be interpreted in terms of threshold functions. The general results of Section 10.1 give, as one example, that $p = n^{-2/3}$ is a threshold function for containment of a K_4 . That is, when $p \ll n^{-2/3}$, $G(n, p)$ almost surely does not contain a K_4 whereas when $p \gg n^{-2/3}$ it almost surely does contain a K_4 . In between, say, at $p = n^{-2/3}$, the probability is between 0 and 1, in this case $1 - e^{-1/24}$. The (admittedly rough) notion is that at a threshold function the Zero-One Law will not hold and so to say that $p(n)$ satisfies the Zero-One Law is to say that $p(n)$ is not a threshold function - that it is a boring place in the evolution of the random graph, at least through the spectacles of the first-order language. In stark terms: What happens in the evolution of $G(n, p)$ at $p = n^{-\pi/7}$? The answer: Nothing!

Our approach to Zero-One Laws will be through a variant of the Ehrenfeucht Game, which we now define. Let G, H be two vertex disjoint graphs and t a positive integer. We define a perfect information game, denoted $\text{EHR}[G, H, t]$, with two players, denoted Spoiler and Duplicator. The game has t rounds. Each round has two parts. First the Spoiler selects either a vertex $x \in V(G)$ or a vertex $y \in V(H)$. He chooses which graph to select the vertex from. Then the Duplicator must select a vertex in the other graph. At the end of the t rounds t vertices have been selected from

each graph. Let $x_1, \dots, x_t$ be the vertices selected from $V(G)$ and $y_1, \dots, y_t$ be the vertices selected from $V(H)$, where x_i, y_i are the vertices selected in the i th round. Then Duplicator wins if and only if the induced graphs on the selected vertices are order-isomorphic; that is, if for all $1 \leq i \leq j \leq t$,

$$\{x_i, x_j\} \in E(G) \iff \{y_i, y_j\} \in E(H).$$

As there are no hidden moves and no draws one of the players must have a winning strategy and we will say that that player wins $\text{EHR}[G, H, t]$.

Lemma 10.4.3 For every first-order A there is a $t = t(A)$ so that if G, H are any graphs with $G \models A$ and $H \models \neg A$ then Spoiler wins $\text{EHR}[G, H, t]$.

A detailed proof would require a formal analysis of the first-order language so we give only an example. Let A be the property $\forall x \exists y [x \sim y]$ of not containing an isolated point and set $t = 2$. Spoiler begins by selecting an isolated point $y_1 \in V(H)$, which he can do as $H \models \neg A$. Duplicator must pick $x_1 \in V(G)$. As $G \models A$, x_1 is not isolated so Spoiler may pick $x_2 \in V(G)$ with $x_1 \sim x_2$ and now Duplicator cannot pick a "duplicating" y_2 .

Theorem 10.4.4 A function $p = p(n)$ satisfies the Zero-One Law if and only if for every t , letting $G(n, p(n)), H(m, p(m))$ be independently chosen random graphs on disjoint vertex sets,

$$\lim_{m, n \rightarrow \infty} \Pr[\text{Duplicator wins } \text{EHR}[G(n, p(n)), H(m, p(m)), t]] = 1$$

Remark. For any given choice of G, H somebody must win $\text{EHR}[G, H, t]$. (That is, there is no random play, the play is perfect.) Given this probability distribution over (G, H) there will be a probability that $\text{EHR}[G, H, t]$ will be a win for Duplicator, and this must approach 1.

Proof. We prove only the "if" part. Suppose $p = p(n)$ did not satisfy the Zero-One Law. Let A satisfy

$$\lim_{n \rightarrow \infty} \Pr[G(n, p(n)) \models A] = c,$$

with $0 < c < 1$. Let $t = t(A)$ be as given by the lemma. With limiting probability $2c(1 - c) > 0$ exactly 1 of $G(n, p(n)), H(n, p(n))$ would satisfy A and thus Spoiler would win, contradicting the assumption. This is not a full proof since when the Zero-One Law is not satisfied $\lim_{n \rightarrow \infty} \Pr[G(n, p(n)) \models A]$ might not exist. If there is a subsequence n_i on which the limit is $c \in (0, 1)$ we may use the same argument. Otherwise there will be two subsequences

n_i, m_i on which the limit is zero and one, respectively. Then letting $n, m \rightarrow \infty$ through n_i, m_i respectively, Spoiler will win $\text{EHR}[G, H, t]$ with probability approaching 1.

Theorem 10.4.4 provides a bridge from logic to random graphs. To prove that $p = p(n)$ satisfies the Zero-One Law we now no longer need to know anything about logic - we just have to find a good strategy for the Duplicator.

We say that a graph G has the full level s extension property if for every distinct $u_1, \dots, u_a, v_1, \dots, v_b \in G$ with $a + b \leq s$ there is an $x \in V(G)$ with $\{x, u_i\} \in E(G), 1 \leq i \leq a$ and $\{x, v_j\} \notin E(G), 1 \leq j \leq b$. Suppose that G, H both have the full level $s - 1$ extension property. Then Duplicator wins $\text{EHR}[G, H, s]$ by the following simple strategy. On the i th round, with $x_1, \dots, x_{i-1}, y_1, \dots, y_{i-1}$ already selected, and Spoiler picking, say, x_i , Duplicator simply picks y_i having the same adjacencies to the $y_j, j < i$ as x_i has to the $x_j, j < i$. The full extension property says that such a y_i will surely exist.

Theorem 10.4.5 For any fixed $p, 0 < p < 1$, and any $s, G(n, p)$ almost always has the full level s extension property.

Proof. For every distinct $u_1, \dots, u_a, v_1, \dots, v_b, x \in G$ with $a + b \leq s$ we define $E_{u_1, \dots, u_a, v_1, \dots, v_b, x}$ to be the event that $\{x, u_i\} \in E(G), 1 \leq i \leq a$ and $\{x, v_j\} \notin E(G), 1 \leq j \leq b$. Then

$$\Pr[E_{u_1, \dots, u_a, v_1, \dots, v_b, x}] = p^a(1 - p)^b.$$

Now define

$$E_{u_1, \dots, u_a, v_1, \dots, v_b} = \bigwedge_x \overline{E_{u_1, \dots, u_a, v_1, \dots, v_b, x}}$$

the conjunction over $x \neq u_1, \dots, u_a, v_1, \dots, v_b$. These events are mutually independent over x since they involve different edges. Thus

$$\Pr \left[\bigwedge_x \overline{E_{u_1, \dots, u_a, v_1, \dots, v_b, x}} \right] = [1 - p^a(1 - p)^b]^{n-a-b}.$$

Set $\epsilon = \min\{p, 1 - p\}^s$ so that

$$\Pr \left[\bigwedge_x \overline{E_{u_1, \dots, u_a, v_1, \dots, v_b, x}} \right] \leq (1 - \epsilon)^{n-s}.$$

The key here is that ϵ is a fixed (dependent on p, s) positive number. Set

$$E = \bigvee E_{u_1, \dots, u_a, v_1, \dots, v_b}$$

the disjunction over all distinct $u_1, \dots, u_a, v_1, \dots, v_b \in G$ with $a + b \leq s$. There are less than $s^2 n^s = O(n^s)$ such choices as we can choose a, b and then the vertices. Thus

$$\Pr[E] \leq s^2 n^s (1 - \epsilon)^{n-s}$$

But

$$\lim_{n \rightarrow \infty} s^2 n^s (1 - \epsilon)^{n-s} = 0$$

and so E holds almost never. Thus $\neg E$, which is precisely the statement that $G(n, p)$ has the full level s extension property, holds almost always.

But now we have proved Theorem 10.4.1. For any $p \in (0, 1)$ and any fixed s as $m, n \rightarrow \infty$ with probability approaching one both $G(n, p)$ and $H(m, p)$ will have the full level s extension property and so Duplicator will win $\text{EHR}[G(n, p), H(m, p), s]$.

Why can't Duplicator use this strategy when $p = n^{-\alpha}$? We illustrate the difficulty with a simple example. Let $0.5 < \alpha < 1$ and let Spoiler and Duplicator play a three move game on G, H . Spoiler thinks of a point $z \in G$ but doesn't tell Duplicator about it. Instead he picks $x_1, x_2 \in G$, both adjacent to z . Duplicator simply picks $y_1, y_2 \in H$, either adjacent or not adjacent dependent on whether $x_1 \sim x_2$. But now wily Spoiler picks $x_3 = z.H \sim H(m, m^{-\alpha})$ does not have the full level 2 extension property. In particular, most pairs y_1, y_2 do not have a common neighbor. Unless Duplicator was lucky, or shrewd, he then cannot find $y_3 \sim y_1, y_2$ and so he loses. This example does not say that Duplicator will lose with perfect play - indeed, we will show that he almost always wins with perfect play - it only indicates that the strategy used need be more complex.

We begin our proof of the Zero-One Law, Theorem 10.4.2. Let $\alpha \in (0, 1)$, α irrational, be fixed. A rooted graph is a pair (R, H) where H is a graph on vertex set, say, $V(H) = \{X_1, \dots, X_r, Y_1, \dots, Y_v\}$ and $R = \{X_1, \dots, X_r\}$ is a specified subset of $V(H)$, called the roots. For example, (R, H) might consist of one vertex Y_1 adjacent to the two roots X_1, X_2 . Let $v = v(R, H)$ denote the number of vertices that are not roots and let $e = e(R, H)$ denote the number of edges, excluding those edges between two roots. We say (R, H) is dense if $v - e\alpha < 0$ and sparse if $v - e\alpha > 0$. The irrationality of α assures us that all (R, H) are in one of these categories. We call (R, H) rigid if for all S with $R \subseteq S \subset V(H)$, (S, H) is dense. We call (R, H) safe if for all S with $R \subset S \subseteq V(H)$, $(R, H|_S)$ is sparse. Several elementary properties of these concepts are given as Exercise 4. We sometimes write (R, S) for $(R, H|_S)$ when the graph H is understood.

We think of rooted graphs as on abstract points. In a graph G we say that vertices $y_1, \dots, y_v$ form an (R, H) extension of $x_1, \dots, x_r$ if whenever X_i is adjacent to Y_j in H , x_i is adjacent to y_j in G and also whenever Y_i and Y_j are adjacent in H , y_i and y_j are adjacent in G . Note that we allow G to have more edges than H and that the edges between the roots "don't count."

Lemma 10.4.6 [Generic Extension] Let (R, H) , as given above, be safe. Let $t \geq 0$ be an arbitrary, but fixed, integer. Then in $G \sim G(n, n^{-\alpha})$ almost surely for all $x_1, \dots, x_r$ there exist $y_1, \dots, y_v$ such that

- (i) $y_1, \dots, y_v$ form an (R, H) extension of $x_1, \dots, x_r$.
- (ii) x_i, y_j are adjacent in G if and only if X_i, Y_j are adjacent in H and y_i, y_j are adjacent in G if and only if Y_i, Y_j are adjacent in H .
- (iii) (For $t > 0$) If $z_1, \dots, z_u$ with $u \leq t$ form a rigid (R', H') extension over $x_1, \dots, x_r, y_1, \dots, y_v$ then there are no adjacencies between any pair z_k, y_j .

Example. Let $\alpha \in (\frac{1}{2}, 1)$, $t = 2$, and let (R, H) have root X_1 , nonroot Y_1 and edge $\{X_1, Y_1\}$. Note that (R', H') consisting of two roots X_1, X_2 with a common neighbor Y_1 has $v = 1, e = 2$ and is rigid. Generic Extension in this instance says that every x_1 has a neighbor y_1 such that x_1, y_1 do not have a common neighbor z_1 .

Proof. From Exercise 5 almost surely every $x_1, \dots, x_r$ has $\Theta(n^v p^e)(R, H)$ extensions $y_1, \dots, y_v$. Our rough notion will be that the number of these $y_1, \dots, y_v$ that fail to be generic, in any of the bounded number of ways that could occur, would be bounded by a smaller power of n .

Call y special if $y \in \text{cl}_{t+v}(x_1, \dots, x_r)$ (as defined below), otherwise nonspecial. Let K , from the Finite Closure Lemma 10.4.7 below, be an almost sure bound on the number of special y , uniform over all choices of the x 's. Extend (R, H) to (R^+, H^+) by adding K new roots and no new edges. This is still safe and of the same type as (R, H) so again by Exercise 5 almost surely every $x_1, \dots, x_r, z_1, \dots, z_K$ has $\Theta(n^v p^e)(R^+, H^+)$ extensions $y_1, \dots, y_v$. Letting the z 's include all the special vertices we have that almost surely every $x_1, \dots, x_r$ has $\Theta(n^v p^e)(R, H)$ extensions $y_1, \dots, y_v$ with all y_i nonspecial. Now we bound from above the number of those nonspecial (R, H) extensions that fail condition (ii) or (iii).

Consider those extensions (R, H') with an additional edge y_i, y_j or x_i, y_j . This cannot contain a rigid subextension as that would make some y_i special. Hence by Exercise 4 it must

be a safe extension. Applying Exercise 5 there are $\Theta(n^v p^{e+1}) = o(n^v p^e)$ such extensions.

Consider extensions by $y_1, \dots, y_v$ and $z_1, \dots, z_u$ as in condition (iii) with some z_j, y_k adjacent. We can further assume the z 's form a minimal rigid extension over the x 's and y 's. Let the z 's have type (v_1, e_1) as an extension over the x 's and y 's so that $v_1 - e_1 \alpha$ is negative. If the y 's and z 's together formed a safe extension over the x 's there would be $\Theta(n^{v+v_1} p^{e+e_1}) = o(n^v p^e)$ such extensions and hence at most that many choices for the y 's. Otherwise, by Exercise 4, there would be a rigid subextension. It could not overlap the nonspecial y 's. From the minimality it must be precisely all of the z 's. Given the x 's from the Finite Closure Lemma 10.4.7 there are $O(1)$ choices for the z 's. Then the y 's form a (v, e') extension over the x 's and y 's with $e' > e$. This extension has no rigid subextensions (again as the y 's are nonspecial) and hence is safe. Again applying Exercise 5 there are $\Theta(n^v p^{e'})$ such y 's for each choice of the z 's and so $O(n^v p^{e'}) = o(n^v p^e)$ total choices of such y 's.

In all cases the number of y 's that fail conditions (ii) or (iii) is $o(n^v p^e)$. Hence there exist y 's, indeed most choices of nonspecial y 's, that are (R, H) extensions and satisfy conditions (ii) and (iii).

A rigid t -chain in G is a sequence $X = X_0 \subset X_1 \subset \dots \subset X_K$ with all (X_{i-1}, X_i) rigid and all $|X_{i+1} - X_i| \leq t$. The t -closure of X , denoted by $\text{cl}_t(X)$, is the maximal Y for which there exists a rigid t -chain (of arbitrary length) $X = X_0 \subset X_1 \subset \dots \subset X_K = Y$. When there are no such rigid t -chains we define $\text{cl}_t(X) = X$. To see this is well defined we note (using Exercise 4) that if $X = X_0 \subset X_1 \subset \dots \subset X_K = Z$ and $X = X_0 \subset Y_1 \subset \dots \subset Y_L = Y$ are rigid t -chains then so is $X = X_0 \subset X_1 \subset \dots \subset X_K \subset Z \cup Y_1 \subset \dots \subset Z \cup Y_L = Z \cup Y$. Alternatively, the t -closure $\text{cl}_t(X)$ is the minimal set containing X that has no rigid extensions of $\leq t$ vertices. We say $x_1, \dots, x_r \in G, y_1, \dots, y_r \in H$ have the same t -type if their t -closures are isomorphic as graphs, the isomorphism sending each x_i to the corresponding y_i .

The t -closure is a critical definition, describing the possible special properties of the roots. Suppose, for example, $\alpha \in (\frac{1}{2}, 1)$ and consider $\text{cl}_1(x_1, x_2)$. The only rigid extension with $t = 1$ in this range is a nonroot adjacent to two (or more) roots. A sample 1-type would be: x_1, x_2 have common neighbors y_1, y_2 and then x_1, y_1 have common neighbor y_3 and there are no further edges among these vertices and no pairs have common neighbors other than those described. A randomly chosen x_1, x_2 would have type: x_1, x_2 have no common neighbors and are not adjacent.

We can already describe the nature of Duplicator's strategy. At the end of the r th move, with $x_1, \dots, x_r$ and $y_1, \dots, y_r$ having been selected from the two graphs, Duplicator will assure that these sets have the same a_r -type. We shall call this the $(a_1, \dots, a_t)$ lookahead strategy. Here a_r must depend only on t , the total number of moves in the game and α . We shall set $a_t = 0$ so that at the end of the game, if Duplicator can stick to the $(a_1, \dots, a_t)$ lookahead strategy then he has won. If, however, Spoiler picks, say, x_{r+1} so that there is no corresponding y_{r+1} with $x_1, \dots, x_{r+1}$ and $y_1, \dots, y_{r+1}$ having the same a_{r+1} -type, then the strategy fails and we say that Spoiler wins. The values a_r give the "lookahead" that Duplicator uses but before defining them we need some preliminary results.

Lemma 10.4.7 [Finite Closure] Let $\alpha, r > 0$ be fixed. Set ε equal to the minimal value of $(e\alpha - v)/v$ over all integers v, e with $1 \leq v \leq t$ and $e\alpha - v > 0$. Let K be such that $r - K\varepsilon < 0$. Then in $G(n, n^{-\alpha})$ almost surely,

$$|\text{cl}_t(X)| \leq K + r$$

for all $X \subset G$ with $|X| = r$.

Proof. If not there would be a rigid t -chain $X = X_0 \subset X_1 \subset \dots \subset X_L = Y$ with $K + r < |Y| < K + r + t$. Letting (X_{i-1}, X_i) have type (v_i, e_i) the restriction of G to Y would have $r + \sum v_i$ vertices and at least $\sum e_i$ edges. But

$$\left(r + \sum v_i\right) - \alpha \left(\sum e_i\right) = r + \sum (v_i - \alpha e_i) \leq r - \varepsilon \sum v_i < r - K\varepsilon < 0$$

and G almost surely has no such subgraph.

Remark. The bound on $|\text{cl}_t(X)|$ given by this proof depends strongly on how close α may be approximated by rationals of denominator at most t . This is often the case. If, for example,

$$\frac{1}{2} + \frac{1}{s-1} > \alpha > \frac{1}{2} + \frac{1}{s}$$

then almost surely there will be two points $x_1, x_2 \in G(n, n^{-\alpha})$ having s common neighbors so that $|\text{cl}_1(x_1, x_2)| \geq s + 2$.

Now we define the $a_1, \dots, a_t$ of the lookahead strategy by reverse induction. We set $a_t = 0$. If at the end of the game Duplicator can assure that the 0-types of $x_1, \dots, x_t$ and $y_1, \dots, y_t$ are the same then they have the same induced subgraphs and he has won. Suppose, inductively, that $b = a_{r+1}$ has been defined. We define $a = a_r$ to be any integer satisfying

1. $a \geq b$.
2. Almost surely $|\text{cl}_b(W)| - r \leq a$ for all sets W of size $r + 1$.

Now we need to show that almost surely this strategy works. Let $G_1 \sim G(n, n^{-\alpha})$, $G_2 \sim G(m, m^{-\alpha})$ and Duplicator tries to play the $(a_1, \dots, a_t)$ lookahead strategy on $\text{EHR}(G_1, G_2, t)$.

Consider the $(r + 1)$ st move. We have $b = a_{r+1}, a = a_r$ as above. Points $x_1, \dots, x_r \in G_1, y_1, \dots, y_r \in G_2$ have already been selected. Set, for notational convenience, $X = \{x_1, \dots, x_r\}$ and $Y = \{y_1, \dots, y_r\}$. We assume Duplicator has survived thus far so that $\text{cl}_a(X) \cong \text{cl}_a(Y)$, the isomorphism sending each x_i to the corresponding y_i . Spoiler picks, say, $x = x_{r+1} \in G_1$. Set $X^+ = X \cup \{x\}$ and $Y^+ = Y \cup \{y\}$, where y is Duplicator's as yet undetermined countermove. We distinguish two cases.

We say Spoiler has moved inside if $x \in \text{cl}_a(X)$. Then as $b \leq a$, $\text{cl}_b(X^+) \subseteq \text{cl}_a(X)$. Duplicator looks at the isomorphism $\Psi : \text{cl}_a(X) \rightarrow \text{cl}_a(Y)$ and selects $y = \Psi(x)$.

We say Spoiler has moved outside if $x \notin \text{cl}_a(X)$. Let NEW be those vertices of $\text{cl}_b(X^+)$ that do not lie in $\text{cl}_a(X)$. $NEW \neq \emptyset$ as $x \in NEW$. $|NEW| \leq a$ as $NEW \subseteq \text{cl}_b(X^+) - X$. Consider NEW as an (R, H) extension of $\text{cl}_a(X)$. This extension must be safe as otherwise it would have a rigid subextension NEW^- but that subextension would then be in $\text{cl}_a(X)$. Duplicator now goes to G_2 and, applying the Generic Extension Lemma 10.4.6 with $t = b$, finds an (R, H) extension of $\text{cl}_a(Y)$. That is, he finds an edge preserving injection $\Psi : \text{cl}_a(X) \cup NEW \rightarrow H$ extending the isomorphism between $\text{cl}_a(X)$ and $\text{cl}_a(Y)$. Duplicator selects $y = \Psi(x)$.

Why does this work? Set $NEW' = \Psi(NEW)$ and $CORE = \Psi(\text{cl}_b(X^+))$. We can reach $\text{cl}_b(X^+)$ by a rigid b -chain from X^+ and the isomorphism gives the same chain from Y^+ to $CORE$ so that $\text{cl}_b(Y^+)$ contains $CORE$. But can it have additional vertices? We use the genericity to say no. Suppose there was a rigid extension $MORE$ over $CORE$ with at most b nonroots. We can't have $MORE$ entirely inside $\Psi[\text{cl}_a(X) \cup NEW]$ as then $\Psi^{-1}[MORE]$ would be in $\text{cl}_b(X^+)$ as well. Let $MORE^+$ be the vertices of $MORE$ lying outside $\Psi[\text{cl}_a(X) \cup NEW]$. $MORE^+$ is then a rigid extension of $\Psi[\text{cl}_a(X) \cup NEW]$. By the genericity $MORE^+$ would have no adjacencies to NEW' and so would be a rigid extension of $\Psi[\text{cl}_a(X)] = \text{cl}_a(Y)$. As $a \geq b$ the a -closure of a set cannot have rigid extensions with $\leq b$ vertices. Hence there is no $MORE$.

The first move follows the same pattern but is somewhat simpler. Set $b = a_1$ and let a satisfy $a \geq b$ and $a \geq |\text{cl}_b(x)|$ for any x . Spoiler plays $x \in G_1$. (Effectively, there is no inside move as $X = \emptyset$ is the set of previous moves and $\text{cl}_a(\emptyset) = \emptyset$.) Duplicator calculates the graph

$H = \text{cl}_b(x)$ that has, say, v vertices (including x) and e edges. Since H is a subgraph of G_1 the threshold function for the appearance of H must come before $n^{-\alpha}$. In particular, for every subgraph H' of H with v' vertices

and e' edges we cannot have $v' - \alpha e' < 0$ and therefore must have $v' - \alpha e' > 0$. The conditions of Theorem 4.4.5 then apply and G_2 almost surely has $\Theta(m^{e-v\alpha})$ copies of H . Consider any graph H^+ consisting of H together with a rigid extension of H with at most b vertices. Such H^+ would have $v + v^+$ vertices and $e + e^+$ edges with $v^+ - \alpha e^+ < 0$. The expected number of copies of H^+ is then $\Theta\left(m^{e-v\alpha+(v^+-\alpha e^+)}\right)$, which is $o(m^{e-v\alpha})$. Hence there will be in G_2 a copy of H that is not part of any such H^+ . (Effectively, this is generic extension over the empty set.) Duplicator finds the edge preserving injection $\Psi : \text{cl}_b(x) \rightarrow G_2$ giving such a copy of H and selects $y = \Psi(x)$.

We have shown that the $(a_1, \dots, a_t)$ lookahead strategy almost surely results in a win for Duplicator. By Theorem 10.4.4 this implies the Zero-One Law, Theorem 10.4.2.

18.1.5 10.5 Exercises

1. Show that there is a graph on n vertices with minimum degree at least $n/2$ in which the size of every dominating set is at least $\Omega(\log n)$.
2. Find a threshold function for the property: $G(n, p)$ contains a copy of the graph consisting of a complete graph on four vertices plus an extra vertex joined to one of its vertices.
3. Let X be the number of cycles in the random graph $G(n, p)$ with $p = c/n$. Give an exact formula for $E[X]$. Find the asymptotics of $E[X]$ when $c < 1$. Find the asymptotics of $E[X]$ when $c = 1$.
4. Here we write (R, S) for $(R, H|_S)$, where H is some fixed graph.
 - Let $R \subset S \subset T$. Show that if $(R, S), (S, T)$ are both dense then so is (R, T) . Show that if $(R, S), (S, T)$ are both sparse then so is (R, T) .
 - Let $R \subset S$. Show that if (R, S) is rigid then $(X \cup R, X \cup S)$ is rigid for any X .
 - $R \subset U$ with (R, U) not sparse. Show there is a T with $R \subset T \subset U$ with (R, T) dense. Show further there is an S with $R \subset S \subset T$ with (R, S) rigid.
 - Show that any (R, T) is either rigid or sparse itself or there exists S with $R \subset S \subset T$ such that (R, S) is rigid and (S, T) is sparse.
5. We call (R, H) hinged if it is safe but there is no S with $R \subset S \subset V(H)$ such that (S, H) is safe. For $x_1, \dots, x_r \in G$ let $N(x_1, \dots, x_r)$ denote the number of (R, H) extensions. Set $\mu = E[N] \sim n^v p^e$.
 - Let (R, H) be hinged and fix $x_1, \dots, x_r \in G$. Following the model of Section 8.5, especially Theorem 8.5.4, show that

$$\Pr[|N(x_1, \dots, x_r) - \mu| > \epsilon\mu] = o(n^{-r}).$$

- Deduce that almost surely all $N(x_1, \dots, x_r) \sim \mu$.
- Show that $N(x_1, \dots, x_r) \sim \mu$ holds for any safe (R, H) , by decomposing (R, H) into hinged extensions.

The Probabilistic Lens: Counting Subgraphs

A graph $G = (V, E)$ on n vertices has 2^n induced subgraphs but some will surely be isomorphic. How many different subgraphs can G have? Here we show that there are graphs G with $2^n(1 - o(1))$ different subgraphs. The argument we give is fairly coarse. It is typical of those situations where a probabilistic approach gives fairly quick answers to questions otherwise difficult to approach.

Let G be a random graph on n vertices with edge probability $1/2$. Let $S \subseteq V$, $|S| = t$ be fixed. For any one to one $\rho : S \rightarrow V$, $\rho \neq id$, let A_ρ be the event that ρ gives a graph isomorphism - that is, for $x, y \in S$, $\{x, y\} \in E \Leftrightarrow \{\rho x, \rho y\} \in E$. Set $M_\rho = \{x \in S : \rho x \neq x\}$. We split the set of ρ by $g = g(\rho) = |M_\rho|$.

Consider the $g(t - g) + \binom{g}{2}$ pairs x, y with $x, y \in S$ and at least one of x, y in M . For all but at most $g/2$ of these pairs $\{x, y\} \neq \{\rho x, \rho y\}$. (The exceptions are when $\rho x = y, \rho y = x$.) Let E_ρ be the set of pairs $\{x, y\}$ with $\{x, y\} \neq \{\rho x, \rho y\}$. Define a graph H_ρ with vertices E_ρ and vertex $\{x, y\}$ adjacent to $\{\rho x, \rho y\}$. In H_ρ each vertex has degree at most two ($\{x, y\}$ may also be adjacent to $\{\rho^{-1}x, \rho^{-1}y\}$) and so it decomposes into isolated vertices, paths and circuits. On each such component there is an independent set of size at least one-third the number of elements, the extreme case being a triangle. Thus there is a set $I_\rho \subseteq E_\rho$ with

$$|I_\rho| \geq |E_\rho| \geq \frac{g(t - g) + \binom{g}{2} - g/2}{3},$$

so that the pairs $\{x, y\}, \{\rho x, \rho y\}$ with $\{x, y\} \in I_\rho$ are all distinct.

For each $\{x, y\} \in I_\rho$ the event $\{x, y\} \in E \Leftrightarrow \{\rho x, \rho y\} \in E$ has probability $1/2$. Moreover these events are mutually independent over $\{x, y\} \in I_\rho$ since they involve distinct pairs. Thus we bound

$$\Pr[A_\rho] \leq 2^{-|I_\rho|} \leq 2^{-[g(t-g) + \binom{g}{2} - g/2]/3}.$$

For a given g the function ρ is determined by $\{x : \rho x \neq x\}$ and the values ρx for those x so that there are less than n^{2g} such ρ . We bound

$$\sum_{\rho \neq id} \Pr[A_\rho] = \sum_{g=1}^t \sum_{g(\rho)=g} \Pr[A_\rho] \leq \sum_{g=1}^t n^{2g} 2^{-(g(t-g) + \binom{g}{2} - g/2)/3}.$$

We make the rough bound

$$g(t - g) + \binom{g}{2} - \frac{g}{2} = g \left(t - \frac{g}{2} - 1 \right) \geq g \left(\frac{t}{2} - 1 \right)$$

since $g \leq t$. Then

$$\sum_{\rho \neq id} \Pr[A_\rho] \leq \sum_{g=1}^t \left(n^2 2^{(-t/2+1)/3} \right)^g$$

For, again being rough, $t > 50 \ln n$, $2^{1/2-t/6} < n^{-3}$ and $\sum_{\rho \neq id} \Pr[A_\rho] = o(1)$. That is, almost surely there is no isomorphic copy of $G|_S$.

For all $S \subseteq V$ with $|S| > 50 \ln n$ let I_S be the indicator random variable for there being no other subgraph isomorphic to $G|_S$. Set $X = \sum I_S$. Then $E[I_S] = 1 - o(1)$ so, by linearity of expectation - there being $2^n(1 - o(1))$ such S -

$$E[X] = 2^n(1 - o(1))$$

Hence there is a specific G with $X > 2^n(1 - o(1))$.

18.2 11 The Erdős-rényi Phase Transition

Bach, Mozart, Schubert - they will never fail you. When you perform their work properly it will have the character of the inevitable, as in great mathematics, which seems always to be made of pre-existing truths.

\author{

- E. L. Doctorow
- }

In their great work *On the Evolution of Random Graphs*, Erdős and Rényi (1960) expressed a special interest in the behavior of $\Gamma_{n,N(n)}$, the random graph with n vertices and $N(n)$ edges, when $N(n)$ was near $n/2$:

Thus the situation may be summarized as follows: the largest component of $\Gamma_{n,N(n)}$ is of order $\log n$ for $N(n)/n \rightarrow c < \frac{1}{2}$, of order $n^{2/3}$ for $N(n)/n \rightarrow c \sim \frac{1}{2}$, and of order n for $N(n)/n \rightarrow c > \frac{1}{2}$. This double "jump" of the size of the largest component when $N(n)/n$ passes the value $\frac{1}{2}$ is one of the most striking facts concerning random graphs.

Striking, indeed. The past half century has certainly confirmed the excitement that Erdős and Rényi expressed in their discovery.

18.2.1 11.1 AN Overview

We favor the more modern viewpoint, examining the random graph $G(n, p)$. The behavior of Erdős and Rényi's $\Gamma_{n,N(n)}$ then corresponds to that of $G(n, p)$ with $p = N(n)/\binom{n}{2}$. We shall assume $p = \Theta(n^{-1})$ throughout this chapter.

We shall call

$$p = \frac{c}{n}$$

the coarse parametrization. The value $\frac{1}{2}$ in the Erdős-Rényi formulation corresponds to the value $c = 1$ in our parametrization. Values $c < 1$ and $c > 1$ give $G(n, p)$ that are essentially different. We shall call

$$p = \frac{1}{n} + \lambda n^{-4/3}$$

the fine parametrization. The importance of this parametrization is not a priori at all obvious. Indeed, its "discovery" was one of the great advances in the field. In Section 11.7 we give a heuristic argument why this is the appropriate fine parametrization. Along with the fine parametrization we also define

$$\varepsilon = \lambda n^{-1/3} \text{ so that } p = \frac{1 + \varepsilon}{n}. \quad (11.1)$$

We shall express various results in terms of either λ or ε (or both), whichever best illustrates the result. We shall think of ε, λ as functions of n . To avoid negative numbers we shall sometimes parametrize $p = (1 - \varepsilon)/n$ with $\varepsilon = \lambda n^{-1/3}$. This includes functions such as $p = 1/n - 100n^{0.01}n^{-4/3}$. Of course, for n small this would give $p < 0$ and so would be nonsense. For n sufficiently large we will have $p \in [0, 1]$. As our results are always asymptotic we shall allow this slight abuse of notation and consider $G(n, p)$ defined only for n appropriately large.

In describing the nature of $G(n, p)$ we shall refer to the complexity of components, as defined below. Observe that complexity zero and one correspond to tree components and unicyclic components, respectively.

Definition 5 A connected component of a graph G with v vertices and e edges is said to have complexity $e - v + 1$. Components with complexity zero or one are called simple; components with complexity greater than one are called complex.

Let $C(v)$ denote the component containing a given vertex v . Its size $|C(v)|$ has a distribution. From the symmetry of $G(n, p)$ the distributions of all $|C(v)|$ are the same. We shall be concerned with the sizes of the largest components. We shall let C_i denote the i th largest component and L_i denote its number of vertices. Thus $L_1 = \max_v |C(v)|$. We shall be particularly interested in L_1, L_2 and whether or not they are close together.

The study of $G(n, p)$ when $p = \Theta(n^{-1})$ splits into five regions. We describe them in order of increasing p , thus giving some sense of the evolution.

Very Subcritical. Here we employ the coarse parametrization $p = c/n$ and assume c is a constant with $c < 1$.

Example: $p = 1/2n$.

- All components are simple.
- $L_1 = \Theta(\ln n)$.
- $L_k \sim L_1$ for all fixed k .

Barely Subcritical. Here we employ the fine parametrization: $p = (1 - \varepsilon)/n$ with $\varepsilon = \lambda n^{-1/3}$. We assume $\varepsilon = o(1)$ and $\lambda \rightarrow \infty$.

Example: $p = 1/n - n^{-4/3}n^{0.01}$.

- All components are simple.
- $L_1 = \Theta(\varepsilon^{-2} \ln \lambda) = \Theta(n^{2/3} \lambda^{-2} \ln \lambda)$.
- $L_k \sim L_1$ for all fixed k .

The Critical Window. Here λ is a real constant. The value $\lambda = 0$, perhaps surprisingly, has no special status.

Example: $p = 1/n \pm 2n^{-4/3}$.

- The largest k components (k fixed) all have size $L_k = \Theta(n^{2/3})$.
- Parametrizing $L_k = c_k n^{2/3}$ and letting d_k denote the complexity of C_k there is a nontrivial joint distribution for $c_1, \dots, c_k, d_1, \dots, d_k$.

Barely Supercritical. Here we employ the fine parametrization: $p = (1 + \varepsilon)/n$ with $\varepsilon = \lambda n^{-1/3}$. We assume $\varepsilon = o(1)$ and $\lambda \rightarrow \infty$.

Example: $p = 1/n + n^{-4/3}n^{0.01}$.

- $L_1 \sim 2\varepsilon n = 2\lambda n^{2/3}$.
- The largest component has complexity approaching infinity.
- All other components are simple.
- $L_2 = \Theta(\varepsilon^{-2} \ln \lambda) = \Theta(n^{2/3} \lambda^{-2} \ln \lambda)$.

Note that the ratio L_1/L_2 goes to infinity. For this reason, in this regime we call the largest component the dominant component.

Very Supercritical. We employ the coarse parametrization and assume $c > 1$.

- $L_1 \sim yn$, where $y = y(c)$ is the positive real satisfying

$$e^{-cy} = 1 - y. \quad (11.2)$$

- The largest component has complexity approaching infinity.
- All other components are simple.
- $L_2 = \Theta(\ln n)$.

Following the terminology made famous by Erdős and Rényi, we call the largest component the giant component.

We shall give arguments for only some of the above statements, and then often in limited form. Other results are given in the exercises. Full arguments for these results, and much much more, can be found in the classic texts of Bollobás (2001) and of Janson et al. (2000).

18.2.2 11.2 Three Processes

We place here in concise form three classes of probability spaces that we shall contrast and analyze. Our goal is to analyze the graph branching model. It is estimated by the binomial branching model, and thence by the Poisson branching model, which has a particularly nice analysis.

- The Poisson Branching Model.
- Parameter: Nonnegative real c .
- Underlying Space: An infinite sequence $Z_t, t = 1, 2, \dots$ of independent identically distributed random variables, each having Poisson distribution with mean c .
- Auxiliary $Y_t, t \geq 0$: Given by initial value $Y_0 = 1$ and recursion $Y_t = Y_{t-1} + Z_t - 1$.
- Auxiliary T : T is that minimal t with $Y_t = 0$. If no such t exists we write $T = \infty$.
- Nomenclature: Z_t is the number of nodes born at time t , Y_t is the queue size at time t , T is the total size.
- Interpretation: T is the total size of a Galton-Watson process, as described in Section 11.3, using a Poisson distribution with mean c .
- The Binomial Branching Model
- Parameters: Positive integer m , real $p \in [0, 1]$.
- Underlying Space: An infinite sequence $Z_t, t = 1, 2, \dots$ of independent identically distributed random variables, each having binomial distribution $B(m, p)$.
- Auxiliary $Y_t, t \geq 0$: Given by initial value $Y_0 = 1$ and recursion $Y_t = Y_{t-1} + Z_t - 1$.
- Auxiliary T : T is that minimal t with $Y_t = 0$. If no such t exists we write $T = \infty$.
- Nomenclature: Z_t is the number of nodes born at time t , Y_t is the queue size at time t , T is the total size.

- Interpretation: T is the total size of a Galton-Watson process, as described in Section 11.3 using a binomial distribution with parameters m, p .
- The Graph Branching Model
- Parameters: Positive integer n , real $p \in [0, 1]$.
- Underlying Space: A sequence $Z_1, \dots, Z_n, Z_t$ has binomial distribution with parameters N_{t-1}, p , with N_{t-1} as given below.
- Auxiliary $Y_t, t \geq 0$: Given by initial value $Y_0 = 1$ and recursion $Y_t = Y_{t-1} + Z_t - 1$.
- Auxiliary $N_t, t \geq 0$: Given by initial value $N_0 = n - 1$ and recursion $N_t = N_{t-1} - Z_t$. Equivalently, $N_t = n - t - Y_t$.
- Auxiliary T : T is that minimal t with $Y_t = 0$ or, equivalently, $N_t = n - t$. $1 \leq T \leq n$ always.
- Nomenclature: Z_t is the number of nodes born at time t , Y_t is the queue size at time t , N_t is the number of neutral vertices at time t , T is total size.
- Interpretation: T is the size of the component $C(v)$ of a given vertex v in $G(n, p)$, as found by the Breadth First Search process described in Section 11.5.

We use the superscripts *po* (Poisson), *bin* (binomial) and *gr* (graph) to distinguish these three processes when necessary.

18.2.3 11.3 The Galton-watson Branching Process

Let Z be a distribution over the nonnegative integers. The Galton-Watson process begins with a single root node, we can call her Eve. Eve has Z children. Each of her children (if there are any) now independently has Z children. The process continues, each new offspring having an independent number Z of children. Let T be the total number of nodes (including Eve herself) created in the process. It is possible that the process goes on forever, in which case we write $T = \infty$.

Our analysis of the Galton-Watson process uses fictional continuation. Let $Z_t, t = 1, 2, \dots$, be a countable sequence of independent identically distributed variables, each having distribution Z . This defines our probability space. We think of the children being born in a Breadth First Search manner; that is, Eve has her children, which are ordered in some way. Now the children, in order, have children. Each child's children are ordered in some way and this gives an ordering of Eve's grandchildren. Now the grandchildren have children in order, and the process continues. We count Eve as node number 1, her children have node numbers $2, \dots, 1 + Z_1$ and, more generally, each node is given a distinct positive integer as its node number. We let Z_t be the number of children of the t th node. Since the Z_t are independent and have distribution Z this corresponds to the Galton-Watson process. Imagine the t th node having Z_t children and then dying. By time t we mean the process after the t th node has had her children and died. Let Y_t be the number of living children at time t . We set initial value $Y_0 = 1$, corresponding to the node Eve. We have the recursion

$$Y_t = Y_{t-1} + Z_t - 1 \text{ for all } t \geq 1$$

There are two essentially different cases.

- $Y_t > 0$ for all $t \geq 0$. In this case the Galton-Watson process goes on forever and $T = \infty$.

- $Y_t = 0$ for some $t \geq 0$. In this case let T be the least integer for which $Y_T = 0$. Then the Galton-Watson process stops with the death of the T th node and T is the total number of nodes in the process.

Our fictional continuation enables us to consider the Y_t as an infinite random walk, with step size $Z - 1$. When $c < 1$ the walk has negative drift and so tends to $-\infty$. When $c > 1$ the walk has positive drift and tends to $+\infty$. The process when $c < 1$ is called subcritical and the process when $c > 1$ is called supercritical. When $c = 1$ the walk has zero drift and the situation is especially delicate.

The above is quite general. When Z is Poisson or binomial (the only cases of interest to us) this yields the Poisson branching process and the binomial branching process of Section 11.2.

18.2.4 11.4 Analysis of the Poisson Branching Process

In this section we study $T = T_c^{po}$. We often drop the value c and the superscript po for notational simplicity.

Theorem 11.4.1 If $c < 1$, T is finite with probability one. If $c = 1$, T is finite with probability one. If $c > 1$ then T is infinite with probability $y = y(c)$, where y is that unique positive real satisfying the equation (11.2).

Proof. Suppose $c < 1$. If $T > t$ then $Y_t > 0$ so that $Z_1 + \cdots + Z_t \geq t$. Chernoff bounds give that $\Pr[Y_t > 0] < e^{-kt}$ for a constant k . In particular, $\Pr[Y_t > 0] \rightarrow 0$ so that $\Pr[T > t] \rightarrow 0$ and T is finite with probability one.

Suppose $c \geq 1$. Set $z = 1 - y = \Pr[T < \infty]$. Given that Eve has i children the probability that the branching process is finite is z^i as all i branches must be finite. Thus

$$z = \sum_{i=0}^{\infty} \Pr[Z_1 = i] z^i = \sum_{i=0}^{\infty} e^{-c} \frac{c^i z^i}{i!} = e^{c(z-1)}.$$

Setting $y = 1 - z$ gives the equation (11.2). For $c = 1$, $e^{-y} > 1 - y$ for $y > 0$ so the solution must be $y = 0$. For $c > 1$ the function $f(y) = 1 - y - e^{-cy}$ has $f(0) = 1$, $f(1) < 0$ and $f'(0) = c - 1 > 0$ so there is a $y \in (0, 1)$ with $f(y) = 0$. Furthermore, as f is convex, there is precisely one y . We have shown that either $\Pr[T < \infty] = 1$ or $\Pr[T < \infty] = 1 - y > 0$. The argument that $\Pr[T < \infty] \neq 1$ (not surprising as the walk has positive drift) is left for the exercises.

Theorem 11.4.2 For any positive real c and any integer k , setting $T = T_c^{po}$,

$$\Pr[T = k] = \frac{e^{-ck} (ck)^{k-1}}{k!}.$$

We defer the proof of this classic result to Section 11.6 when we will give a probabilistic proof!

When $c = 1$ Stirling's formula gives

$$\Pr[T_1 = k] = \frac{e^{-k} k^{k-1}}{k!} \sim \frac{1}{\sqrt{2\pi}} k^{-3/2}. \quad (11.3)$$

This perforce approaches zero but it does so only at polynomial speed. In general,

$$\Pr[T_c = k] \sim \frac{1}{\sqrt{2\pi}} k^{-3/2} c^{-1} (ce^{1-c})^k.$$

For any $c \neq 1$ (whether larger or smaller than one) $ce^{1-c} < 1$ and therefore $\Pr[T_c = k]$ approaches zero at exponential speed. This gives a bound on the tail distribution

$$\Pr[T_c \geq u] < e^{-u(\alpha+o(1))}, \quad (11.4)$$

where $\alpha = c - 1 - \ln c > 0$.

We are particularly interested in the Poisson branching process when c is near one. Let us parametrize

$$c = 1 + \varepsilon.$$

When $\varepsilon > 0$, $\Pr[T_{1+\varepsilon} = \infty]$ is the $y = y(\varepsilon) \in (0, 1)$ satisfying $f(y) = 1 - y - e^{-(1+\varepsilon)y} = 0$. Some fun calculus gives

$$\Pr[T_{1+\varepsilon} = \infty] \sim 2\varepsilon \text{ as } \varepsilon \rightarrow 0^+. \quad (11.5)$$

Suppose $c \rightarrow 1^+$ so that $\varepsilon \rightarrow 0^+$. We have

$$\ln(ce^{1-c}) = \ln(1 + \varepsilon) - \varepsilon \sim -\frac{\varepsilon^2}{2}.$$

Thus

$$\Pr[T_{1+\varepsilon} = u] \sim \frac{1}{\sqrt{2\pi}} u^{-3/2} \text{ for } u = o(\varepsilon^{-2}).$$

Note that $\Pr[T_{1+\varepsilon} = u] \sim \Pr[T_1 = u]$ in this range. When u reaches order ε^{-2} there is a change. For $u = A\varepsilon^{-2}$ and fixed A ,

$$\Pr[T_{1+\varepsilon} = A\varepsilon^{-2}] \sim \frac{1}{\sqrt{2\pi}} \varepsilon^3 A^{-3/2} e^{-A/2}.$$

When $A \rightarrow \infty$ we absorb smaller factors into the exponential term:

$$\Pr[T_{1+\varepsilon} = A\varepsilon^{-2}] = \varepsilon^3 e^{-(1+o(1))A/2}.$$

When c is slightly less than one we can write $c = 1 - \varepsilon$, where $\varepsilon \rightarrow 0^+$. We have $\ln(ce^{-c}) \sim -\frac{1}{2}\varepsilon^2$, the same as for $c = 1 + \varepsilon$. Indeed when $u = o(\varepsilon^{-3})$,

$$\Pr[T_{1-\varepsilon} = u] \sim \Pr[T_{1+\varepsilon} = u].$$

For $A \rightarrow \infty$,

$$\Pr[T_{1-\varepsilon} = A\varepsilon^{-2}] = \varepsilon^3 e^{-(1+o(1))A/2}.$$

The Poisson branching processes with means $1 + \varepsilon$ and $1 - \varepsilon$ look almost the same, with the (important!) distinction that the mean $1 + \varepsilon$ process is sometimes infinite while the mean $1 - \varepsilon$ process never is.

In short, the Poisson branching process with mean $1 \pm \varepsilon$ acts as if it had mean 1 until reaching size on the order ε^{-2} . Until then $\Pr[T_{1\pm\varepsilon} = u]$ is dropping at a polynomial rate. Upon reaching order ε^{-2} , $\Pr[T_{1\pm\varepsilon} = u]$ drops exponentially in u .

We are particularly interested in the tail distribution. For $\varepsilon \rightarrow 0^+$ and $A \rightarrow \infty$,

$$\Pr[T_{1-\varepsilon} > A\varepsilon^{-2}] < e^{-(1+o(1))A/2} \varepsilon. \quad (11.6)$$

The same holds for the finite part of $T_{1+\varepsilon}$:

$$\Pr[\infty > T_{1+\varepsilon} > A\varepsilon^{-2}] < e^{-(1+o(1))A/2} \varepsilon$$

When $A \rightarrow \infty$ this quantity is $o(\varepsilon)$ so (11.5) gives

$$\Pr[T_{1+\varepsilon} > A\varepsilon^{-2}] \sim 2\varepsilon \text{ when } \varepsilon \rightarrow 0^+ \text{ and } A \rightarrow \infty. \quad (11.7)$$

18.2.5 11.5 The Graph Branching Model

Abbreviation: We use BFS as an abbreviation for Breadth First Search. BFS algorithms are a mainstay of computer science and central to our approach.

Let $C(v)$ denote the component, in $G(n, p)$, containing a designated vertex v . We generate $C(v)$ using the (standard) BFS algorithm to find $C(v)$. We begin with root v . In this procedure all vertices will be live, dead, or neutral. The live vertices will be contained in a queue. Initially, at time zero, v is live, the queue consists of one vertex, v itself, and all other vertices are neutral. At each time t we remove a live vertex w from the top of the queue (in computer science parlance we "pop the queue") and check all pairs $\{w, w'\}$, w' neutral, for adjacency in G . The popped vertex w is now dead. Those neutral w' (if any) adjacent to w are added to the bottom of the queue and are now live. (They can be placed in any particular order.) The procedure ends when the queue is empty. We let T denote that time. At time T all vertices are neutral or dead and the set of dead vertices is precisely the component $C(v)$. That is, $T = |C(v)|$.

Let Z_t denote the number of vertices added to the queue at time t . Let Y_t denote the size of the queue at the conclusion of time t . We set $Y_0 = 1$, reflecting the initial size of the queue. At time t we remove one vertex and add Z_t vertices to the queue so we have the recursion $Y_t = Y_{t-1} - 1 + Z_t$. Let N_t denote the number of neutral vertices at time t . As Z_t vertices switched from neutral to live at time t , N_t satisfies the recursion $N_0 = n - 1$, $N_t = N_{t-1} - Z_t$. Equivalently, as there are t dead and Y_t live vertices at time t , $N_t = n - t - Y_t$. Z_t is found by checking N_{t-1} pairs for adjacency. As these pairs have not yet been examined they remain adjacent with independent probability p . That is,

$$Z_t \sim B(N_{t-1}, p) \sim B(n - (t - 1) - Y_{t-1}, p) \quad (11.8)$$

The graph branching process of Section 11.2 mirrors the above analysis until time T and then continues until time n . This fictional continuation shall be useful in the analysis of $C(v)$. The graph branching process is similar to a binomial branching process in that the Z_t have binomial distributions but dissimilar in that the parameter N_{t-1} in the graph branching process depends on previous values Z_i .

As $N_t = N_{t-1} - Z_t$, (11.8) yields $N_t \sim B(N_{t-1}, 1 - p)$. By induction we find the distributions

$$N_t \sim B(n - 1, (1 - p)^t) \quad \text{for } 0 \leq t \leq n.$$

If $T = t$ it is necessary (though not sufficient, due to fictitious continuation) that $N_t = n - t$. This yields the useful inequalities.

Theorem 11.5.1 In $G(n, p)$

$$\Pr[|C(v)| = t] \leq \Pr[B(n - 1, (1 - p)^t) = n - t]$$

or, equivalently,

$$\Pr[|C(v)| = t] \leq \Pr[B(n - 1, 1 - (1 - p)^t) = t - 1]. \quad (11.9)$$

An Alternate Analysis. The following analysis of $C(v)$ on $G(n, p)$ has been explored by van der Hofstad and Spencer (2006). Each $w \neq v$ flips a coin, heads with probability p , repeatedly until getting a head. Let X_w denote that flip on which w gets a head. Suppose $X_w = j$. Then w enters the BFS at time j . (However, it may have missed the boat if the BFS has already terminated.) This reverses the usual randomness; we are here imagining the $w \neq v$ trying to get into the BFS tree, rather than the BFS tree trying to expand by finding neutral vertices. Suppose $t = |C(v)|$. Every $w \neq v$ that is in $C(v)$ must have entered by time t so $X_w \leq t$. Every $w \neq v$ that is not in $C(v)$ had t opportunities to enter $C(v)$ and so $X_w > t$. Thus $\Pr[|C(v)| = t]$ is at most the probability that $X_w \leq t$ for precisely $t - 1$ $w \neq v$. For each

$w \neq v$, $\Pr[X_w = t] = 1 - (1 - p)^t$ and these events are independent over w , yielding (11.9). In van der Hofstad and Spencer (2006) this analysis is extended to give more precise bounds on $\Pr[|C(v)| = t]$.

18.2.6 11.6 The Graph and Poisson Processes Compared

Set $p = c/n$. A key observation is that $Z_1 \sim B(n - 1, c/n)$ approaches (in n) the Poisson distribution with mean c . Furthermore, in a more rough sense, the same holds for Z_t as long as $N_{t-1} \sim o(n)$ or, equivalently, the number of live and dead vertices is $o(n)$. That is, the generation of $C(v)$ mimics the Poisson branching process with mean c as long as the number of vertices found is not too large. This allows for a very accurate description in the very subcritical regime $c < 1$. But in the very supercritical regime $c > 1$ the relationship between the generation of $C(v)$ and the Poisson branching process breaks down. As the number N_{t-1} of neutral vertices drops so does the expected number $E[Z_t]$ of vertices added to the queue. Eventually the drift of the walk Y_t lowers from positive to negative, and this eventually causes the process to halt. We call this phenomenon the ecological limitation. Indeed, there must be an ecological limitation. The Poisson branching process becomes infinite with positive probability; the component $C(v)$ tautologically cannot be greater than n .

Theorem 11.6.1 For any positive real c and any fixed integer k

$$\lim_{n \rightarrow \infty} \Pr[|C(v)| = k \text{ in } G(n, c/n)] = \Pr[T_c = k]$$

Proof. Let Z_t^{po}, T^{po} and Z_t^{gr}, T^{gr} denote the values in the Poisson branching process with parameter c and the graph branching process with parameters n, p , respectively. Let Γ denote the set of k -tuples $\vec{z} = (z_1, \dots, z_k)$ of nonnegative integers such that the recursion $y_0 = 1, y_t = y_{t-1} + z_t - 1$ has $y_t > 0$ for $t < k$ and $y_k = 0$. Then

$$\begin{aligned} \Pr[T^{gr} = k] &= \sum \Pr[Z_i^{gr} = z_i, 1 \leq i \leq k], \\ \Pr[T^{po} = k] &= \sum \Pr[Z_i^{po} = z_i, 1 \leq i \leq k], \end{aligned}$$

where both sums are over $\vec{z} \in \Gamma$. Fix such a $\vec{z}$.

$$\Pr[Z_i^{gr} = z_i, 1 \leq i \leq k] = \prod_{i=1}^k \Pr[B(N_{i-1}^{gr}, p) = z_i].$$

As i, y_{i-1}, z_i are fixed $Z_{i-1} = n - O(1)$ and $B(Z_{i-1}, p)$ approaches the Poisson distribution. More precisely,

$$\lim_{n \rightarrow \infty} \Pr[B(Z_{i-1}, p) = z_i] = \Pr[Z_i^{po} = z_i]$$

Furthermore, as the products are of a fixed number of terms,

$$\lim_{n \rightarrow \infty} \Pr[Z_i^{gr} = z_i, 1 \leq i \leq k] = \Pr[Z_i^{po} = z_i, 1 \leq i \leq k]$$

Proof [Theorem 11.4.2]. By Theorem 11.6.1,

$$\Pr[T_c^{po} = k] = \lim_{n \rightarrow \infty} \Pr[|C(v)| = k]$$

where the second probability is in $G(n, p)$ with $p = c/n$ and v is an arbitrary vertex of that graph. There are $\binom{n}{k-1}$ choices for $C(v)$. On any particular such S there is probability $O(p^k) = O(n^{-k})$ that $G(n, p)$ has more than $k - 1$ edges. If $G(n, p)$ has precisely $k - 1$ edges on S they must form a tree. There are k^{k-2} such trees. Each occurs with probability

$p^{k-1}(1-p)^{\binom{k}{2}-k+1} \sim p^{k-1} = c^{k-1}n^{1-k}$. Thus the total probability that $G(n, p)$ restricted to S forms a connected graph is $\sim k^{k-2}c^{k-1}n^{1-k}$. As a connected component, we must further have no edges between $C(v)$ and its complement; this has probability $(1-p)^{k(n-k)} \sim e^{-ck}$. Thus

$$\Pr[C(v) = k] \sim \binom{n}{k-1} k^{k-2} c^{k-1} n^{1-k} e^{-ck} \rightarrow \frac{e^{-ck} (ck)^{k-1}}{k!},$$

as desired.

The graph branching process can be compared to the binomial branching process in both directions. An important cautionary note: the event $T_{n-1,p}^{\text{bin}} \geq u$ in Theorem 11.6.2 (and similarly $T_{n-u,p}^{\text{bin}} \geq u$ in Theorem 11.6.3) includes the possibility that the binomial branching process is infinite. Indeed, in application this will be the critical term.

Theorem 11.6.2 For any u

$$\Pr[T_{n,p}^{\text{gr}} \geq u] \leq \Pr[T_{n-1,p}^{\text{bin}} \geq u]$$

Proof. We modify the graph branching process by constantly replenishing the supply of neutral vertices. That is, when we pop the vertex w and there are $n-1-s$ neutral vertices, we create s fictional vertices w' and allow w, w' to be adjacent with probability p . This gives a component of size $T_{n-1,p}^{\text{bin}}$; the actual $C(v)$ will be a subset of it. Thus $T_{n-1,p}^{\text{bin}}$ dominates $T_{n,p}^{\text{gr}}$.

Theorem 11.6.3 For any u

$$\Pr[T_{n,p}^{\text{gr}} \geq u] \geq \Pr[T_{n-u,p}^{\text{bin}} \geq u]$$

Proof. We halt the graph branching process when the number of found (live plus dead) vertices reaches u . This does not affect the probability of finding at least u vertices. In this truncated graph process we diminish the number of neutral vertices to $n-u$. That is, when we pop the vertex w and there are $n-1-s \geq n-u$ neutral vertices, we select $n-u$ of them and only allow adjacencies w, w' to them. The truncated graph process dominates this truncated binomial $n-u, p$ process and so has a greater or equal probability of reaching u .

The Poisson Approximation. We are working in the range $p = \Theta(n^{-1})$. There the binomial $B(n-1, p)$ distribution and the Poisson distribution with mean np are very close. The Poisson branching process is precisely understood and, we feel, the "purest" branching process. Our goal in this chapter is to give the reader a picture for the "why" of the various regimes. To do this we shall often avoid the technical calculations and simply assume that the binomial $n-1, p$ branching process is very close to the Poisson branching process with mean np .

18.2.7 11.7 The Parametrization Explained

In the parametrization (11.1) for the critical window why is the exponent $-\frac{1}{3}$ as opposed to, say, $-\frac{1}{4}$ or $-\frac{2}{7}$ or something completely different? In the experience of the authors this is the question most frequently asked about the Erdős-Rényi phase transition. Here is a heuristic that may be helpful.

Parametrize $p = (1 + \varepsilon)/n$ with $\varepsilon = \varepsilon(n)$ positive and approaching zero. We look for the following picture. Consider the Poisson branching process $T = T_{1+\varepsilon}^{\text{po}}$. It is infinite with probability $\sim 2\varepsilon$, otherwise its probability of exceeding $A\varepsilon^{-2}$ drops exponentially in A . The graph branching process mimics the Poisson branching process as long as it is not too successful. The cases when the Poisson branching process is finite are mimicked, yielding components of size up to roughly ε^{-2} . The cases when the Poisson branching process is infinite are mimicked by components that "escape" until the ecological limitation sets in. These components all join together. They form a single component, the dominant component, of size $2\varepsilon n$.

In order for the above (admittedly rough) picture to hold there needs to be a distinction between the small components, up to size ε^{-2} , and the dominant component of size $2\varepsilon n$. That is, we need $2\varepsilon n \gg \varepsilon^{-2}$. This heuristic leads us to $\varepsilon = n^{-1/3}$ as the breakpoint. When $\varepsilon \gg n^{-1/3}$ we have the distinction between small and dominant and are in the supercritical regime. When $\varepsilon = O(n^{-1/3})$ there is no effective analogy to the Poisson branching process being infinite, and there is no dominant component.

18.2.8 11.8 The Subcritical Regimes

Let $p = c/n$ with $c < 1$. Theorem 11.6.2 gives

$$\Pr[T_{n,p}^{gr} \geq u] \leq \Pr[T_{n-1,p}^{bin} \geq u]$$

With the Poisson approximation,

$$\Pr[|C(v)| \geq u] \leq (1 + o(1)) \Pr[T_c \geq u].$$

From (11.4) this drops exponentially in u . Taking $u = K \ln n$ for appropriately large K , $\Pr[|C(v)| \geq u] < n^{-1.01}$. As this holds for each of the n vertices v , the probability that any v has $|C(v)| \geq u$ is less than $nn^{-1.01} \rightarrow 0$. That is, $L_1 = O(\ln n)$ with probability tending to one.

Let's push this argument into the barely subcritical regime $p = (1 - \varepsilon)/n$ with $\varepsilon = \lambda n^{-1/3}$. Let I_v be the indicator random variable for $C(v)$ having at least u vertices, u to be determined below. As above Theorem 11.6.2 and our Poisson approximation give the bound

$$\Pr[|C(v)| \geq u] \leq (1 + o(1)) \Pr[T_{1-\varepsilon} \geq u]$$

We now parametrize

$$u = K\varepsilon^{-2} \ln \lambda = Kn^{2/3} \lambda^{-2} \ln \lambda$$

For an appropriately large constant K the bound (11.6) gives

$$\Pr[T_{1-\varepsilon} \geq u] \leq \varepsilon e^{-3.1\lambda} = \varepsilon \lambda^{-3.1}$$

Let $X = \sum_v I_v$ be the number of vertices v in components of size at least u and let Y be the number of components of $G(n, p)$ of size at least u . Linearity of expectation gives

$$\mathbb{E}[X] = n\mathbb{E}[I_v] \leq n\varepsilon \lambda^{-3.1} = n^{2/3} \lambda^{-2.1}$$

As $Y \leq Xu^{-1}$,

$$\mathbb{E}[Y] \leq u^{-1}\mathbb{E}[X] \leq K^{-1} \lambda^{-0.1} \rightarrow 0$$

With probability approaching one, $Y = 0$ and so

$$L_1 \leq u = K\varepsilon^{-2} \ln \lambda = Kn^{2/3} \lambda^{-2} \ln \lambda.$$

18.2.9 11.9 The Supercritical Regimes

In the supercritical regimes there are two salient points about the giant or dominant component; First, it exists. Second, it is unique. Neither is trivial.

The Very Supercritical Phase. We start with the very supercritical region, $p = c/n$, with $c > 1$ constant. The ideas here will carry into the barely supercritical region. Let $y = y(c)$ be

the positive real solution of the equation $e^{-cy} = 1 - y$. Let δ be an arbitrarily small constant and let K be an appropriately large constant. Set $S = K \ln n$, $L^- = (y - \delta)n$ and $L^+ = (y + \delta)n$. Call a component $C(v)$ and its size $|C(v)|$ small if $|C(v)| < S$, giant if $L^- < |C(v)| < L^+$ and awkward otherwise.

No Middle Ground. We claim that the probability of having any awkward component is $o(n^{-20})$. (We could make '20' arbitrarily large by changing K .) There are n choices for v and n choices for $t = |C(v)|$. Thus it suffices to show that for any v and for any awkward t that $\Pr[|C(v)| = t] = o(n^{-18})$. From Theorem 11.5.1 it suffices to bound $\Pr[B(n-1, 1 - (1 - c/n)^t) = t - 1]$. We indicate the technical calculations. When $t = o(n)$, $1 - (1 - c/n)^t \sim cn/t$ and $c > 1$ so $\Pr[B(n-1, 1 - (1 - c/n)^t) \leq t - 1]$ is exponentially small in t . As $t \geq K \ln n$ this is polynomially small in n . When $t \sim xn$, $1 - (1 - c/n)^t \sim 1 - e^{-cx}$. For $x \neq y$, $1 - e^{-cx} \neq x$ so the mean of the binomial is not near t and the probability that it is equal to t is exponentially small in n . In all cases the bounds on $\Pr[|C(v)| = t]$ follow from basic Chernoff bounds.

Escape Probability. Set $\alpha = \Pr[C(v) \text{ is not small}]$. [When this happens we like to think that the BFS on $G(n, p)$ starting with root v has escaped an early death.] Theorems 11.6.2 and 11.6.3 sandwich

$$\Pr[T_{n-S,p}^{\text{bin}} \geq S] \leq \alpha \leq \Pr[T_{n-1,p}^{\text{bin}} \geq S]$$

From our Poisson approximation both $\Pr[T_{n-S,p}^{\text{bin}} \geq S]$ and $\Pr[T_{n-1,p}^{\text{bin}} \geq S]$ are asymptotic to $\Pr[T_c \geq S]$. Thus $\alpha \sim \Pr[T_c \geq S]$. As c is assumed fixed and $S \rightarrow \infty$,

$$\alpha \sim \Pr[T_c \geq S] \sim \Pr[T_c = \infty] = y,$$

with y as in (11.2).

Because there is no middle ground, not small is the same as giant. $C(v)$ is giant with probability $\sim y$. Thus the expected number of vertices in giant components is $\sim yn$. Each giant component has size between $(y - \delta)n$ and $(y + \delta)n$. Our goal is a single giant component of size $\sim yn$. We are almost there - but maybe with probability $\frac{1}{2}$ there are two giant components.

Sprinkling. Set $p_1 = n^{-3/2}$. (Any p_1 with $n^{-2} \ll p_1 \ll n^{-1}$ would do here.) Let $G_1 \sim G(n, p_1)$ be selected independently from $G \sim G(n, p)$ on the same vertex set and let $G^+ = G \cup G_1$ so that $G^+ \sim G(n, p^+)$ with $p^+ = p + p_1 - pp_1$. (We "sprinkle" the relatively few edges of G_1 on G to make G^+ .) Suppose $G(n, p)$ had more than one giant component and let V_1, V_2 be the vertex sets of two of those components. There are $\Omega(n^2)$ pairs $\{v_1, v_2\}$ with $v_1 \in V_1, v_2 \in V_2$. We have selected p_1 large enough so that with probability $1 - o(1)$ at least one of these pairs is in the sprinkling G_1 . Adding this edge merges components V_1, V_2 into a component of size at least $2y(1 - \delta)n$ in G^+ . We have selected p_1 small enough so that $p^+ \sim p = c/n$. The probability that G^+ has a component so large, and hence awkward, is therefore $o(n^{-20})$. Hence the probability that G had more than one giant component is $o(n^{-20})$.

Finally, we make δ arbitrarily small. $G(n, p)$ has an expected number $\sim yn$ of points in giant components and giant components all have size $\sim yn$. Furthermore, by the sprinkling argument, the contribution to this expectation from the possibility of G having more than one giant component is $o(nn^{-20})$, which is negligible. Thus with probability $1 - o(1)$ there is precisely one giant component. This gives the salient features of the very supercritical phase. There is a giant component so $L_1 \sim yn$. There is only one giant component and no middle ground so $L_2 \leq S = O(\ln n)$.

The sprinkling for complexity argument given below in the barely supercritical phase can easily be modified to show that the giant component has high complexity, indeed, complexity $\Omega(n)$.

The Barely Supercritical Phase. Set $p = (1 + \varepsilon)/n$ with $\varepsilon = \lambda n^{-1/3}$ and $\lambda \rightarrow \infty$. Note $\varepsilon^{-2} = \lambda^{-2} n^{2/3} \ll 2\varepsilon n$. The analysis of the barely supercritical region becomes more difficult

as $\lambda = \lambda(n)$ approaches infinity more slowly. We shall add the simplifying assumption that $\lambda \gg \ln n$. Furthermore, we shall find somewhat weaker bounds than stated on L_2 .

Bollobás (1984) showed the existence of the dominant component when $\lambda > K\sqrt{\ln n}$, K constant. That paper was the first indication of the appropriate scaling for the critical window. Łuczak (1990a) tightened the result to "best possible," showing that if $\lambda \rightarrow +\infty$ then the dominant component exists.

Let δ be an arbitrarily small constant and let K be an appropriately large constant. Set $S = K\varepsilon^{-2} \ln n$, $L^- = (1 - \delta)2\varepsilon n$ and $L^+ = (1 + \delta)2\varepsilon n$. Call a component $C(v)$ and its size $|C(v)|$ small if $|C(v)| < S$, dominant if $L^- < |C(v)| < L^+$ and awkward otherwise.

No Middle Ground. We claim that the probability of having any awkward component is $o(n^{-20})$. (We could make '20' arbitrarily large by changing K .) There are n choices for v and n choices for $t = |C(v)|$. Thus it suffices to show that for any v and for any awkward t that $\Pr[|C(v)| = t] = o(n^{-18})$. Again we bound $\Pr[B(n-1, 1 - (1-p)^t) = t-1]$. We indicate the technical calculations. Let μ and σ^2 denote the mean and the variance of the binomial. Then $\mu = (n-1)(1 - (1-p)^t)$ and, in this range, $\sigma^2 \sim \mu$. When $t = o(n\varepsilon)$ we estimate $1 - (1-p)^t$ by $pt = t + t\varepsilon$. Then $\mu - t \sim -t\varepsilon$ and $\sigma^2 \sim t$. This probability is roughly $\exp[-(t\varepsilon)^2/2t] = \exp[-t\varepsilon^2/2]$. As $t \geq S$ the exponent is $o(n^{-18})$ for $K > 36$. [To push S down to $K\varepsilon^{-2} \ln \lambda$ requires a finer bound on $\Pr[|C(v)| = t]$.] Now suppose $t \sim xn\varepsilon$, where $x \neq 2$. The ecological limitation now has an effect and we estimate $1 - (1-p)^t$ by $pt - \frac{1}{2}p^2t^2$ so

$$\mu - t \sim t\varepsilon - \frac{1}{2}t^2n^{-2} \sim (n\varepsilon) \left(x - \frac{1}{2}x^2 \right)$$

[Observe that when $x = 2$ the mean of the binomial is very close to t and so we do not get a small bound on $\Pr[|C(v)| = t]$. This is natural when we consider that there will be a dominant component of size $\sim 2\varepsilon n$.] Again $\sigma^2 \sim t$ so the probability is $\exp[-\Omega((n\varepsilon)^2/t)]$, which is extremely small. When $t \gg n\varepsilon$ the probability is even smaller.

Escape Probability. Set $\alpha = \Pr[C(v) \text{ is not small}]$. Theorems 11.6.2 and 11.6.3 sandwich

$$\Pr[T_{n-S,p}^{\text{bin}} \geq S] \leq \alpha \leq \Pr[T_{n-1,p}^{\text{bin}} \geq S]$$

The Poisson approximation for $T_{n-1,p}^{\text{bin}}$ is $T_{1+\varepsilon}$. As $S \gg \varepsilon^{-2}$, bound (11.7) gives

$$\alpha \leq \Pr[T_{1+\varepsilon} \geq S] \sim \Pr[T_{1+\varepsilon} = \infty] \sim 2\varepsilon.$$

Replacing $n-1$ by $n-S$ lowers the mean by $\sim Sn^{-1}$. But $Sn^{-1}/\varepsilon \sim (\ln n)/(n\varepsilon^3) = \lambda^{-3} \ln n$ and we have made λ large enough that this is $o(1)$. That is, $Sn^{-1} = o(\varepsilon)$. Therefore $T_{n-S,p}^{\text{bin}}$ is approximated by $T_{1+\varepsilon-o(\varepsilon)}$ and

$$\alpha \geq \Pr[T_{1+\varepsilon+o(\varepsilon)} \geq S] \sim \Pr[T_{1+\varepsilon+o(\varepsilon)} = \infty] \sim 2\varepsilon$$

α has been sandwiched and $\alpha \sim 2\varepsilon$.

Because there is no middle ground, not small is the same as dominant. $C(v)$ is dominant with probability $\sim 2\varepsilon$. Thus the expected number of vertices in dominant components is $\sim 2n\varepsilon$. Each dominant component has size between $(1 - \delta)2n\varepsilon$ and $(1 + \delta)2n\varepsilon$. As in the very supercritical case, we need worry about having more than one dominant component.

Sprinkling. Set $p_1 = n^{-4/3}$. Let $G_1 \sim G(n, p_1)$ be selected independently from $G \sim G(n, p)$ on the same vertex set and let $G^+ = G \cup G_1$ so that $G^+ \sim G(n, p^+)$ with $p^+ = p + p_1 - pp_1 = 1 + \varepsilon + o(\varepsilon)$. Suppose $G(n, p)$ had more than one giant component and let V_1, V_2 be the vertex sets of two of those components. There are $\gg n^{4/3}$ pairs $\{v_1, v_2\}$ with $v_1 \in V_1, v_2 \in V_2$. With probability $1 + o(1)$ at least one of these pairs is in the sprinkling G_1 . Adding this edge merges components V_1, V_2 into a component of size at least $(1 - \delta)4\varepsilon n$ in G^+ . The probability G^+ has such a large, and hence awkward, component is $o(n^{-20})$. Thus the probability G had had two

(or more) dominant components is $o(n^{-20})$. Taking δ arbitrarily small, as in the supercritical case, G has with probability $1 - o(1)$ precisely one dominant component. Thus $L_1 \sim 2n\varepsilon$ and, as there is no middle ground, $L_2 \leq K\varepsilon^{-2} \ln n$.

Sprinkling for Complexity. Take $p_1 = (1 + \varepsilon/2)/n$ and $p_2 \sim \varepsilon/2n$ so that $p_1 + p_2 - p_1p_2 = (1 + \varepsilon)/n$. Let $G_1 \sim G(n, p_1)$, $G_2 \sim G(n, p_2)$ and $G_3 = G_1 \cup G_2$ so that $G_3 \sim G(n, (1 + \varepsilon)/n)$. G_1, G_3 will have dominant components V_1, V_3 of sizes $\sim n\varepsilon$ and $\sim 2n\varepsilon$. As G_3 has no middle ground in its component sizes, $V_1 \subseteq V_3$. Now the sprinkling G_2 adds $\sim p_2 \binom{n\varepsilon}{2} \sim n\varepsilon^3/2$ edges internal to V_1 . Thus V_3 will have complexity at least $n\varepsilon^3/2 = \lambda^3/2$, which approaches infinity.

18.2.10 11.10 The Critical Window

We now fix a real λ and set $p = 1/n + \lambda n^{-4/3}$. There has been massive study of this critical window, Łuczak (1990b) and the monumental Janson, Knuth, Łuczak and Pittel (1993) being only two examples. Calculations in this regime are remarkably delicate.

Fix $c > 0$ and let X be the number of tree components of size $k = cn^{2/3}$. Then

$$\mathbb{E}[X] = \binom{n}{k} k^{k-2} p^{k-1} (1-p)^{k(n-k) + \binom{k}{2} - (k-1)}$$

Recall that $\ln(1+x) = x - \frac{1}{2}x^2 + O(x^3)$ and watch the terms cancel!

$$\binom{n}{k} = \frac{(n)_k}{k!} \sim \frac{n^k e^k}{k^k \sqrt{2\pi k}} \prod_{i=1}^{k-1} \left(1 - \frac{i}{n}\right).$$

For $i = 0, \dots, k-1$,

$$-\ln\left(1 - \frac{i}{n}\right) = \frac{i}{n} + \frac{i^2}{2n^2} + O\left(\frac{i^3}{n^3}\right),$$

so that

$$\sum_{i=1}^{k-1} -\ln\left(1 - \frac{i}{n}\right) = \frac{k^2}{2n} + \frac{k^3}{6n^2} + o(1) = \frac{k^2}{2n} + \frac{c^3}{6} + o(1).$$

Considering $p^{k-1} = n^{1-k} (1 + \lambda n^{-1/3})^{k-1}$,

$$(k-1) \ln(1 + \lambda n^{-1/3}) = k\lambda n^{-1/3} - \frac{1}{2}c\lambda^2 + o(1)$$

Last but not least,

$$\ln(1-p) = -p + O(n^{-2}) = -\frac{1}{n} - \frac{\lambda}{n^{4/3}} + O(n^{-2})$$

and

$$k(n-k) + \binom{k}{2} - (k-1) = kn - \frac{k^2}{2} + O(n^{2/3})$$

so that

$$\left[k(n-k) + \binom{k}{2} - (k-1) \right] \ln(1-p) = -k + \frac{k^2}{2n} - \frac{\lambda k}{n^{1/3}} + \frac{\lambda c^2}{2} + o(1)$$

Putting it all together,

$$\mathbb{E}[X] \sim \frac{n^k k^{k-2}}{k^k \sqrt{2\pi k n^{k-1}}} e^A = n k^{-5/2} (2\pi)^{-1/2} e^A$$

where

$$A = k - \frac{k^2}{2n} - \frac{c^3}{6} + \frac{\lambda k}{n^{1/3}} - \frac{\lambda^2 c}{2} - k + \frac{k^2}{2n} - \frac{\lambda k}{n^{1/3}} + \frac{\lambda c^2}{2}.$$

The k and n terms cancel and we can give A the intriguing form

$$A = A(c) = \frac{(\lambda - c)^3 - \lambda^3}{6}$$

Writing k in terms of n then yields

$$\mathbb{E}[X] \sim n^{-2/3} e^{A(c)} c^{-5/2} (2\pi)^{-1/2}$$

For any particular such k , $\mathbb{E}[X] \rightarrow 0$ but if we sum k between $cn^{2/3}$ and $(c + dc)n^{2/3}$ we multiply by $n^{2/3}dc$. Going to the limit gives an integral: For any fixed a, b, λ let X be the number of tree components of size between $an^{2/3}$ and $bn^{2/3}$. Then

$$\lim_{n \rightarrow \infty} \mathbb{E}[X] = \int_a^b e^{A(c)} c^{-5/2} (2\pi)^{-1/2} dc$$

The large components are not all trees. Wright (1977) proved that for fixed $l \geq 0$ there are asymptotically $c_l k^{k-2+(3l/2)}$ connected graphs on k points with $k - 1 + l$ edges, where $c_0 = 1$, $c_1 = \sqrt{\pi/8}$ and c_l was given by a specific recurrence. Asymptotically in l , $c_l = l^{-l/2(1+o(1))}$. The calculation for $X^{(l)}$, the number of such components on k vertices, leads to extra factors of $c_l k^{3l/2}$ and n^{-l} , which gives $c_l c^{3l/2}$. For fixed a, b, λ, l the number $X^{(l)}$ of components of size between $an^{2/3}$ and $bn^{2/3}$ with $l - 1$ more edges than vertices satisfies

$$\lim_{n \rightarrow \infty} \mathbb{E}[X^{(l)}] = \int_a^b e^{A(c)} c^{-5/2} (2\pi)^{-1/2} (c_l c^{3l/2}) dc$$

and letting X^* be the total number of components of size between $an^{2/3}$ and $bn^{2/3}$

$$\lim_{n \rightarrow \infty} \mathbb{E}[X^*] = \int_a^b e^{A(c)} c^{-5/2} (2\pi)^{-1/2} g(c) dc$$

where

$$g(c) = \sum_{l=0}^{\infty} c_l c^{3l/2}$$

a sum convergent for all c . A component of size $\sim cn^{2/3}$ will have probability $c_l c^{3l/2}/g(c)$ of having complexity l , independent of λ . As $\lim_{c \rightarrow 0} g(c) = 1$, most components of size $\epsilon n^{2/3}$, $\epsilon \ll 1$, are trees but as c gets bigger the distribution on l moves inexorably higher.

An Overview. For any fixed λ the sizes of the largest components are of the form $cn^{2/3}$ with a distribution over the constant. This distribution has support the positive reals. Thus, for example, for $\lambda = -4$ there is some positive limiting probability that the largest component is bigger than $10n^{2/3}$ and for $\lambda = +4$ there is some positive limiting probability that the largest component is smaller than $0.1n^{2/3}$, though both these probabilities are miniscule. The $c^{-5/2}$ term dominates the integral as $c \rightarrow 0^+$, reflecting the notion that for any fixed λ there should be many components of size near $\epsilon n^{2/3}$ for $\epsilon = \epsilon(\lambda)$ appropriately small. When λ is large

negative (e.g., $\lambda = -4$) the largest component is likely to be $\epsilon n^{2/3}$, ϵ small, and there will be many components of nearly that size. The nontree components will be a negligible fraction of the tree components. When λ is large positive (e.g., $\lambda = +4$) the dominant component will have begun to emerge. The largest component is likely to be $\sim 2\lambda n^{2/3}$ and of moderately high (not zero or one) complexity and the second largest component will be considerably smaller and simple.

Now consider the evolution of $G(n, p)$ in terms of λ . Suppose that at a given λ there are components of size $c_1 n^{2/3}$ and $c_2 n^{2/3}$. When we move from λ to $\lambda + d\lambda$ there is a probability $c_1 c_2 d\lambda$ that they will merge. Components have a peculiar gravitation in which the probability of merging is proportional to their sizes. With probability $(c_1^2/2) d\lambda$ there will be a new internal edge in a component of size $c_1 n^{2/3}$ so that large components rarely remain trees. Simultaneously, big components are eating up other vertices.

With $\lambda = -4$, say, we have feudalism. Many small components (castles) are each vying to be the largest. As λ increases the components increase in size and a few large components (nations) emerge. An already large France has much better chances of becoming larger than a smaller Andorra. The largest components tend strongly to merge and by $\lambda = +4$ it is very likely that a dominant component, a Roman Empire, has emerged. With high probability this component is nevermore challenged for supremacy but continues absorbing smaller components until full connectivity One World - is achieved.

18.2.11 11.11 Analogies to Classical Percolation Theory

The study of percolation has involved the intense efforts of both mathematicians and physicists for many years. A central object of that study has been bond percolation on $\mathbb{Z}^d$, as described below. Here we explore, without proofs, the fruitful analogies between that percolation and the Erdős-Rényi phase transition. A classic text in this field is Percolation by Grimmett (1999) and we shall follow its treatment.

Let $d \geq 2$. (All parameters below shall depend on the choice of d .) Let $\mathbb{Z}^d$, as usual, represent the set of $\vec{a} = (a_1, \dots, a_d)$ with a_i integers. The d -dimensional cubic lattice, written L^d , is that graph with vertices $\mathbb{Z}^d$, two vertices $\vec{a}, \vec{b}$ being adjacent if they agree on $d - 1$ coordinates and differ by one on the other coordinate. Let $p \in [0, 1]$. The random subgraph $L^d(p)$ contains each edge of L^d (and no others) with independent probability p . We let $C(\vec{a})$ denote the connected component of $L^d(p)$ containing the vertex $\vec{a}$. We generally examine $C(\vec{0})$ as, by symmetry, all $C(\vec{a})$ look the same. [In Grimmett (1999) and elsewhere the edges of L^d are called bonds and they are open with probability p and closed otherwise. The word cluster is used in place of connected component.] Naturally, as p becomes larger $L^d(p)$ will have more adjacencies. There is a critical probability, denoted by p_c , at which $L^d(p)$ undergoes a macroscopic change.

- For $p < p_c$, the subcritical region, all connected components are finite.
- For $p > p_c$, the supercritical region, there is precisely one infinite component.
- For $p = p_c$, at the critical point, the situation is particularly delicate, as discussed below.

The constant probabilities of bond percolation correspond to the parametrized probabilities $p = c/n$ in Erdős-Rényi's $G(n, p)$. The value $c = 1$ is then the critical probability in the Erdős-Rényi model. The infinite component in the bond percolation model is analogous to giant components, components of size $\Omega(n)$, in the Erdős-Rényi model. The finite components in the bond percolation model are analogous to components of size $O(\ln n)$ in the Erdős-Rényi model.

The uniqueness of the infinite component in bond percolation was an open question (though the physicists "knew" it was true!) for many years. It was proved by Aizenman, Kesten and

Newman (1987) and the Book Proof is given by Burton and Keane (1989). This corresponds to the uniqueness of the giant component in $G(n, p)$.

In the bond percolation model there are only three choices for p : it can be less than, greater than, or equal to p_c . The barely subcritical and barely supercritical phases of the Erdős-Rényi model correspond to an asymptotic study of the bond percolation model as p approaches p_c from below and from above, respectively. This study is done through the use of critical exponents as described below.

Set $\theta(p) = \Pr[C(\vec{0}) \text{ is infinite}]$. For $p < p_c$, $\theta(p) = 0$ as with probability one there are no infinite components. For $p > p_c$, $\theta(p) > 0$. This corresponds to the infinite component having positive density, strengthening the analogy to the giant components of the Erdős-Rényi model. When p is barely greater than p_c there will be an infinite component but its density will be very small. The critical exponent β is that real number so that

$$\theta(p) = (p - p_c)^{\beta+o(1)} \text{ as } p \rightarrow p_c^+$$

[As mathematicians, we are aware that $\theta(p)$ could behave erratically as $p \rightarrow p_c^+$ and β might not exist. This holds for all critical exponents we discuss. For a physicist, there is no doubt that the critical exponents do exist, and they can tell you the values to a few decimal places!] Analogously, in the Erdős-Rényi model $\theta(c)$ is the proportion of points in the giant component, that $y = y(c) > 0$ satisfying (11.2). From (11.5), $y(1 + \varepsilon) \sim 2\varepsilon$ as $\varepsilon \rightarrow 0^+$. Therefore $\beta = 1$.

The susceptibility, denoted by $\chi(p)$ (not to be confused with chromatic number), is given by $\chi(p) = E[|C(\vec{0})|]$. For $p > p_c$, $\chi(p) = \infty$ as with positive probability $C(\vec{0})$ is infinite. For $p < p_c$, $\chi(p)$ is finite and $\chi(p) \rightarrow \infty$ as $p \rightarrow p_c^-$. That the susceptibility approaches infinity at the same critical value for which an infinite component appears is not at all obvious and was one of the great developments of the field, due independently to Aizenman and Barsky (1987) and Men'shikov (1986). When p is barely less than p_c , $\chi(p)$ will be finite but large. The critical number γ is that real number so that

$$\chi(p) = (p_c - p)^{-\gamma+o(1)} \text{ as } p \rightarrow p_c^-$$

Analogously, in the Erdős-Rényi model we examine $E[|C(v)|]$ in $G(n, (1 - \varepsilon)/n)$. In the subcritical region this is well mirrored by $T_{1-\varepsilon}$, the total size of a subcritical Poisson branching process. We find $E[T_{1-\varepsilon}]$ by looking at each generation. There is one root Eve, who has an expected number $1 - \varepsilon$ children. They behave similarly and so Eve has an expected number $(1 - \varepsilon)^2$ grandchildren. This continues; there are an expected number $(1 - \varepsilon)^i$ nodes in the i th generation so that

$$E[T_{1-\varepsilon}] = \sum_{i=0}^{\infty} (1 - \varepsilon)^i = \varepsilon^{-1}$$

precisely. Therefore $\gamma = 1$.

While $\chi(p)$ is infinite in the supercritical region we can examine the "finite portion" of $L^d(p)$. The finite susceptibility χ^f is given by

$$\chi^f(p) = E[|C(\vec{0})| \mid C(\vec{0}) \text{ is finite}]$$

When p is barely greater than p_c , $\chi^f(p)$ will be finite but large. The critical number γ' is the real number satisfying

$$\chi^f(p) = (p - p_c)^{-\gamma'+o(1)} \text{ as } p \rightarrow p_c^+$$

The Erdős-Rényi analogue is $E[|C(v)|]$ in $G(n, (1 + \varepsilon)/n)$, conditional on v not being in the giant component. In $G(n, (1 + \varepsilon)/n)$, $|C(v)|$ has basically distribution $T_{1+\varepsilon}^{po}$, with the value

$T_{1+\varepsilon}^{po} = \infty$ corresponding to being in the giant component. The finite analogue then corresponds to $T_{1+\varepsilon}^{po}$, conditional on it being finite. The probability that $T_{1+\varepsilon}^{po}$ is finite approaches one as $\varepsilon \rightarrow 0^+$. The Poisson branching processes $T_{1+\varepsilon}^{po}, T_{1-\varepsilon}^{po}$ have nearly the same finite distribution. Conditioning on v not being in the giant component, $|C(v)|$ then behaves like $T_{1-\varepsilon}^{po}$. Therefore $\gamma' = 1$.

At the critical value $p = p_c$, all components are finite. The distribution of $|C(\vec{0})|$ will have a heavy tail. The critical number δ is that real number so that at $p = p_c$

$$\Pr[|C(\vec{0})| \geq s] = s^{-1/\delta+o(1)} \text{ as } s \rightarrow \infty$$

For the Erdős-Rényi analogue we consider $|C(v)|$ in $G(n, 1/n)$. One needs to be cautious about the double limit. For any fixed s ,

$$\lim_{n \rightarrow \infty} \Pr[|C(v)| \geq s] = \Pr[T_1^{po} \geq s] = \Theta(s^{-1/2})$$

from (11.3). Therefore $\delta = 2$.

We further examine the gap exponent, denoted by Δ . In the subcritical region the distribution of $|C(\vec{0})|$ drops off exponentially. For each $k \geq 1$ it has a finite k th moment. The hypothetical quantity Δ is such that

$$\frac{\mathbb{E}[|C(\vec{0})|^{k+1}]}{\mathbb{E}[|C(\vec{0})|^k]} = (p_c - p)^{-\Delta+o(1)}$$

The belief is that Δ does not depend on the choice of k . In the supercritical region the belief is that the same asymptotics hold when the infinite component is erased. More precisely, the belief is that

$$\frac{\mathbb{E}[|C(\vec{0})|^{k+1} \mid C(\vec{0}) \text{ is finite}]}{\mathbb{E}[|C(\vec{0})|^k \mid C(\vec{0}) \text{ is finite}]} = (p_c - p)^{-\Delta+o(1)}$$

for all $k \geq 1$. In the Erdős-Rényi analogue the distribution of $C(v)$ in $G(n, (1 - \varepsilon)/n)$ mirrors that of $T_{1-\varepsilon}^{po}$. [The supercritical $G(n, (1 + \varepsilon)/n)$, with its giant component erased, behaves similarly.] From Section 11.4, $\Pr[T_{1-\varepsilon}^{po} = s]$ drops like $s^{-3/2}$ until k reaches $\Theta(\varepsilon^{-2})$ when it begins its exponential drop-off. The region of exponential drop-off has negligible effect on the finite moments. The k th moment of $T_{1-\varepsilon}^{po}$ is basically the sum of $s^{-3/2}s^k$ for $s = O(\varepsilon^{-2})$, which is of order $(\varepsilon^{-2})^{k+1/2}$, or ε^{-2k-1} . The ratio of the $(k + 1)$ st and k th moments is then $\Theta(\varepsilon^{-2})$. Therefore $\Delta = 2$.

For bond percolation in $\mathbb{Z}^d$ define the triangle function

$$T(p) = \sum_{\vec{x}, \vec{y} \in \mathbb{Z}^d} \Pr[\vec{0} \leftrightarrow \vec{x}] \Pr[\vec{0} \leftrightarrow \vec{y}] \Pr[\vec{x} \leftrightarrow \vec{y}],$$

where $\vec{x} \leftrightarrow \vec{y}$ means that $\vec{x}, \vec{y}$ lie in the same component. In Aizenman and Newman (1984) the following condition was introduced.

The Triangle Condition: $T(p_c) < \infty$.

They showed that when the triangle condition holds some of the conditions for mean field theory (as discussed below) are valid. For the Erdős-Rényi percolation analogue we fix a vertex v of $G(n, p)$ and define the discrete triangle function

$$T(p) = \sum_{x, y} \Pr[v \leftrightarrow x] \Pr[v \leftrightarrow y] \Pr[x \leftrightarrow y]$$

The critical probability p_c is replaced by $p = n^{-1}$. Finiteness is replaced by the boundedness giving the following.

The Discrete Triangle Condition: $T(p) = O(1)$.

The contribution to $T(p)$ when two or three of v, x, y are equal is easily bounded, leaving the contribution from all triples v, x, y of distinct vertices. As all pairs behave the same and there are $(n-1)(n-2) \sim n^2$ terms,

$$T(p) \sim O(1) + n^2 \Pr[v \leftrightarrow x]^3$$

and

$$\Pr[v \leftrightarrow x] = \sum_t \Pr[|C(v)| = t] \frac{t-1}{n-1} \sim n^{-1} \sum_t t \Pr[|C(v)| = t].$$

We know that $\Pr[|C(v)| = t]$ behaves like $t^{-3/2}$ until t reaches $\Theta(n^{2/3})$ and then drops off exponentially. Ignoring constants,

$$\sum_t t \Pr[|C(v)| = t] = \Theta \left(\sum_{t=\Theta(n^{2/3})} t t^{-3/2} \right) = \Theta \left((n^{2/3})^{1/2} \right)$$

Now $\Pr[v \leftrightarrow x] = \Theta(n^{-2/3})$. [Basically, the main contribution to $\Pr[v \leftrightarrow x]$ comes when v lies in a component of size $\Theta(n^{2/3})$, even though that rarely occurs.] The triangle condition does hold as

$$T(p) = O(1) + O(n^2) \Theta(n^{-2/3})^3 = O(1)$$

The discrete triangle condition does not hold in the barely supercritical region. There $\Pr[v \leftrightarrow x]$ is dominated by the probability that both v, w lie in the dominant component. As the dominant component has size $\gg n^{2/3}$, $\Pr[v \leftrightarrow x] \gg n^{-2/3}$, and $T(p) \gg 1$. This is not mere serendipity. Rather, the boundedness of $T(p)$ provides a natural boundary between the critical window and the barely supercritical region for discrete random structures. This connection is explored in depth in Borgs, Chayes, van der Hofstad, Slade and Spencer (2005) and the recent lecture notes of Slade (2006).

Hara and Slade (1990) [see also the survey by Hara and Slade (1994)] proved that the triangle condition holds in the bond percolation model for sufficiently high dimensions d . [More precisely, they showed that $T(p)$ could be made very small by taking p slightly less than p_c . Their argument works for $d \geq 19$ and for all $d > 6$ with a somewhat different model. It is strongly believed that the condition does hold for all $d > 6$.] Building on that, they found that the critical exponent values $\beta = 1$, $\gamma = \gamma' = 2$, $\delta = 2$ and $\Delta = 2$ hold for those d . Mathematical physicists have a term mean field, which, quoting Grimmett, "permits several interpretations depending on context." A commonly held requirement is that the critical exponents have the values given above. Thus bond percolation for $d \geq 19$ is regarded as exhibiting mean field behavior. Using the analogues described above it seems reasonable to say that the Erdős-Rényi model exhibits mean field behavior.

18.2.12 11.12 Exercises

1. Consider the Poisson branching model with mean $c = 1$ and root Eve. For $n \geq 3$ let A_n be the event that Eve has precisely two children, Dana and Fan, and that the total tree size $T = n$. Let X be the size of the subtree with root Dana. For each $j \geq 1$ find $\lim_{n \rightarrow \infty} \Pr[X = j \mid A_n]$. Find an asymptotic formula for $\Pr[n/3 < X < 2n/3]$.

2. Consider the binomial branching model with parameters m, p and $mp > 1$. Set $y = y(m, p) = \Pr[T = \infty]$. Give an implicit equation for y analogous to (11.2). With m fixed set $mp = (1 + \epsilon)$. Find

$$\lim_{\epsilon \rightarrow 0^+} \frac{y(m, p)}{\epsilon}$$

3. Let $c > 1$. Let $Z_i, i = 1, 2, \dots$ be independent Poisson variables with mean c . For $a > 1$ consider the walk defined by initial condition $Y_1 = a$ and recursion $Y_t = Y_{t-1} + Z_i - 1$ for $t \geq 2$. Use Chernoff bounds to show

$$\lim_{a \rightarrow \infty} \sum_{t \geq 2} \Pr[Y_t \leq 0] = 0$$

Use this to show that the random walk defined by initial condition $Y_0 = 1$ and recursion $Y_t = Y_{t-1} + Z_t - 1$ for $t \geq 1$ has a positive probability of being positive for all t .

4. An Openended Computer Experiment. Begin with vertices $1, \dots, n$ ($n = 10^6$ is very quick when done right) and no edges. Each round pick two random vertices and add an edge between them. Use a Union-Find algorithm to keep track of the components and the component sizes. Parametrize round number E by $E/\binom{n}{2} = 1/n + \lambda n^{-4/3}$ and concentrate on the region $-4 \leq \lambda \leq +4$. Update the ten largest component sizes, noting particularly when two of the ten largest components merge. Watch the barely subcritical picture at $\lambda = -4$

turn into a barely supercritical picture at $\lambda = +4$ as the bulk of the moderate size components merge to form a dominant component.

The Probabilistic Lens: the Rich Get Richer

Consider two bins, each of which initially has one ball. At each time $u = 1, 2, \dots$ we add one ball to one of the bins. The ball is placed randomly, in proportion to the square of the number of balls already in the bin. [For example, if the bins have 5 balls and 3 balls, respectively, the next ball is placed in the bin with 5 balls with probability $25/(25 + 9)$.]

Theorem 1 With probability one, one of the bins will get all but a finite number of the balls.

Proof. We move to a continuous time model. Let X_i be independent random variables, X_i having the exponential distribution with mean i^{-2} . (That is, X_i has density function $i^2 e^{-ti^2}$ for $t \geq 0$.) At time zero the first bin has one ball. It receives its second ball at time X_1 . In general, it receives its i th ball time X_i after receiving its $(i - 1)$ st ball. Let X'_i also be independent exponential distributions with mean i^{-2} , independently chosen from the X_i . The second bin receives its balls according to the X'_i . The process ends when an infinite number of balls have been placed. The fictitious continuation, of defining the X_i, X'_i for all $i \geq 1$, shall be helpful in the analysis.

We use two basic properties of exponential distributions. Both are easy calculus exercises.

Proposition 2 Let X be exponential with mean μ and let $a > 0$. Then $X - a$, conditional on $X \geq a$, is also exponential with mean μ .

This is often called the forgetfulness property.

Proposition 3 Let X, X' be independent exponentials with means μ, ν , respectively. Then

$$\Pr[\min\{X, X'\} = X] = \frac{\mu^{-1}}{\mu^{-1} + \nu^{-1}}$$

The continuous time process mirrors the sequential process. Clearly the first ball is equally likely to go into either of the two bins. Suppose at some time $t > 0$ the first (say) bin has just

received its i th ball and the second bin last received its j th ball at time $t' < t$. (When the second bin has not yet received its second ball set $j = 1$ and $t' = 0$.) The waiting time for the first bin is then X_i . The waiting time for the second was X_j at time t' . By the forgetfulness property its conditional waiting time at time t is X_j^* , exponential with mean j^{-2} . The next ball goes into the first bin if and only if $\min\{X_i, X_j^*\} = X_i$, which occurs with probability $i^2/(i^2 + j^2)$ as desired.

Let $T = \sum_{i=1}^{\infty} X_i, T' = \sum_{i=1}^{\infty} X'_i$ be the total times for the bins to receive (under fictitious continuation) an infinite number of balls. As $E[X_i] = E[X'_i] = i^{-2}$ and (critically!) $\sum_{i=1}^{\infty} i^{-2}$ converges, both T, T' have finite means and so are finite with probability one. As sums of independent continuous distributions $\Pr[T = T'] = 0$. Suppose $T < T'$, the other case being identical. At time T the first bin has received an infinite number of balls. The second bin has not. Therefore the second bin has received only a finite number of balls!

18.3 12 Circuit Complexity

Abstract

It is not knowledge, but the act of learning, not possession but the act of getting there, which grants the greatest enjoyment. When I have clarified and exhausted a subject, then I turn away from it, in order to go into darkness again; the never-satisfied man is so strange - if he has completed a structure then it is not in order to dwell in it peacefully, but in order to begin another. I imagine the world conqueror must feel thus, who, after one kingdom is scarcely conquered, stretches out his arms for another.

- Karl Friedrich Gauss

18.3.1 12.1 Preliminaries

A Boolean function $f = f(x_1, \dots, x_n)$ on the n variables $x_1, x_2, \dots, x_n$ is simply a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$. In particular, $0, 1, x_1 \wedge \dots \wedge x_n, x_1 \vee \dots \vee x_n$ and $x_1 \oplus \dots \oplus x_n$ denote, as usual, the two constant functions, the And function (whose value is 1 iff $x_i = 1$ for all i), the OR function (whose value is 0 iff $x_i = 0$ for all i) and the parity function (whose value is 0 iff an even number of variables x_i is 1), respectively. For a function f , we let $\bar{f} = f \oplus 1$ denote its complement Not f . The functions x_i and $\bar{x}_i$ are called atoms. In this section we consider the problem of computing various Boolean functions efficiently. A circuit is a directed, acyclic graph, with a special vertex with no outgoing edges called the output vertex. Every vertex is labeled by a Boolean function of its immediate parents and the vertices with no parents (i.e., those with no ingoing edges) are labeled either by one of the variables x_i or by a constant 0 or 1. For every assignment of binary values to each variable x_i one can compute, recursively, the corresponding value of each vertex of the circuit by applying the corresponding function labeling it to the already computed values of its parents. We say that the circuit computes the function $f = f(x_1, \dots, x_n)$ if for each $x_i \in \{0, 1\}$, the corresponding value of the output vertex of the circuit equals $f(x_1, \dots, x_n)$. For example, Figure 12.1 presents a circuit computing $f(x_1, x_2, x_3) = (x_1 \oplus (x_2 \wedge x_3)) \wedge x_1$.




Fig. 12.1 A binary circuit for $\left(x_1 \oplus (x_2 \wedge x_3)\right) \wedge x_1$.

If every fanout in a circuit is at most one (i.e., the corresponding graph is a tree) the circuit is called a formula. If every fanin in a circuit is at most two the circuit is called a binary circuit. Therefore the circuit in Figure 12.1 is binary, but it is not a formula. The size of a circuit is the number of vertices in it and its depth is the maximum length (number of edges) of a directed path in it. The binary circuit complexity of a Boolean function is the size of the smallest binary circuit computing it. An easy counting argument shows that for large n the binary circuit complexity of almost all the functions of n variables is at least $(1 + o(1))2^n/n$. This is because the number of binary circuits of size s on n variables can be shown to be less than $(c(s + n))^s$, whereas the total number of Boolean functions on n variables is 2^{2^n} . On the other hand, there is no known nonlinear, not to mention exponential (in n), lower bound for the binary circuit complexity of any "explicit" function. By "explicit" here we mean an NP-function, that is, one of a family $\{f_{n_i}\}_{i \geq 1}$ of Boolean functions, where f_{n_i} has n_i variables, $n_i \rightarrow \infty$, and there is a nondeterministic Turing machine that, given n_i and $x_1, \dots, x_{n_i}$, can decide in (nondeterministic) polynomial time (in n_i) if $f_{n_i}(x_1, \dots, x_{n_i}) = 1$. [An example for such a family is the $(n/2)$ -clique function; here $n_i = \binom{i}{2}$, the n_i variables $x_1, \dots, x_{n_i}$ represent the edges of a graph on i vertices and $f_{n_i}(x_1, \dots, x_{n_i}) = 1$ iff the corresponding graph contains a clique on at least $i/2$ vertices.] Any nonpolynomial lower bound for the binary circuit complexity of an explicit function would imply (among other things) that $\mathbf{P} \neq \mathbf{NP}$ and thus solve the arguably most important open problem in theoretical computer science. Unfortunately, at the moment, the best known lower bound for the binary circuit complexity of an explicit function of n variables is only $3n$ [see Blum (1984), Paul (1977)]. However, several nontrivial lower bounds are known when we impose certain restrictions on the structure of the circuits. Most of the known proofs of these bounds rely heavily on probabilistic methods. In this chapter we describe some of these results. We note that there are many additional beautiful known results about circuit complexity; see, for example, Wegener (1987) and Karchmer and Wigderson (1990), but those included here are not only among the crucial ones, but also represent the elegant methods used in this field. Since most results in this chapter are asymptotic we assume, throughout the chapter, whenever it is needed, that the number of variables we have is sufficiently large.

18.3.2 12.2 Random Restrictions and Bounded-depth Circuits

Let us call a Boolean function G a t -And-Or if it can be written as an And of an arbitrary number of functions, each being an Or of at most t atoms; that is, $G = G_1 \wedge \cdots \wedge G_w$, where $G_i = y_{i1} \vee \cdots \vee y_{ia_i}$, $a_i \leq t$ and each y_{ij} is an atom. Similarly, we call a Boolean function an s -Or-And, if it can be written as an Or of And gates each containing at most s atoms. A minterm of a function is a minimal assignment of values to some of the variables that forces the function to be 1. Its size is the number of variables whose values are set. Note that a function is an s -OrAnd if and only if each of its minterms is of size at most s . A restriction is a map ρ of the set of indices $\{1, \dots, n\}$ to the set $\{0, 1, *\}$. The restriction of the function $G = G(x_1, \dots, x_n)$ by ρ , denoted by $G|_\rho$, is the Boolean function obtained from G by setting the value of each x_i for $i \in \rho^{-1}\{0, 1\}$ to $\rho(i)$, and leaving each x_j for $j \in \rho^{-1}\{*\}$ as a variable. Thus, for example, if $G(x_1, x_2, x_3) = (x_1 \wedge x_2) \vee x_3$ and $\rho(1) = 0, \rho(2) = \rho(3) = *$, then $G|_\rho = x_3$. For $0 \leq p \leq 1$, a random p -restriction is a random restriction ρ defined by choosing for each $1 \leq i \leq n$ independently the value of $\rho(i)$ according to the following distribution:

$$\Pr[\rho(i) = *] = p, \quad \Pr[\rho(i) = 0] = \Pr[\rho(i) = 1] = (1 - p)/2. \quad (12.1)$$

Improving the results of Furst, Saxe and Sipser (1984), Ajtai (1983) and Yao (1985), Håstad (1988) proved the following result, which is very useful in establishing lower bounds for bounded-depth circuits.

Lemma 12.2.1 [The Switching Lemma] Let $G = G(x_1, \dots, x_n)$ be a t -And-Or, that is, $G = G_1 \wedge G_2 \wedge \cdots \wedge G_w$, where each G_i is an Or of at most t atoms. Let ρ be the random restriction defined by (12.1). Then

$$\begin{aligned} & \Pr \left[G|_\rho \text{ is not an } (s-1)\text{-Or-AND} \right] \\ &= \Pr \left[G|_\rho \text{ has a minterm of size } \geq s \right] \leq (5pt)^s. \end{aligned}$$

Proof. Let E_s be the event that $G|_\rho$ has a minterm of size at least s . To bound $\Pr[E_s]$, we prove a stronger result; for any Boolean function F ,

$$\Pr[E_s | F|_\rho \equiv 1] \leq (5pt)^s \quad (12.2)$$

Here we agree that if the condition is unsatisfied then the conditional probability is 0. Lemma 12.2.1 is obtained from (12.2) by taking $F \equiv 1$. We prove (12.2) by induction on w . For $w = 0$, $G \equiv 1$ and there is nothing to prove. Assuming (12.2) holds whenever the number of G_i is less than w , we prove it for w . Put $G = G_1 \wedge G^*$, where $G^* = G_2 \wedge \cdots \wedge G_w$, and let E_s^* be the event that $G^*|_\rho$ has a minterm of size at least s . By interchanging, if necessary, some of the variables with their complements we may assume, for convenience, that $G_1 = \bigvee_{i \in T} x_i$, where $|T| \leq t$. Either $G_1|_\rho \equiv 1$ or $G_1|_\rho \not\equiv 1$. In the former case, E_s holds if and only if E_s^* holds and hence, by induction,

$$\Pr[E_s | F|_\rho \equiv 1, G_1|_\rho \equiv 1] = \Pr[E_s^* | (F \wedge G_1)|_\rho \equiv 1] \leq (5pt)^s \quad (12.3)$$

The case $G_1|_\rho \not\equiv 1$ requires more work. In this case, any minterm of $G|_\rho$ must assign a value 1 to at least one x_i , for $i \in T$. For a nonempty $Y \subseteq T$ and for a function $\sigma : Y \rightarrow \{0, 1\}$ that is not identically 0, let $E_s(Y, \sigma)$ be the event that $G|_\rho$ has a minterm of size at least s which assigns the value $\sigma(i)$ to x_i for each $i \in Y$ and does not assign any additional values to variables x_j with $j \in T$. By the preceding remark,

$$\Pr[E_s | F|_\rho \equiv 1, G_1|_\rho \not\equiv 1] \leq \sum_{Y, \sigma} \Pr[E_s(Y, \sigma) | F|_\rho \equiv 1, G_1|_\rho \not\equiv 1] \quad (12.4)$$

Observe that the condition $G_1|_\rho \not\equiv 1$ means precisely that $\rho(i) \in \{0, *\}$ for all $i \in T$ and hence, for each $i \in T$,

$$\Pr \left[\rho(i) = * | G_1|_\rho \not\equiv 1 \right] = \frac{p}{p + (1-p)/2} = \frac{2p}{1+p}$$

Thus, if $|Y| = y$,

$$\Pr \left[\rho(Y) = * | G_1|_\rho \not\equiv 1 \right] \leq \left(\frac{2p}{1+p} \right)^y$$

The further condition $F|_\rho \equiv 1$ can only decrease this probability. This can be shown using the FKG Inequality (see Chapter 6). It can also be shown directly as follows. For any fixed $\rho' : N - Y \rightarrow \{0, 1, *\}$, where $N = \{1, \dots, n\}$, we claim that

$$\Pr \left[\rho(Y) = * | F|_\rho \equiv 1, G_1|_\rho \not\equiv 1, \rho|_{N-Y} = \rho' \right] \leq \left(\frac{2p}{1+p} \right)^y$$

Indeed, the given ρ' has a unique extension ρ with $\rho(Y) = *$. If that ρ does not satisfy the above conditions, then the conditional probability is zero. If it does, then so do all extensions p with $p(i) \in \{0, *\}$ for $i \in Y$, and so the inequality holds in this case too. As this holds for all fixed ρ' we conclude that indeed

$$\Pr \left[\rho(Y) = * | F|_\rho \equiv 1, G_1|_\rho \not\equiv 1 \right] \leq \left(\frac{2p}{1+p} \right)^y \leq (2p)^y \quad (12.5)$$

Let $\rho' : T \rightarrow \{0, *\}$ satisfy $\rho(Y) = *$ and consider all possible restrictions ρ satisfying $\rho|_T = \rho'$. Under this condition, ρ may be considered as a random restriction on $N - T$. The event $F|_\rho \equiv 1$ reduces to the event $\tilde{F}|_{\rho|_{N-T}} \equiv 1$, where $\tilde{F}$ is the And of all functions obtained from F by substituting the values of x_i according to ρ' for those $i \in T$ with $\rho'(i) = 0$, and by taking all possibilities for all the other variables x_j for $j \in T$. If the event $E_s(Y, \sigma)$ occurs then $G^*|_{\rho\sigma}$ has a minterm of size at least $s - y$ that does not contain any variable x_i with $i \in T - Y$. But this happens if and only if $\tilde{G}|_{\rho|_{N-T}}$ has a minterm of size at least $s - y$, where $\tilde{G}$ is the function obtained from G^* by substituting the values of x_j for $j \in Y$ according to σ , the values of x_i for $i \in T - Y$ and $\rho'(i) = 0$ according to ρ' and by removing all the variables x_k with $k \in T - Y$ and $\rho'(k) = *$. Denoting this event by $\tilde{E}_{s-y}$ we can apply induction and obtain $\Pr \left[E_s(Y, \sigma) | F|_\rho \equiv 1, G_1|_\rho \not\equiv 1, \rho|_T = \rho' \right] \leq \Pr \left[\tilde{E}_{s-y} | \tilde{F}|_\rho \equiv 1 \right] \leq (5pt)^{s-y}$. Since any ρ with $F|_\rho \equiv 1, G_1|_\rho \equiv 1, \rho(Y) = *$ must have $\rho|_T = \rho'$ for some ρ' of this form, and since the event $E_s(Y, \sigma)$ may occur only if $\rho(Y) = *$ we conclude that

$$\Pr \left[E_s(Y, \sigma) | F|_\rho \equiv 1, G_1|_\rho \not\equiv 1, \rho(Y) = * \right] \leq (5pt)^{s-y}$$

and, by (12.5),

$$\begin{aligned} & \Pr \left[E_s(Y, \sigma) | F|_\rho \equiv 1, G_1|_\rho \not\equiv 1 \right] \\ &= \Pr \left[\rho(Y) = * | F|_\rho \equiv 1, G_1|_\rho \not\equiv 1 \right] \\ & \quad \cdot \Pr \left[E_s(Y, \sigma) | F|_\rho \equiv 1, G_1|_\rho \not\equiv 1, \rho(Y) = * \right] \\ & \leq (2p)^y (5pt)^{s-y} \end{aligned}$$

Substituting in (12.4) and using the fact that $|T| \leq t$ and that

$$\sum_{y=1}^t \frac{(2^y - 1) 2^y}{5^y y!} \leq \frac{2}{5} + \sum_{y=2}^{\infty} \frac{(4/5)^y}{y!} = \frac{2}{5} + e^{4/5} - 1 - \frac{4}{5} < 1$$

we obtain

$$\begin{aligned} & \Pr [E_s | F|_{\rho} \equiv 1, G_1|_{\rho} \not\equiv 1] \\ & \leq \sum_{y=1}^{|T|} \binom{|T|}{y} (2^y - 1) (2p)^y (5pt)^{s-y} \leq (5pt)^s \sum_{y=1}^t \frac{t^y}{y!} (2^y - 1) \left(\frac{2}{5t}\right)^y \\ & = (5pt)^s \sum_{y=1}^t (2^y - 1) \cdot \frac{2^y}{5^y \cdot y!} \leq (5pt)^s \end{aligned}$$

This, together with (12.3), gives

$$\Pr [E_s | F|_{\rho} \equiv 1] \leq (5pt)^s$$

completing the induction and the proof.

By taking the complement of the function G in Lemma 12.2.1 and applying De Morgan's rules one clearly obtains its dual form: If G is a t -OR-AND and ρ is the random restriction given by (12.1) then

$$\Pr [G|_{\rho} \text{ is not an } (s-1)\text{-AND-OR}] \leq (5pt)^s$$

We now describe an application of the switching lemma that supplies a lower bound to the size of circuits of small depth that compute the parity function $x_1 \oplus \cdots \oplus x_n$. We consider circuits in which the vertices are arranged in levels; those in the first level are atoms (i.e., variables or their complements) and each other gate is either an Or or an AND of an arbitrary number of vertices from the previous level. We assume that the gates in each level are either all And gates or all Or gates, and that the levels alternate between And levels and Or levels. A circuit of this form is called a $C(s, s', d, t)$ -circuit if it contains at most s gates, at most s' of which are above the second level, its depth is at most d and the fanin of each gate in its second level is at most t . Thus, for example, the circuit that computes the parity function by computing an Or of the 2^{n-1} terms $x_1^{\varepsilon_1} \wedge \cdots \wedge x_n^{\varepsilon_n}$, where $(\varepsilon_1, \dots, \varepsilon_n)$ ranges over all even binary vectors and $x_i^{\varepsilon_i} = x_i \oplus \varepsilon_i$, is a $C(2^{n-1} + 1, 1, 2, n)$ -circuit.

Theorem 12.2.2 Let $f = f(x_1, \dots, x_n)$ be a function and let C be a $C(\infty, s, d, t)$ circuit computing f , where $s(1/2)^t \leq 0.5$. Then either f or its complement $\bar{f}$ has a minterm of size at most $n - n / \lceil 2(10t)^{d-2} \rceil + t$.

Proof. Let us apply to C , repeatedly, $d-2$ times a random $(1/10t)$ -restriction. Each of these random restrictions, when applied to any bottom subcircuit of depth 2, transforms it by Lemma 12.2.1 with probability at least $1 - (1/2)^t$ from a t -Or-And to a t -And-OR (or conversely). If all these transformations succeed we can merge the new And gates with these from the level above them and obtain a circuit with a smaller depth. As the total size of the circuit is at most s and $s(1/2)^t \leq 0.5$, we conclude that with probability at least $\frac{1}{2}$, all transformations succeed and C is transformed into a $C(\infty, 1, 2, t)$ -circuit. Each variable x_i , independently, is still a variable (i.e., has not been assigned a value) with probability $1/(10t)^{d-2}$. Thus the number of remaining variables is a binomial random variable with expectation $n/(10t)^{d-2}$ and a little smaller variance. By the standard estimates for binomial distributions (see Appendix A) the probability that at least $n/2(10t)^{d-2}$ variables are still variables is more than $\frac{1}{2}$. Therefore, with positive probability, at most $n - n/2(10t)^{d-2}$ of the variables have been fixed and the resulting

restriction of f has a $C(\infty, 1, 2, t)$ -circuit; that is, its value can be fixed by assigning values to at most t additional variables. This completes the proof.

Corollary 12.2.3 For any $d \geq 2$, there is no

$$C\left(\infty, \frac{1}{2} \cdot 2^{(1/10)n^{1/(d-1)}}, d, \frac{1}{10}n^{1/(d-1)}\right)\text{-circuit}$$

that computes the parity function $f(x_1, \dots, x_n) = x_1 \oplus \dots \oplus x_n$.

Proof. Assuming there is such a circuit we obtain, by Theorem 12.2.2, that the value of f can be fixed by assigning values to at most $n - \frac{1}{2}n^{1/(d-1)} + \frac{1}{10}n^{1/(d-1)} < n$ variables. This is false, and hence there is no such circuit.

The estimate in Corollary 12.2.3 is, in fact, nearly the best possible. Since every $C(s, s', d, t)$ -circuit can be transformed into a $C((t+1)s, s, d+1, 2)$ -circuit (by replacing each atom by an Or or And of two copies of itself), Corollary 12.2.3 easily implies that the depth d of any $C(s, s', d, t)$ -circuit of polynomial size that computes the parity of n bits is at least $\Omega(\log n / \log \log n)$. This lower bound is also optimal.

18.3.3 12.3 More on Bounded-depth Circuits

In the previous section we saw that the parity function is hard to compute in small depth using And, Or and Not gates. It turns out that even if we allow the use of parity gates (in addition to the And, Or and Not gates) there are still some relatively simple functions that are hard to compute. Such a result was first proved by Razborov (1987). His method was modified and strengthened by Smolensky (1987). For an integer $k \geq 2$, let $\text{Mod}_k(x_1, x_2, \dots, x_n)$ be the Boolean function whose value is 1 iff $\sum x_i \not\equiv 0 \pmod{k}$. Smolensky showed that for every two powers p and q of distinct primes, the function Mod_p cannot be computed in a bounded-depth polynomial-size circuit that uses And, Or, Not and Mod_q gates. Here we present the special case of this result in which $q = 3$ and $p = 2$.

Let C be an arbitrary circuit of depth d and size s consisting of And, Or, Not and Mod_3 gates. A crucial fact, due to Razborov, is the assertion that the output of C can be approximated quite well (depending on d and s) by a polynomial of relatively small degree over $GF(3)$. This is proved by applying the probabilistic method as follows. Let us replace each gate of the circuit C by an approximate polynomial operation, according to the following rules, which guarantee that in each vertex in the new circuit we compute a polynomial over $GF(3)$, whose values are all 0 or 1 (whenever the input is a 0–1 input).

- (i) Each NOT gate $\bar{y}$ is replaced by the polynomial gate $(1 - y)$.
- (ii) Each Mod_3 gate $\text{Mod}_3(y_1, \dots, y_m)$ is replaced by the polynomial gate $(y_1 + y_2 + \dots + y_m)^2$.

The rule for replacement of Or and And gates is a little more complicated. Observe that in the two previous cases (i) and (ii) there was no approximation; the new gates compute precisely what the old ones did, for all possible Boolean values of the variables. This can, in principle, be done here too. An And gate $y_1 \wedge \dots \wedge y_m$ should simply be replaced by the product $y_1 \dots y_m$. An Or gate $y_1 \vee \dots \vee y_m$ can then be computed by De Morgan's rules. Since $y_1 \vee \dots \vee y_m = \overline{(\bar{y}_1 \wedge \dots \wedge \bar{y}_m)}$ and $\bar{y}$ is realized by $(1 - y)$, this would give

$$1 - (1 - y_1)(1 - y_2) \dots (1 - y_m) \quad (12.6)$$

The trouble is that this procedure would increase the degree of our polynomials too much. Hence we need to be a little more tricky. Let ℓ be an integer, to be chosen later. Given an Or gate $y_1 \vee \dots \vee y_m$, we choose ℓ random subsets $I_1, \dots, I_\ell$ of $\{1, \dots, m\}$, where for each $1 \leq i \leq \ell$ and for each $1 \leq j \leq m$ independently $\Pr[j \in I_i] = 1/2$. Observe that for each fixed i , $1 \leq i \leq \ell$,

the sum $\left(\sum_{j \in I_i} y_j\right)^2$ over $GF(3)$ is certainly 0 if $y_1 \vee \cdots \vee y_m = 0$ and is 1 with probability at least $1/2$ if $y_1 \vee \cdots \vee y_m = 1$. Hence, if we compute the OR function of the ℓ expressions $\left(\sum_{j \in I_i} y_j\right)^2$, $1 \leq i \leq \ell$, this function is 0 if $y_1 \vee \cdots \vee y_m = 0$ and is 1 with probability at least $1 - (1/2)^\ell$ if $y_1 \vee \cdots \vee y_m = 1$. We thus compute the Or and write it as a polynomial, in the way explained in equation (12.6). This gives

$$1 - \prod_{i=1}^{\ell} \left(1 - \left(\sum_{j \in I_i} y_j\right)^2\right). \quad (12.7)$$

Therefore in our new circuit we replace each Or gate by an approximation polynomial gate of the form described in (12.7). Once we have an approximation to an Or gate we can obtain the corresponding one for an And gate by applying De Morgan's rules. Since $y_1 \wedge \cdots \wedge y_m = (\overline{y_1} \vee \cdots \vee \overline{y_m})$ we replace each AND gate of the form $y_1 \wedge \cdots \wedge y_m$ by

$$\prod_{i=1}^{\ell} \left(1 - \left[\sum_{j \in I_i} (1 - y_j)\right]^2\right). \quad (12.8)$$

Observe that the polynomials in (12.7) and (12.8) are both of degree at most 2ℓ . Given the original circuit C of depth d and size s , we can now replace all its gates by our approximative polynomial gates and get a new circuit CP , which depends on all the random choices made in each replacement of each of the And/Or gates. The new circuit CP computes a polynomial $P(x_1, \dots, x_n)$ of degree at most $(2\ell)^d$. Moreover, for every fixed Boolean assignment to $x_1, x_2, \dots, x_n$, the probability that all the new gates compute exactly what the corresponding gates in C computed is at least $1 - s/2^\ell$. Therefore the expected number of inputs on which $P(x_1, \dots, x_n)$ is equal to the output of C is at least $2^n (1 - s/2^\ell)$. We have thus proved the following.

Lemma 12.3.1 For any circuit C of depth d and size s on n Boolean variables that uses Not, Or, And and Mod₃ gates and for any integer ℓ , there is a polynomial $P = P(x_1, \dots, x_n)$ of degree at most $(2\ell)^d$ over $GF(3)$ whose value is equal to the output of C on at least $2^n (1 - s/2^\ell)$ inputs.

In order to apply this lemma for obtaining lower bounds for the size of any circuit of the above type that computes the parity function, we need the following additional combinatorial result.

Lemma 12.3.2 For $n \geq 20$, there is no polynomial $P(x_1, \dots, x_n)$ over $GF(3)$ of degree at most $\sqrt{n}$ which is equal to the parity of $x_1, \dots, x_n$ for a set S of at least $0.9 \cdot 2^n$ distinct binary vectors $(x_1, \dots, x_n)$.

Proof. Suppose this is false, and suppose $S \subset \{0, 1\}^n$, $|S| \geq 0.9 \cdot 2^n$ and $P(x_1, \dots, x_n) = x_1 \oplus \cdots \oplus x_n$ for all $(x_1, \dots, x_n) \in S$. Define a polynomial $Q = Q(y_1, \dots, y_n)$ by

$$Q(y_1, \dots, y_n) = P(y_1 + 2, \dots, y_n + 2) - 2$$

and let

$$T = \{(y_1, \dots, y_n) \in \{1, -1\}^n : (y_1 + 2, \dots, y_n + 2) \in S\},$$

where all additions are in $GF(3)$. Clearly Q has degree at most $\sqrt{n}$ and satisfies $Q(y_1, \dots, y_n) = \prod_{i=1}^n y_i$ for all $(y_1, \dots, y_n) \in T$. Let now $G = G(y_1, \dots, y_n)$ be an arbitrary function from T to $GF(3)$. Extend it in an arbitrary way to a function from $(GF(3))^n$ to $GF(3)$ and write this function as a polynomial in n variables. [Trivially, any function from $(GF(3))^n$ to $GF(3)$ is a polynomial. This follows from the fact that it is a linear combination

of functions of the form $\prod_{i=1}^n (y_i - \varepsilon_i)(y_i - \varepsilon_i - 1)$, where $\varepsilon_i \in GF(3)$.] Replace each occurrence of y_i^2 in this polynomial by 1 to obtain a multilinear polynomial $\tilde{G}$ that agrees with G on T . Now replace each monomial $\prod_{i \in U} y_i$, where $|U| > n/2 + \sqrt{n}/2$ by $\prod_{i \notin U} y_i \cdot Q(y_1, \dots, y_n)$ and replace this new polynomial by a multilinear one, $\tilde{G}'$, again by replacing each y_i^2 by 1. Since for $y_i \in \{\pm 1\}$, $\prod_{i \notin U} y_i \cdot \prod_{i=1}^n y_i = \prod_{i \in U} y_i$, $\tilde{G}'$ is equal to G on T and its degree is at most $n/2 + \sqrt{n}/2$. However, the number of possible $\tilde{G}'$ is $3^{\sum_{i=0}^{n/2+\sqrt{n}/2} \binom{n}{i}} < 3^{0.88 \cdot 2^n}$, whereas the number of possible G is $3^{|T|} \geq 3^{0.9 \cdot 2^n}$. This is impossible, and hence the assertion of the lemma holds.

Corollary 12.3.3 There is no circuit of depth d and size $s \leq 0.1 \cdot 2^{(1/2)n^{1/2d}}$ computing the parity of $x_1, x_2, \dots, x_n$ using Not, And, Or and Mod 3 gates.

Proof. Suppose this is false and let C be such a circuit. Put $\ell = \frac{1}{2}n^{1/2d}$. By Lemma 12.3.1 there is a polynomial $P = P(x_1, \dots, x_n)$ over $GF(3)$, whose degree is at most $(2\ell)^d = \sqrt{n}$, which is equal to the parity of $x_1, \dots, x_n$ on at least $2^n \left(1 - s/2^{(1/2)n^{1/2d}}\right) \geq 0.9 \cdot 2^n$ inputs. This contradicts Lemma 12.3.2 and hence completes the proof.

18.3.4 12.4 Monotone Circuits

A Boolean function $f = f(x_1, \dots, x_n)$ is monotone if $f(x_1, \dots, x_n) = 1$ and $x_i \leq y_i$ imply $f(y_1, \dots, y_n) = 1$. A binary monotone circuit is a binary circuit that contains only binary And and Or gates. It is easy to see that a function is monotone if and only if there is a binary monotone circuit that computes it. The monotone complexity of a monotone function is the smallest size of a binary monotone circuit that computes it. Until 1985, the largest known lower bound for the monotone complexity of a monotone NP-function of n variables was $4n$. This was considerably improved in the fundamental paper of Razborov (1985), where a bound of $n^{\Omega(\log n)}$ to the k -clique function (which is 1 iff a given graph contains a clique of size k) is established. Shortly afterwards, Andreev (1985) used similar methods to obtain an exponential lower bound to a somewhat unnatural NP-function. Alon and Boppana (1987) strengthened the combinatorial arguments of Razborov and proved an exponential lower bound for the monotone circuit complexity of the clique function. In this section we describe a special case of this bound by showing that there are no linear size monotone circuits that decide if a given graph contains a triangle. Although this result is much weaker than the ones stated above, it illustrates nicely all the probabilistic considerations in the more complicated proofs and avoids some of the combinatorial subtleties, whose detailed proofs can be found in the above mentioned papers.

Put $n = \binom{m}{2}$, and let $x_1, x_2, \dots, x_n$ be n Boolean variables representing the edges of a graph on the set of vertices $\{1, 2, \dots, m\}$. Let $T = T(x_1, \dots, x_n)$ be the monotone Boolean function whose value is 1 if the corresponding graph contains a triangle. Clearly there is a binary monotone circuit of size $O(m^3)$ computing T . Thus the following theorem is tight, up to a polylogarithmic factor.

Theorem 12.4.1 The monotone circuit complexity of T is at least $\Omega(m^3 / \log^4 m)$.

Before we present the proof of this theorem we introduce some notation and prove a simple lemma. For any Boolean function $f = f(x_1, \dots, x_n)$ define

$$A(f) = \{(x_1, \dots, x_n) \in \{0, 1\}^n : f(x_1, \dots, x_n) = 1\}.$$

Clearly $A(f \vee g) = A(f) \cup A(g)$ and $A(f \wedge g) = A(f) \cap A(g)$. Let C be a monotone circuit of size s computing the function $f = f(x_1, \dots, x_n)$. Clearly C supplies a monotone straight-line program of length s computing f ; that is, a sequence of functions $x_1, x_2, \dots, x_n, f_1, \dots, f_s$, where $f_s = f$ and each f_i , for $1 \leq i \leq s$, is either an OR or an AND of two of the previous functions. By applying the operation A we obtain a sequence $A(C)$ of subsets of $\{0, 1\}^n$: $A_{-n} = A_{x_n}, \dots, A_{-1} = A_{x_1}, A_1, \dots, A_s$, where $A_{x_i} = A(x_i)$, $A_s = A(f)$ and each A_i for $1 \leq i \leq s$ is either

a union or an intersection of two of the previous subsets. Let us replace the sequence $A(C)$ by an approximating sequence $M(C) : M_{-n} = M_{x_n} = A_{x_n}, \dots, M_{-1} = M_{x_1} = A_{x_1}, M_1, \dots, M_s$ defined by replacing the union and intersection operations in $A(C)$ by the approximating operations $\sqcup$ and $\sqcap$, respectively. The exact definition of these two operations will be given later, in such a way that for all admissible M and L the inclusions

$$M \sqcup L \supseteq M \cup L \quad \text{and} \quad M \sqcap L \subseteq M \cap L \quad (12.9)$$

will hold. Thus $M_{x_i} = A_{x_i}$ for all $1 \leq i \leq n$ and if for some $1 \leq j \leq s$ we have $A_j = A_\ell \cup A_k$ then $M_j = M_\ell \sqcup M_k$, whereas if $A_j = A_\ell \cap A_k$ then $M_j = M_\ell \sqcap M_k$. In the former case put $\delta_\sqcup^j = M_j - (M_\ell \cup M_k)$ and $\delta_\sqcap^j = \emptyset$, and in the latter case put $\delta_\sqcap^j = (M_\ell \cap M_k) - M_j$ and $\delta_\sqcup^j = \emptyset$.

Lemma 12.4.2 For all members M_i of $M(C)$,

$$A_i \setminus \bigcup_{j \leq i} \delta_\sqcap^j \subseteq M_i \subseteq A_i \cup \bigcup_{j \leq i} \delta_\sqcup^j. \quad (12.10)$$

Proof. We apply induction on i . For $i < 0$, $M_i = A_i$ and thus (12.10) holds. Assuming (12.10) holds for all M_j with $j < i$ we prove it for i . If $A_i = A_\ell \cup A_k$, then, by the induction hypothesis,

$$M_i = M_\ell \cup M_k \cup \delta_\sqcup^i \subseteq A_\ell \cup A_k \cup \bigcup_{j \leq i} \delta_\sqcup^j = A_i \cup \bigcup_{j \leq i} \delta_\sqcup^j$$

and

$$\begin{aligned} M_i &= M_\ell \sqcup M_k \supseteq M_\ell \cup M_k \\ &\supseteq \left(A_\ell \setminus \bigcup_{j \leq \ell} \delta_\sqcap^j \right) \cup \left(A_k \setminus \bigcup_{j \leq k} \delta_\sqcap^j \right) \supseteq A_i \setminus \bigcup_{j \leq i} \delta_\sqcap^j \end{aligned}$$

as needed. If $A_i = A_\ell \cap A_k$ the proof is similar.

Lemma 12.4.2 holds for any choice of the operations $\sqcup$ and $\sqcap$ that satisfies (12.9). In order to prove Theorem 12.4.1 we define these operations as follows. Put $r = 100 \log^2 m$. For any set R of at most r edges on $V = \{1, 2, \dots, m\}$, let $\lceil R \rceil$ denote the set of all graphs on V containing at least one edge of R . In particular, $\lceil \emptyset \rceil$ is the empty set. We also let $\lceil * \rceil$ denote the set of all graphs. The elements of $M(C)$ will all have the form $\lceil R \rceil$ or $\lceil * \rceil$. Note that $A_{x_i} = M_{x_i}$ is simply the set $\lceil R \rceil$, where R is a singleton containing the appropriate single edge. For two sets R_1 and R_2 of at most r edges each, we define $\lceil R_1 \rceil \sqcap \lceil R_2 \rceil = \lceil R_1 \cap R_2 \rceil$, $\lceil R_1 \rceil \sqcap \lceil * \rceil = \lceil R_1 \rceil$ and $\lceil * \rceil \sqcap \lceil * \rceil = \lceil * \rceil$. Similarly, if $|R_1 \cup R_2| \leq r$ we define $\lceil R_1 \rceil \sqcup \lceil R_2 \rceil = \lceil R_1 \cup R_2 \rceil$, whereas if $|R_1 \cup R_2| > r$ then $\lceil R_1 \rceil \sqcup \lceil R_2 \rceil = \lceil * \rceil$. Finally $\lceil * \rceil \sqcup \lceil R_1 \rceil = \lceil * \rceil \sqcup \lceil * \rceil = \lceil * \rceil$.

Proof [Theorem 12.4.1]. We now prove Theorem 12.4.1 by showing that there is no monotone circuit of size $s < \binom{m}{3}/2r^2$ computing the function T . Indeed, suppose this is false and let C be such a circuit. Let $M(C) = M_{x_n}, \dots, M_{x_1}, M_1, \dots, M_s$ be an approximating sequence of length s obtained from C as described above. By Lemma 12.4.2,

$$A(T) \setminus \bigcup_{j \leq s} \delta_\sqcap^j \subseteq M_s \subseteq A(T) \cup \bigcup_{j \leq s} \delta_\sqcup^j \quad (12.11)$$

We consider two possible cases.

1. $M_s = \lceil R \rceil$, where $|R| \leq r$.

Let us choose a random triangle Δ on $\{1, 2, \dots, m\}$. Clearly

$$\Pr [\Delta \in M_s] \leq \frac{r(m-2)}{\binom{m}{3}} < \frac{1}{2}$$

Moreover, for each fixed $j \leq s$,

$$\Pr [\Delta \in \delta_{\square}^j] \leq \frac{r^2}{\binom{m}{3}}$$

This is because if $\delta_{\square}^j \neq \emptyset$, then $\delta_{\square}^j = ([R_1] \cap [R_2]) \setminus [R_1 \cap R_2]$ for some two sets of edges R_1, R_2 , each of cardinality at most r . The only triangles in this difference are those containing an edge from R_1 and another edge from R_2 (and no edge of both). Since there are at most r^2 such triangles the last inequality follows. Since $s < \binom{m}{3}/2r^2$ the last two inequalities imply that $\Pr [\Delta \notin M_s \text{ and } \Delta \notin \bigcup_{j \leq s} \delta_{\square}^j] > 0$ and thus there is such a triangle Δ . Since this triangle belongs to $A(T)$ this contradicts (12.11), showing that this case is impossible.

2. $M_s = [*]$.

Let B be a random spanning complete bipartite graph on $V = \{1, 2, \dots, m\}$ obtained by coloring each vertex in V randomly and independently by 0 or 1 and taking all edges connecting vertices with distinct colors. Since M_s is the set of all graphs, $B \in M_s$. Also $B \notin A(T)$, as it contains no triangle. We claim that for every fixed $j, j \leq s$,

$$\Pr [B \in \delta_{\square}^j] \leq 2^{-\sqrt{r}/2} < \frac{1}{m^5}. \quad (12.12)$$

Indeed, if $\delta_{\square}^j \neq \emptyset$, then $\delta_{\square}^j = [*] \setminus ([R_1] \cup [R_2])$, where $|R_1 \cup R_2| > r$. Consider the graph whose set of edges is $R_1 \cup R_2$. Let d be its maximum degree. By Vizing's Theorem the set of its edges can be partitioned into at most $d + 1$ matchings. Thus either $d > \sqrt{r}/2$ or the size of the maximum matching in this graph is at least $\sqrt{r}/2$. It follows that our graph contains a set of $k = \sqrt{r}/2$ edges $e_1, \dots, e_k$ that form either a star or a matching. In each of these two cases $\Pr [e_i \in B] = \frac{1}{2}$ and these events are mutually independent. Hence

$$\Pr [B \notin [R_1] \cup [R_2]] \leq 2^{-\sqrt{r}/2},$$

implying (12.12). Note that a similar estimate can be established without Vizing's Theorem by observing that B does not belong to $[R_1] \cup [R_2]$ if and only if the vertices in any connected component of the graph whose edges are $R_1 \cup R_2$ belong to the same color class of B .

Since $s < \binom{m}{3}/2r^2 < m^5$, inequality (12.12) implies that there is a bipartite B such that $B \in M_s$, $B \notin A(T)$ and $B \notin \bigcup_{j \leq s} \delta_{\square}^j$. This contradicts (12.11), shows that this case is impossible as well and hence completes the proof of Theorem 12.4.1. $\square$

18.3.5 12.5 Formulae

Recall that a formula is a circuit in which every fanout is at most 1. Unlike in the case of circuits, there are known superlinear lower bounds for the minimum size of formulae computing various explicit NP-functions over the full binary basis. For a Boolean function $f = f(x_1, \dots, x_n)$, let us denote by $L(f)$ the minimum number of And and Or gates in a formula that uses And, Or and Not gates and computes f . By De Morgan's rules we may assume that all Not gates appear in the first level of this formula. We conclude this chapter with a simple result of Subbotovskaya (1961), which implies that for the parity function $f = x_1 \oplus \dots \oplus x_n$, $L(f) \geq \Omega(n^{3/2})$. This bound was improved later by Khrapchenko (1971) to $L(f) = n^2 - 1$. However, we present here only the weaker $\Omega(n^{3/2})$ lower bound, not only because it demonstrates, once more, the power of relatively simple probabilistic arguments, but also because a modification of this proof

enabled Andreev (1987) to obtain an $\Omega(n^{5/2}/(\log n)^{O(1)})$ lower bound for $L(g)$ for another NP-function $g = g(x_1, \dots, x_n)$. Håstad (1998) later improved this lower bound to $\Omega(n^{3-o(1)})$. This is at present the largest known lower bound for the formula complexity of an NP-function of n variables over a complete basis.

The method of Subbotovskaya is based on random restrictions similar to the ones used in Section 12.2. The main lemma is the following.

Lemma 12.5.1 Let $f = f(x_1, \dots, x_n)$ be a nonatom Boolean function of n variables. Then there is an i , $1 \leq i \leq n$, and an $\varepsilon \in \{0, 1\}$ such that for the function $g = f(x_1, \dots, x_{i-1}, \varepsilon, x_{i+1}, \dots, x_n)$ of $n - 1$ variables obtained from f by substituting $x_i = \varepsilon$, the following inequality holds:

$$(L(g) + 1) \leq \left(1 - \frac{3}{2n}\right) (L(f) + 1) \leq \left(1 - \frac{1}{n}\right)^{3/2} (L(f) + 1).$$

Proof. Fix a formula F computing f with $l = L(f)$ And and Or gates. F can be represented by a binary tree each of whose $l+1$ leaves is labeled by an atom x_i or $\bar{x}_i$. Let us choose, randomly, a variable x_i , $1 \leq i \leq n$, according to a uniform distribution, and assign to it a random binary value $\varepsilon \in \{0, 1\}$. When we substitute the values ε and $1 - \varepsilon$ to x_i and $\bar{x}_i$, respectively, the number of leaves in F is reduced; the expected number of leaves omitted in this manner is $(l+1)/n$. However, further reduction may occur.

Indeed, suppose a leaf is labeled x_i and it feeds, say, an And gate $x_i \wedge H$ in F . Observe that we may assume that the variable x_i does not appear in the subformula H , as otherwise F can be simplified by substituting $x_i = 1$ in H . If $x_i = \varepsilon = 0$, then H can be deleted once we substitute the value for x_i , thus further decreasing the number of leaves.

Since the behavior of this effect is similar for an Or gate (and also for $\bar{x}_i$ instead of x_i), it follows that the expected number of additional leaves omitted is at least $(l+1)/2n$. Hence the expected number of remaining leaves in the simplified formula is at most $(l+1)(1 - 3/2n)$, as claimed.

By repeatedly applying Lemma 12.5.1 we obtain the following.

Corollary 12.5.2 If $f = f(x_1, \dots, x_n)$ and $L(f) \leq (n/k)^{3/2} - 1$, then one can assign values to $n - k$ variables so that the resulting function g is an atom.

Proof. Repeated application of Lemma 12.5.1 $n - k$ times yields a g with

$$(L(g) + 1) \leq \prod_{i=k+1}^n \left(1 - \frac{1}{i}\right)^{3/2} (L(f) + 1) = (k/n)^{3/2} (L(f) + 1) \leq 1$$

Hence g is either x_i or $\bar{x}_i$ for some i .

Corollary 12.5.3 For the parity function $f = x_1 \oplus \dots \oplus x_n$,

$$L(f) > \left(\frac{n}{2}\right)^{3/2} - 1$$

18.3.6 12.6 Exercises

1. Show that there exists a constant c such that the number of binary Boolean circuits of size s on n variables is at most $(c(s+n))^s$.
2. Let f be a Boolean formula in the n variables $x_1, x_2, \dots, x_n$, where f is an AND of an arbitrary (finite) number of clauses; each clause is an Or of 10 literals, where each literal is either a variable or its negation, and suppose each variable appears (negated or unnegated) in at most 10 clauses. Prove that f is satisfiable.

3. (*) Prove that there is a bounded-depth, polynomial-size, monotone circuit of n Boolean inputs $x_1, x_2, \dots, x_n$ computing a function f whose value is 1 if $\sum_{i=1}^n x_i \geq n/2 + n/\log n$ and is 0 if $\sum_{i=1}^n x_i \leq n/2 - n/\log n$.

The Probabilistic Lens: Maximal Antichains

A family $\mathcal{F}$ of subsets of $\{1, \dots, n\}$ is called an antichain if no set of $\mathcal{F}$ is contained in another.

Theorem 1 Let $\mathcal{F}$ be an antichain. Then

$$\sum_{A \in \mathcal{F}} \frac{1}{\binom{n}{|A|}} \leq 1.$$

Proof. Let σ be a uniformly chosen permutation of $\{1, \dots, n\}$ and set

$$\mathcal{C}_\sigma = \{\{\sigma(j) : 1 \leq j \leq i\} : 0 \leq i \leq n\}$$

(The cases $i = 0$ and $i = n$ give $\emptyset \in \mathcal{C}_\sigma$ and $\{1, \dots, n\} \in \mathcal{C}_\sigma$, respectively.) Define a random variable

$$X = |\mathcal{F} \cap \mathcal{C}_\sigma|$$

We decompose

$$X = \sum_{A \in \mathcal{F}} X_A,$$

where X_A is the indicator random variable for $A \in \mathcal{C}$. Then

$$\mathbb{E}[X_A] = \Pr[A \in \mathcal{C}_\sigma] = \frac{1}{\binom{n}{|A|}}$$

since $\mathcal{C}_\sigma$ contains precisely one set of size $|A|$, which is distributed uniformly among the $|A|$ -sets. By linearity of expectation,

$$\mathbb{E}[X] = \sum_{A \in \mathcal{F}} \frac{1}{\binom{n}{|A|}}$$

For any σ , $\mathcal{C}_\sigma$ forms a chain - every pair of sets is comparable. Since $\mathcal{F}$ is an antichain we must have $X = |\mathcal{F} \cap \mathcal{C}_\sigma| \leq 1$. Thus $\mathbb{E}[X] \leq 1$.

Corollary 2 [Sperner's Theorem] Let $\mathcal{F}$ be an antichain. Then

$$|\mathcal{F}| \leq \binom{n}{\lfloor n/2 \rfloor}.$$

Proof. The function $\binom{n}{x}$ is maximized at $x = \lfloor n/2 \rfloor$ so that

$$1 \geq \sum_{A \in \mathcal{F}} \frac{1}{\binom{n}{|A|}} \geq \frac{|\mathcal{F}|}{\binom{n}{\lfloor n/2 \rfloor}}$$

18.4 13 Discrepancy

The mystery, as well as the glory, of mathematics lies not so much in the fact that abstract theories do turn out to be useful in solving problems but in that wonder of wonders, the fact that a theory meant for solving one type of problem is often the only way of solving problems of entirely different kinds, problems for which the

theory was not intended. These coincidences occur so frequently that they must belong to the essence of mathematics.

\author{

- Gian-Carlo Rota
- }

18.4.1 13.1 Basics

Suppose we are given a finite family of finite sets. Our object is to color the underlying points red and blue so that all of the sets have nearly the same number of red and blue points. It may be that our cause is hopeless - if the family consists of all subsets of a given set Ω then regardless of the coloring some set, either the red or the blue points, will have size at least half that of Ω and be monochromatic. In the other extreme, should the sets of the family be disjoint then it is trivial to color so that all sets have the same number of red and blue points or, at worst if the cardinality is odd, the number of red and blue points differ by only one. The discrepancy will measure how good a coloring we may find.

To be formal, let a family $\mathcal{A}$ of subsets of Ω be given. Rather than using red and blue we consider colorings as maps

$$\chi : \Omega \rightarrow \{-1, +1\}.$$

For any $A \subset \Omega$ we set

$$\chi(A) = \sum_{a \in A} \chi(a).$$

Define the discrepancy of $\mathcal{A}$ with respect to χ by

$$\text{disc}(\mathcal{A}, \chi) = \max_{A \in \mathcal{A}} |\chi(A)|$$

and the discrepancy of $\mathcal{A}$ by

$$\text{disc}(\mathcal{A}) = \min_{\chi: \Omega \rightarrow \{-1, +1\}} \text{disc}(\mathcal{A}, \chi)$$

Other equivalent definitions of discrepancy reveal its geometric aspects. Let $\mathcal{A} = \{S_1, \dots, S_m\}$, $\Omega = \{1, \dots, n\}$ and let $B = [b_{ij}]$ be the $m \times n$ incidence matrix: $b_{ij} = 1$ if $j \in S_i$, otherwise $b_{ij} = 0$. A coloring χ may be associated with the vector $u = (\chi(1), \dots, \chi(n)) \in \{-1, +1\}^n$ so that $Bu^T = (\chi(S_1), \dots, \chi(S_m))$ and

$$\text{disc}(\mathcal{A}) = \min_{u \in \{-1, +1\}^n} \|Bu^T\|_{\infty},$$

where $\|v\|_{\infty}$ is the L^{∞} -norm, the maximal absolute value of the coordinates. Similarly, letting v_j denote the j th column vector of B (the profile of point j)

$$\text{disc}(\mathcal{A}) = \min |\pm v_1 \pm \dots \pm v_n|_{\infty},$$

where the minimum ranges over all 2^n choices of sign.

We will generally be concerned with upper bounds to the discrepancy. Unraveling the definitions, $\text{disc}(\mathcal{A}) \leq K$ if and only if there exists a coloring χ for which $|\chi(A)| \leq K$ for all $A \in \mathcal{A}$. Naturally, we try the random coloring.

Theorem 13.1.1 Let $\mathcal{A}$ be a family of n subsets of an m -set Ω . Then

$$\text{disc}(\mathcal{A}) \leq \sqrt{2m \ln(2n)}$$

Proof. Let $\chi : \Omega \rightarrow \{-1, +1\}$ be random. For $A \subset \Omega$ let X_A be the indicator random variable for $|\chi(A)| > \alpha$, where we set $\alpha = \sqrt{2m \ln(2n)}$. If $|A| = a$ then $\chi(A)$ has distribution S_a , so by Theorem A.1.1

$$\mathbb{E}[X_A] = \Pr[|\chi(A)| > \alpha] < 2e^{-\alpha^2/2a} \leq 2e^{-\alpha^2/2m} = 1/n$$

by our propitious choice of α . Let X be the number of $A \in \mathcal{A}$ with $|\chi(A)| > \alpha$ so that

$$X = \sum_{A \in \mathcal{A}} X_A$$

and linearity of expectation gives

$$\mathbb{E}[X] = \sum_{A \in \mathcal{A}} \mathbb{E}[X_A] < |\mathcal{A}|(1/n) = 1$$

Thus for some χ we must have $X = 0$. This means $\text{disc}(\mathcal{A}, \chi) \leq \alpha$ and therefore $\text{disc}(\mathcal{A}) \leq \alpha$.

18.4.2 13.2 SIX Standard Deviations Suffice

When $\mathcal{A}$ has both n sets and n points Theorem 13.1.1 gives

$$\text{disc}(\mathcal{A}) = O(\sqrt{n \ln(n)})$$

This is improved by the following result. Its proof resembles that of the main result of Beck (1981). The approach via entropy was suggested by R. Boppana.

Theorem 13.2.1 [Spencer (1985b)] Let $\mathcal{A}$ be a family of n subsets of an n -element set Ω . Then

$$\text{disc}(\mathcal{A}) < 6\sqrt{n}$$

With $\chi : \Omega \rightarrow \{-1, +1\}$ random, $A \in \mathcal{A}$, $\chi(A)$ has zero mean and variance at most $\sqrt{n}$. If $|\chi(A)| > 6\sqrt{n}$ then $\chi(A)$ is at least six standard deviations off the mean. The probability of this occurring is very small but a fixed positive constant and the number of sets $A \in \mathcal{A}$ is going to infinity. In fact, a random χ almost always will not work. The specific constant 6 (actually 5.32) was the result of specific calculations that could certainly be further improved and will not concern us here. Rather, we show Theorem 13.2.1 with the constant 11.

A map $\chi : \Omega \rightarrow \{-1, 0, +1\}$ will be called a partial coloring. When $\chi(a) = 0$ we say a is uncolored. We define $\chi(A)$ as before.

Lemma 13.2.2 Let $\mathcal{A}$ be a family of n subsets of an n -set Ω . Then there is a partial coloring χ with at most $10^{-9}n$ points uncolored such that

$$|\chi(A)| \leq 10\sqrt{n}$$

for all $A \in \mathcal{A}$.

Here the values 10 and 10^{-9} are not the best possible. The significant point is that they are absolute constants. Label the sets of $\mathcal{A}$ by $A_1, \dots, A_n$ for convenience. Let

$$\chi : \Omega \rightarrow \{-1, +1\}$$

be random. For $1 \leq i \leq n$ define

$$b_i = \text{nearest integer to } \frac{\chi(A_i)}{20\sqrt{n}}$$

For example, $b_i = 0$ when $-10\sqrt{n} < \chi(A_i) < 10\sqrt{n}$ and $b_i = -3$ when $-70\sqrt{n} < \chi(A_i) < -50\sqrt{n}$. From Theorem A.1.1 (as in Theorem 13.1.1),

$$\begin{aligned}\Pr[b_i = 0] &> 1 - 2e^{-50} \\ \Pr[b_i = 1] &= \Pr[b_i = -1] < e^{-50} \\ \Pr[b_i = 2] &= \Pr[b_i = -2] < e^{-450}\end{aligned}$$

and, in general,

$$\Pr[b_i = s] = \Pr[b_i = -s] < e^{-50(2s-1)^2}$$

Now we bound the entropy $H(b_i)$. This important concept is explored more fully in Section 15.7. Letting $p_j = \Pr[b_i = j]$,

$$\begin{aligned}H(b_i) &= \sum_{j=-\infty}^{+\infty} -p_j \log_2(p_j) \\ &\leq (1 - 2e^{-50}) [-\log_2(1 - 2e^{-50})] + 2e^{-50} [-\log_2 e^{-50}] \\ &\quad + 2e^{-450} [-\log_2 e^{-450}] + \dots\end{aligned}$$

The infinite sum clearly converges and is strongly dominated by the second term. Calculation gives

$$H(b_i) \leq \epsilon = 3 \times 10^{-20}.$$

Now consider the n -tuple $(b_1, \dots, b_n)$. Of course, there may be correlation among the b_i . Indeed, if S_i and S_j are nearly equal then b_i and b_j will usually be equal. But by Proposition 15.7.2 entropy is subadditive. Hence

$$H((b_1, \dots, b_n)) \leq \sum_{i=1}^n H(b_i) \leq \epsilon n$$

If a random variable Z assumes no value with probability greater than 2^{-t} then $H(Z) \geq t$. In contrapositive form, there is a particular n -tuple $(s_1, \dots, s_n)$ so that

$$\Pr[(b_1, \dots, b_n) = (s_1, \dots, s_n)] \geq 2^{-\epsilon n}$$

Our probability space was composed of the 2^n possible colorings χ , all equally likely. Thus, shifting to counting, there is a set $\mathcal{C}'$ consisting of at least $2^{(1-\epsilon)n}$ colorings $\chi : \Omega \rightarrow \{-1, +1\}$, all having the same value $(b_1, \dots, b_n)$.

Let us think of the class $\mathcal{C}$ of all colorings $\chi : \Omega \rightarrow \{-1, +1\}$ as the Hamming cube $\{-1, +1\}^n$ endowed with the Hamming metric

$$\rho(\chi, \chi') = |\{a : \chi(a) \neq \chi'(a)\}|.$$

Kleitman (1966b) has proved that if $\mathcal{D} \subset \mathcal{C}$ and

$$|\mathcal{D}| \geq \sum_{i \leq r} \binom{n}{i}$$

with $r \leq n/2$ then $\mathcal{D}$ has diameter at least $2r$. That is, the set of a given size with minimal diameter is the ball. ($\mathcal{D}$ has diameter at least r trivially, which would suffice to prove Lemma

13.2.2 and Theorem 13.2.1 with weaker values for the constants.)

Proof. In our case we may take $r = \alpha n$ as long as $\alpha < \frac{1}{2}$ and

$$2^{H(\alpha)} \leq 2^{1-\epsilon}$$

Calculation gives that we may take $\alpha = \frac{1}{2}(1 - 10^{-9})$ with room to spare. [The Taylor series expansion give

$$H\left(\frac{1}{2} - x\right) \sim 1 - \frac{2}{\ln 2}x^2$$

for x small.] Thus $\mathcal{C}'$ has diameter at least $n(1 - 10^{-9})$. Let $\chi_1, \chi_2 \in \mathcal{C}'$ be at maximal distance. We set

$$\chi = \frac{\chi_1 - \chi_2}{2}$$

χ is a partial coloring of Ω . $\chi(a) = 0$ if and only if $\chi_1(a) = \chi_2(a)$, which occurs for $n - \rho(\chi_1, \chi_2) \leq 10^{-9}n$ coordinates a . Finally, and crucially, for each $1 \leq i \leq n$ the colorings χ_1, χ_2 yield the same value b_i , which means that $\chi_1(A_i)$ and $\chi_2(A_i)$ lie on a common interval of length $20\sqrt{n}$. Thus

$$|\chi(A_i)| = \left| \frac{\chi_1(A_i) - \chi_2(A_i)}{2} \right| \leq 10\sqrt{n},$$

as desired.

Theorem 13.2.1 requires a coloring of all points whereas Lemma 13.2.2 leaves $10^{-9}n$ points uncolored. The idea now is to iterate the procedure of Lemma 13.2.2, coloring all but, say, $10^{-18}n$ of the uncolored points on the second coloration. We cannot apply Lemma 13.2.2 directly since we have an asymmetric situation with n sets and only $10^{-9}n$ points.

Lemma 13.2.3 Let $\mathcal{A}$ be a family of n subsets of an r -set Ω with $r \leq 10^{-9}n$. Then there is a partial coloring χ of Ω with at most $10^{-40}r$ points uncolored so that

$$|\chi(A)| < 10\sqrt{r}\sqrt{\ln(n/r)}$$

for all $A \in \mathcal{A}$.

Proof. We outline the argument which leaves room to spare. Let $A_1, \dots, A_n$ denote the sets of $\mathcal{A}$. Let $\chi : \Omega \rightarrow \{\pm 1\}$ be random. For $1 \leq i \leq n$ define

$$b_i = \text{nearest integer to } \frac{\chi(A_i)}{20\sqrt{r}\sqrt{\ln(n/r)}}$$

Now the probability that $b_i = 1$ is less than $(r/n)^{50}$. The entropy $H(b_i)$ is dominated by this term and is less than

$$3\left(\frac{r}{n}\right)^{50} \left[-\log_2 \left(\left(\frac{r}{n}\right)^{50} \right) \right] < 10^{-100} \frac{r}{n}$$

The entropy of $(b_1, \dots, b_n)$ is then less than $10^{-100}r$; one finds nearly antipodal χ_1, χ_2 with the same b 's and takes $\chi = (\chi_1 - \chi_2)/2$ as before.

Proof [Theorem 13.2.1]. Apply Lemma 13.2.2 to find a partial coloring χ^1 and then apply Lemma 13.2.3 repeatedly on the remaining uncolored points giving $\chi^2, \chi^3, \dots$ until all points have been colored. Let χ denote the final coloring. For any $A \in \mathcal{A}$,

$$\chi(A) = \chi^1(A) + \chi^2(A) + \dots$$

so that

$$|\chi(A)| \leq 10\sqrt{n} + 10\sqrt{10^{-9}n}\sqrt{\ln 10^9} \\ + 10\sqrt{10^{-49}n}\sqrt{\ln 10^{49}} + 10\sqrt{10^{-89}n}\sqrt{\ln 10^{89}} + \dots$$

Removing the common $\sqrt{n}$ term gives a clearly convergent infinite series, strongly dominated by the first term so that

$$|\chi(A)| \leq 11\sqrt{n}$$

with room to spare.

Suppose that $\mathcal{A}$ consists of n sets on r points and $r < n$. We can apply Lemma 13.2.3 repeatedly (first applying Lemma 13.2.2 if $r > 10^{-9}n$) to give a coloring χ with

$$\text{disc}(\mathcal{A}, \chi) < K\sqrt{r}\sqrt{\ln(n/r)}$$

where K is an absolute constant. As long as $r = n^{1-o(1)}$ this improves the random coloring result of Theorem 13.1.1.

18.4.3 13.3 Linear and Hereditary Discrepancy

We now suppose that $\mathcal{A}$ has more points than sets. We write $\mathcal{A} = \{A_1, \dots, A_n\}$ and $\Omega = \{1, \dots, m\}$ and assume $m > n$. Note that $\text{disc}(\mathcal{A}) \leq K$ is equivalent to the existence of a set S ; namely, $S = \{j : \chi(j) = +1\}$, with $|S \cap A|$ within $K/2$ of $|A|/2$ for all $A \in \mathcal{A}$. We define the linear discrepancy $\text{lindisc}(\mathcal{A})$ by

$$\text{lindisc}(\mathcal{A}) = \max_{p_1, \dots, p_m \in [0,1]} \min_{\epsilon_1, \dots, \epsilon_m \in \{0,1\}} \max_{A \in \mathcal{A}} \left| \sum_{i \in A} (\epsilon_i - p_i) \right|.$$

The upper bound $\text{lindisc}(\mathcal{A}) \leq K$ means that given any $p_1, \dots, p_m$ there is a "simultaneous roundoff" $\epsilon_1, \dots, \epsilon_m$ so that, with $S = \{j : \epsilon_j = 1\}$, $|S \cap A|$ is within K of the weighted sum $\sum_{j \in A} p_j$ for all $A \in \mathcal{A}$. Taking all $p_j = \frac{1}{2}$, the upper bound implies $\text{disc}(\mathcal{A}) \leq 2K$. But $\text{lindisc}(\mathcal{A}) \leq K$ is much stronger. It implies, taking all $p_j = \frac{1}{3}$, the existence of an S with all $|S \cap A|$ within K of $|A|/3$, and much more. Linear discrepancy and its companion hereditary discrepancy defined below have been developed in Lovász, Spencer and Vesztergombi (1986). For $X \subset \Omega$ let $\mathcal{A}|_X$ denote the restriction of $\mathcal{A}$ to X , that is, the family $\{A \cap X : A \in \mathcal{A}\}$. The next result "reduces" the bounding of $\text{disc}(\mathcal{A})$ when there are more points than sets to the bounding of $\text{lindisc}(\mathcal{A})$ when the points do not outnumber the sets.

Theorem 13.3.1 Let $\mathcal{A}$ be a family of n sets on m points with $m \geq n$. Suppose that $\text{lindisc}(\mathcal{A}|_X) \leq K$ for every subset X of at most n points. Then $\text{lindisc}(\mathcal{A}) \leq K$.

Proof. Let $p_1, \dots, p_m \in [0, 1]$ be given. We define a reduction process. Call index j fixed if $p_j \in \{0, 1\}$, otherwise call it floating, and let F denote the set of floating indices. If $|F| \leq n$ then halt. Otherwise, let $y_j, j \in F$, be a nonzero solution to the homogeneous system

$$\sum_{j \in A \cap F} y_j = 0, \quad A \in \mathcal{A}.$$

Such a solution exists since there are more variables ($|F|$) than equations (n) and may be found by standard techniques of linear algebra. Now set

$$p'_j = \begin{cases} p_j + \lambda y_j, & j \in F, \\ p_j, & j \notin F, \end{cases}$$

where we let λ be the real number of least absolute value so that for some $j \in F$ the value p'_j becomes zero or one. Critically,

$$\sum_{j \in A} p'_j = \sum_{j \in A} p_j + \lambda \sum_{j \in A \cap F} y_j = \sum_{j \in A} p_j \quad (13.1)$$

for all $A \in \mathcal{A}$. Now iterate this process with the new p'_j . At each iteration at least one floating j becomes fixed and so the process eventually halts at some $p_1^*, \dots, p_m^*$. Let X be the set of floating j at this point. Then $|X| \leq n$. By assumption there exist $\epsilon_j, j \in X$ so that

$$\left| \sum_{j \in A \cap X} p_j^* - \epsilon_j \right| \leq K, \quad A \in \mathcal{A}.$$

For $j \notin X$ set $\epsilon_j = p_j^*$. As (13.1) holds at each iteration,

$$\sum_{j \in A} p_j^* = \sum_{j \in A} p_j$$

and hence

$$\left| \sum_{j \in A} (p_j - \epsilon_j) \right| = \left| \sum_{j \in A} (p_j - p_j^*) + \sum_{j \in A \cap X} (p_j^* - \epsilon_j) \right| \leq K$$

for all $A \in \mathcal{A}$.

We now define the hereditary discrepancy $\text{herdisc}(\mathcal{A})$ by

$$\text{herdisc}(\mathcal{A}) = \max_{X \subseteq \Omega} \text{disc}(\mathcal{A}|_X)$$

Example. Let $\Omega = \{1, \dots, n\}$ and let $\mathcal{A}$ consist of all intervals $[i, j] = \{i, i+1, \dots, j\}$ with $1 \leq i \leq j \leq n$. Then $\text{disc}(\mathcal{A}) = 1$ as we may color Ω alternately $+1$ and -1 . But also $\text{herdisc}(\mathcal{A}) = 1$. For given any $X \subseteq \Omega$, say, with elements $x_1 < x_2 < \dots < x_r$, we may color X alternately by $\chi(x_k) = (-1)^k$. For any set $[i, j] \in \mathcal{A}$ the elements of $[i, j] \cap X$ are alternately colored.

Theorem 13.3.2 $\text{lindisc}(\mathcal{A}) \leq \text{herdisc}(\mathcal{A})$.

Proof. Set $K = \text{herdisc}(\mathcal{A})$. Let $\mathcal{A}$ be defined on $\Omega = \{1, \dots, m\}$ and let $p_1, \dots, p_m \in [0, 1]$ be given. First let us assume that all p_i have finite expansions when written in base two. Let T be the minimal integer so that all $p_i 2^T \in \mathbb{Z}$. Let J be the set of i for which p_i has a one in the T th digit of its binary expansion, that is, for which $p_i 2^{T-1} \notin \mathbb{Z}$. As $\text{disc}(\mathcal{A}|_J) \leq K$ there exist $\epsilon_j \in \{-1, +1\}$, so that

$$\left| \sum_{j \in J \cap A} \epsilon_j \right| \leq K$$

for all $A \in \mathcal{A}$. Write $p_j = p_j^{(T)}$. Now set

$$p_j^{(T-1)} = \begin{cases} p_j^{(T)} & \text{if } j \notin J \\ p_j^{(T)} + \epsilon_j 2^{-T} & \text{if } j \in J \end{cases}$$

That is, the $p_j^{(T-1)}$ are the "roundoffs" of the $p_j^{(T)}$ in the T th place. Note that all $p_j^{(T-1)} 2^{-(T-1)} \in \mathbb{Z}$. For any $A \in \mathcal{A}$,

$$\left| \sum_{j \in A} p_j^{(T-1)} - p_j^{(T)} \right| = \left| \sum_{j \in J \cap A} 2^{-T} \epsilon_j \right| \leq 2^{-T} K$$

Iterate this procedure, finding $p_j^{(T-2)}, \dots, p_j^{(1)}, p_j^{(0)}$. All $p_j^{(0)} 2^{-0} \in \mathbb{Z}$ so all $p_j^{(0)} \in \{0, 1\}$ and

$$\left| \sum_{j \in A} p_j^{(0)} - p_j^{(T)} \right| \leq \sum_{i=1}^T \left| \sum_{j \in A} p_j^{(i-1)} - p_j^{(i)} \right| \leq \sum_{i=1}^T 2^{-i} K \leq K,$$

as desired.

What about general $p_1, \dots, p_m \in [0, 1]$? We can flip and say that, at least to a computer scientist, all real numbers have finite binary expansions. More rigorously, the function

$$f(p_1, \dots, p_m) = \min_{\epsilon_1, \dots, \epsilon_m \in \{0, 1\}} \max_{A \in \mathcal{A}} \left| \sum_{i \in A} (\epsilon_i - p_i) \right|$$

is the finite minimum of finite maxima of continuous functions and thus is continuous. The set of $(p_1, \dots, p_m) \in [0, 1]^m$ with all $p_i 2^T \in \mathbb{Z}$ for some T is a dense subset of $[0, 1]^m$. As $f \leq K$ on this dense set, $f \leq K$ for all $(p_1, \dots, p_m) \in [0, 1]^m$.

Corollary 13.3.3 Let $\mathcal{A}$ be a family of n sets on m points. Suppose $\text{disc}(\mathcal{A}|_X) \leq K$ for every subset X with at most n points. Then $\text{disc}(\mathcal{A}) \leq 2K$.

Proof. For every $X \subseteq \Omega$ with $|X| \leq n$, $\text{herdisc}(\mathcal{A}|_X) \leq K$ so by Theorem 13.3.2 $\text{lindisc}(\mathcal{A}|_X) \leq K$. By Theorem 13.3.1 $\text{lindisc}(\mathcal{A}) \leq K$. But

$$\text{disc}(\mathcal{A}) \leq 2 \text{lindisc}(\mathcal{A}) \leq 2K.$$

Corollary 13.3.4 For any family $\mathcal{A}$ of n sets of arbitrary size

$$\text{disc}(\mathcal{A}) \leq 12\sqrt{n}$$

Proof. Apply Theorem 13.2.1 and Corollary 13.3.3.

18.4.4 13.4 Lower Bounds

We now give two quite different proofs that, up to a constant factor, Corollary 13.3.4 is the best possible. A Hadamard matrix is a square matrix $H = (h_{ij})$ with all $h_{ij} \in \{-1, +1\}$ and with row vectors mutually orthogonal (and hence with column vectors mutually orthogonal). Let H be a Hadamard matrix of order n and let $v = (v_1, \dots, v_n)$, $v_i \in \{-1, +1\}$. Then

$$Hv = v_1 c_1 + \dots + v_n c_n$$

where c_i denotes the i th column vector of H . Writing $Hv = (L_1, \dots, L_n)$ and letting $|\cdot|$ denote the usual Euclidean norm,

$$L_1^2 + \dots + L_n^2 = |Hv|^2 = v_1^2 |c_1|^2 + \dots + v_n^2 |c_n|^2 = n + \dots + n = n^2$$

since the c_i are mutually orthogonal. Hence some $L_i^2 \geq n$ and thus

$$|Hv|_\infty = \max\{|L_1|, \dots, |L_n|\} \geq \sqrt{n}$$

Now we transfer this result to one on families of sets. Let H be a Hadamard matrix of order n with first row and first column all ones. (Any Hadamard matrix can be so "normalized" by multiplying appropriate rows and columns by -1.) Let J denote the all ones matrix of order n . Let $v_1, \dots, L_1, \dots$ be as above. Then

$$L_1 + \dots + L_n = \sum_{i,j=1}^n v_j h_{ij} = \sum_{j=1}^n v_j \sum_{i=1}^n h_{ij} = nv_1 = \pm n,$$

since the first column sums to n but the other columns, being orthogonal to it, sum to zero. Set $\lambda = v_1 + \dots + v_n$ so that $Jv = (\lambda, \dots, \lambda)$ and

$$(H + J)v = (L_1 + \lambda, \dots, L_n + \lambda).$$

We calculate

$$|(H + J)v|^2 = \sum_{i=1}^n (L_i + \lambda)^2 = \sum_{i=1}^n (L_i^2 + 2\lambda L_i + \lambda^2) = n^2 \pm 2n\lambda + n\lambda^2.$$

Assume n is even. (Hadamard matrices don't exist for odd n , except $n = 1$.) Then λ is an even integer. The quadratic (in λ) $n^2 \pm 2n\lambda + n\lambda^2$ has a minimum at ∓ 1 and so under the restriction of being an even integer its minimum is at $\lambda = 0, \mp 2$ and so

$$|(H + J)v|^2 \geq n^2$$

Again, some coordinate must be at least $\sqrt{n}$. Setting $H^* = (H + J)/2$,

$$|H^*v|_\infty \geq \sqrt{n}/2$$

Let $\mathcal{A} = \{A_1, \dots, A_m\}$ be any family of subsets of $\Omega = \{1, \dots, n\}$ and let M denote the corresponding $m \times n$ incidence matrix. A coloring $\chi : \Omega \rightarrow \{-1, +1\}$ corresponds to a vector $v = (\chi(1), \dots, \chi(n)) \in \{-1, +1\}^n$. Then

$$\text{disc}(\mathcal{A}, \chi) = |Mv|_\infty$$

and

$$\text{disc}(\mathcal{A}) = \min_{v \in \{-1, +1\}^n} |Mv|_\infty$$

In our case H^* has entries 0,1. Thus we have the following theorem.

Theorem 13.4.1 If a Hadamard matrix exists of order $n > 1$ then there exists a family $\mathcal{A}$ consisting of n subsets of an n -set with

$$\text{disc}(\mathcal{A}) \geq \sqrt{n}/2$$

While it is not known precisely for which n a Hadamard matrix exists [the Hadamard conjecture is that they exist for $n = 1, 2$ and all multiples of 4; see, for example, Hall (1986)], it is known that the orders of Hadamard matrices are dense in the sense that for all ϵ if n is sufficiently large there will exist a Hadamard matrix of order between n and $n(1 - \epsilon)$. This result suffices to extend Theorem 13.4.1 to an asymptotic result on all n .

Our second argument for the existence of $\mathcal{A}$ with high discrepancy involves turning the probabilistic argument "on its head." Let M be a random 0,1 matrix of order n . Let $v = (v_1, \dots, v_n)$, $v_j = \pm 1$ be fixed and set $Mv = (L_1, \dots, L_n)$. Suppose half of the $v_j = +1$ and half are -1. Then

$$L_i \sim B\left(\frac{n}{2}, \frac{1}{2}\right) - B\left(\frac{n}{2}, \frac{1}{2}\right)$$

which has roughly the normal distribution $N(0, \sqrt{n}/2)$. Pick $\lambda > 0$ so that

$$\int_{-\lambda}^{\lambda} \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt < \frac{1}{2}$$

Then

$$\Pr [|L_i| < \lambda\sqrt{n}/2] < \frac{1}{2}$$

When v is imbalanced the same inequality holds; we omit the details. Now, crucially, the L_i are mutually independent as each entry of M was independently chosen. Thus

$$\Pr [|L_i| < \lambda\sqrt{n}/2 \text{ for all } 1 \leq i \leq n] < \left(\frac{1}{2}\right)^n$$

There are "only" 2^n possible v . Thus the expected number of v for which $|Mv|_\infty < \lambda\sqrt{n}/2$ is less than $2^n 2^{-n} = 1$. For some M this value must be zero, there are no such v . The corresponding family $\mathcal{A}$ thus has

$$\text{disc}(\mathcal{A}) > \lambda\sqrt{n}/2.$$

18.4.5 13.5 The Beck-fiala Theorem

For any family $\mathcal{A}$ let $\deg(\mathcal{A})$ denote the maximal number of sets containing any particular point. The following result due to Beck and Fiala (1981) uses only methods from linear algebra and thus is technically outside the scope we have set for this book. We include it both for the sheer beauty of the proof and because the result itself is very much in the spirit of this chapter.

Theorem 13.5.1 Let $\mathcal{A}$ be a finite family of finite sets, no restriction on either the number of sets or on the cardinality of the sets, with $\deg(\mathcal{A}) \leq t$. Then

$$\text{disc}(\mathcal{A}) \leq 2t - 1$$

Proof. For convenience write $\mathcal{A} = \{A_1, \dots, A_m\}$ with all $A_i \subseteq \Omega = \{1, \dots, n\}$. To each $j \in \Omega$ there is assigned a value x_j that will change as the proof progresses. Initially all $x_j = 0$. At the end all $x_j = \pm 1$. We will have $-1 \leq x_j \leq +1$ at all times and once $x_j = \pm 1$ it "sticks" there and that becomes its final value. A set S_i has value $\sum_{j \in S_i} x_j$. At any time j is called fixed if $x_j = \pm 1$; otherwise it is floating. A set S_i is safe if it has fewer than t floating points; otherwise it is active. Note, crucially, that as points are in at most t sets and active sets contain more than t floating points there must be fewer active sets than floating points.

We insist at all times that all active sets have value zero. This holds initially since all sets have value zero. Suppose this condition holds at some stage. Consider x_j a variable for each floating j and a constant for each fixed j . The condition that S_i has value zero then becomes a linear equation in these variables. This is an underdetermined system: there are fewer linear conditions (active sets) than variables (floating points). Hence we may find a line, parametrized

$$x'_j = x_j + \lambda y_j, j \text{ floating,}$$

on which the active sets retain value zero. Let λ be the smallest value for which some x'_j becomes ± 1 and replace each x_j by x'_j . (Geometrically, follow the line until reaching the boundary of the cube in the space over the floating variables.) This process has left fixed variables fixed and so safe sets stayed safe sets (though active sets may have become safe) and so the condition still holds. In addition, at least one previously floating j has become fixed.

We iterate the above procedure until all j have become fixed. (Toward the end we may have no active sets at which time we may simply set the floating x_j to ± 1 arbitrarily.) Now consider any set S_i . Initially it had value zero and it retained value zero while it contained at least t floating points. Consider the time when it first becomes safe, say, $1, \dots, l$ were its floating points. At this moment its value is zero. The variables $y_1, \dots, y_l$ can now change less than two to their final value since all values are in $[-1, +1]$. Thus, in total, they may change less than $2t$. Hence the final value of S_i is less than $2t$ and, as it is an integer, it is at most $2t - 1$.

Conjecture 13.5.2 If $\deg(\mathcal{A}) \leq t$ then $\text{disc}(\mathcal{A}) \leq K\sqrt{t}$, K an absolute constant.

This conjecture seems to call for a melding of probabilistic methods and linear algebra. The constructions of t sets on t points, described in Section 13.4, show that, if true, this conjecture would be the best possible.

18.4.6 13.6 Exercises

1. Let $\mathcal{A}$ be a family of n subsets of $\Omega = \{1, \dots, m\}$ with m even. Let $\chi(i)$, $1 \leq i \leq m/2$ be independent and uniform in $\{-1, +1\}$ and set $\chi(i+m/2) = -\chi(i)$ for $1 \leq i \leq m/2$. Using this notion of random coloring, improve Theorem 13.1.1 by showing $\text{disc}(\mathcal{A}) \leq \sqrt{m \ln(2n)}$. Show that this can be improved even further by splitting Ω randomly into $m/2$ disjoint pairs.
2. Let $\vec{v}_1, \dots, \vec{v}_s \in \mathbb{R}^n$. Let $x_1, \dots, x_s \in [-1, +1]$ such that $\sum_{i=1}^s x_i \vec{v}_i = \vec{0}$ and such that $x_i \in \{-1, +1\}$ for all but at most n values of i . Let $\vec{v}_{s+1} \in \mathbb{R}^n$. Use the linear ideas of Section 13.5 to find $x'_1, \dots, x'_s, x'_{s+1}$ with the following properties:
 - $\sum_{i=1}^{s+1} x'_i \vec{v}_i = \vec{0}$.
 - All $x'_i \in [-1, +1]$.
 - $x'_i \in \{-1, +1\}$ for all but at most n values of i .
 - $x'_i = x_i$ whenever $x_i \in \{-1, +1\}$.

Use the above to prove the following result of Barány and Grinberg: Let $|\cdot|$ be an arbitrary norm in $\mathbb{R}^n$. Let $\vec{v}_1, \dots, \vec{v}_s \in \mathbb{R}^n$ with all $|\vec{v}_i| \leq 1$. Then there exist $x_1, \dots, x_s \in \{-1, +1\}$ such that

$$\left| \sum_{i=1}^t x_i \vec{v}_i \right| \leq 2n$$

for all $1 \leq t \leq s$.

3. Let $A_1, \dots, A_n \subset \Omega = \{1, \dots, m\}$ with $m \sim n \ln n$. Assume further that all $|A_i| \leq n$. Use the methods of Theorem 13.2.1, including Kleitman's Theorem, to prove that there exists $\chi : \{1, \dots, m\} \rightarrow \{-1, 0, +1\}$ such that all $\chi(A_i) = O(\sqrt{n \ln \ln n})$ and $\chi(x) = 0$ for at most n vertices x . Use Theorem 13.2.1 to deduce the existence of $\chi : \{1, \dots, m\} \rightarrow \{-1, +1\}$ such that all $\chi(A_i) = O(\sqrt{n \ln \ln n})$.

The Probabilistic Lens: Unbalancing Lights

For any $m \times n$ matrix $B = (b_{ij})$ with coefficients $b_{ij} = \pm 1$ set

$$F[B] = \max_{x_i, y_j = \pm 1} \sum_{i=1}^m \sum_{j=1}^n x_i y_j b_{ij}.$$

As in Section 2.5 we may interpret B as an $m \times n$ array of lights, each either on ($b_{ij} = +1$) or off ($b_{ij} = -1$). For each row and each column there is a switch that, when pulled, changes all lights in that line from on to off or from off to on. Then $F[B]$ gives the maximal achievable number of lights on minus lights off. In Section 2.5 we found a lower bound for $F[B]$ when $m = n$. Here we set $n = 2^m$ and find the precise best possible lower bound.

With $n = 2^m$ let A be an $m \times n$ matrix with coefficients ± 1 containing every possible column vector precisely once. We claim $F[A]$ is the minimal value of $F[B]$ over all $m \times n$ matrices B .

For any given B let $x_1, \dots, x_m = \pm 1$ be independently and uniformly chosen and set

$$X_j = \sum_{i=1}^m x_i b_{ij},$$

$$X = |X_1| + \dots + |X_n|,$$

so that

$$F[B] = \max_{y_j = \pm 1} \max_{x_i = \pm 1} \sum_{j=1}^n y_j X_j = \max_{x_i = \pm 1} \sum_{j=1}^n |X_j| = \max X.$$

Regardless of the b_{ij} , X_i has distribution S_m so that $E[|X_i|] = E[|S_m|]$ and, by linearity of expectation,

$$E[X] = nE[|S_m|].$$

With $B = A$, any choices of $x_1, \dots, x_m = \pm 1$ have the effect of permuting the columns - the matrix $(x_i a_{ij})$ also has every column vector precisely once - so that $X = |X_1| + \dots + |X_m|$ is a constant. Note that $E[X]$ is independent of B . In general, fixing $E[X] = \mu$, the minimal possible value for $\max X$ is achieved when X is the constant μ . Thus $F[B]$ is minimized with $B = A$.

18.5 14 Geometry

Few people think more than two or three times a year. I have made an international reputation for myself by thinking once or twice a week.
-George Bernard Shaw

Suppose we choose randomly n points $P_1, \dots, P_n$ on the unit circle, according to a uniform distribution. What is the probability that the origin lies in the convex hull of these points? There is a surprisingly simple (yet clever) way to compute this probability. Let us first choose n random pairs of antipodal points $Q_1, Q_{n+1} = -Q_1, Q_2, Q_{n+2} = -Q_2, \dots, Q_n, Q_{2n} = -Q_n$ according to a uniform distribution. Note that with probability 1 these pairs are all distinct. Next we choose each P_i to be either Q_i or its antipodal $Q_{n+i} = -Q_i$, where each choice is equally likely. Clearly this corresponds to a random choice of the points P_i . The probability that the origin does not belong to the convex hull of the points P_i , given the (distinct) points Q_j , is precisely $x/2^n$, where x is the number of subsets of the points Q_j contained in an open half-plane determined by a line through the origin, which does not pass through any of the points Q_j . It is easy to see that $x = 2n$. This is because if we renumber the points Q_j so that their cyclic order on the circle is $Q_1, \dots, Q_n, Q_{n+1}, \dots, Q_{2n}$ and $Q_{n+i} = -Q_i$ then the subsets contained in such half-planes are precisely $\{Q_i, \dots, Q_{n+i-1}\}$, where the indices are reduced modulo $2n$. Therefore the probability that the origin is in the convex hull of n randomly chosen points on the unit circle is precisely $1 - 2n/2^n$. Observe that the same result holds if we replace the unit circle by any centrally symmetric bounded planar domain with center 0 and that the argument can easily be generalized to higher dimensions.

This result, due to Wendel (1962), shows how in some cases a clever idea can replace a tedious computation. It also demonstrates the connection between probability and geometry. The probabilistic method has recently been used extensively for deriving results in discrete and computational geometry. Some of these results are described in this chapter.

18.5.1 14.1 The Greatest Angle Among Points in Euclidean Spaces

There are several striking examples, in different areas of combinatorics, where the probabilistic method supplies very simple counter examples to long standing conjectures. Here is an example, due to Erdős and Füredi (1983).

Theorem 14.1.1 For every $d \geq 1$ there is a set of at least $\lfloor \frac{1}{2}(2/\sqrt{3})^d \rfloor$ points in the d -dimensional Euclidean space $\mathbb{R}^d$, such that all angles determined by three points from the set are strictly less than $\pi/2$.

This theorem disproves an old conjecture of Danzer and Grünbaum (1962) that the maximum cardinality of such a set is at most $2d - 1$. We note that as proved by Danzer and Grünbaum the maximum cardinality of a set of points in $\mathbb{R}^d$ in which all angles are at most $\pi/2$ is 2^d .

Proof. We select the points of a set $X \subset \mathbb{R}^d$ from the vertices of the d -dimensional cube. As usual, we view the vertices of the cube, which are 0,1 -vectors of length d , as the characteristic vectors of subsets of a d -element set; that is, each 0,1 -vector a of length d is associated with the set $A = \{i : 1 \leq i \leq d, a_i = 1\}$. A simple consequence of Pythagoras' Theorem gives that the three vertices a, b and c of the d -cube, corresponding to the sets A, B and C , respectively, determine a right angle at c if and only if

$$A \cap B \subset C \subset A \cup B \quad (14.1)$$

As the angles determined by triples of points of the d -cube are always at most $\pi/2$, it suffices to construct a set X of cardinality at least the one stated in the theorem, no three distinct members of which satisfy (14.1).

Define $m = \lfloor \frac{1}{2}(2/\sqrt{3})^d \rfloor$, and choose, randomly and independently, $2md$ dimensional $\{0, 1\}$ -vectors $a_1, \dots, a_{2m}$, where each coordinate of each of the vectors independently is chosen to be either 0 or 1 with equal probability. For every fixed triple a, b and c of the chosen points, the probability that the corresponding sets satisfy equation (14.1) is precisely $(3/4)^d$. This is because (14.1) simply means that for each $i, 1 \leq i \leq d$, neither $a_i = b_i = 0, c_i = 1$ nor $a_i = b_i = 1, c_i = 0$ hold. Therefore the probability that, for three fixed indices i, j and k , our chosen points a_i, a_j, a_k form a right angle at a_k is $(3/4)^d$. Since there are $\binom{2m}{3}3$ possible triples that can produce such angles, the expected number of right angles is

$$\binom{2m}{3}3(3/4)^d \leq m,$$

where the last inequality follows from the choice of m . Thus there is a choice of a set X of $2m$ points in which the number of right angles is at most m . By deleting one point from each such angle we obtain a set of at least $2m - m = m$ points in which all angles are strictly less than $\pi/2$. Note that the remaining points are all distinct since (14.1) is trivially satisfied if $A = C$. This completes the proof.

It is worth noting that, as observed by Erdős and Füredi, the proof above can be easily modified to give the following.

Theorem 14.1.2 For every $\epsilon > 0$ there is a $\delta > 0$ such that for every $d \geq 1$ there is a set of at least $(1 + \delta)^d$ points in $\mathbb{R}^d$ so that all the angles determined by three distinct points from the set are at most $\pi/3 + \epsilon$.

We omit the detailed proof of this result.

18.5.2 14.2 Empty Triangles Determined by Points in the Plane

For a finite set X of points in general position in the plane, let $f(X)$ denote the number of empty triangles determined by triples of points of X ; that is, the number of triangles, determined by points of X , that contain no other point of X . Katchalski and Meir (1988) studied the minimum possible value of $f(X)$ for a set X of n points. Define $f(n) = \min f(X)$, where X ranges over all planar sets of n points in general position (i.e., containing no three collinear points). They proved that

$$\binom{n-1}{2} \leq f(n) < 200n^2.$$

These bounds were improved by Barány and Füredi (1987), who showed that as n grows

$$(1 + o(1))n^2 \leq f(n) \leq (1 + o(1))2n^2$$

The construction that establishes the upper bound is probabilistic and is given in the following theorem. See also Valtr (1995) for a slightly better result.

Theorem 14.2.1 Let $I_1, I_2, \dots, I_n$ be parallel unit intervals in the plane, where

$$I_i = \{(x, y) : x = i, 0 \leq y \leq 1\}.$$

For each i let us choose a point p_i randomly and independently from I_i according to a uniform distribution. Let X be the set consisting of these n randomly chosen points. Then the expected number of empty triangles in X is at most $2n^2 + O(n \log n)$.

Clearly, with probability 1, X is a set of points in general position and hence the above theorem shows that $f(n) \leq 2n^2 + O(n \log n)$.

Proof. We first estimate the probability that the triangle determined by the points p_i, p_{i+a} and p_{i+k} is empty, for some fixed i, a and $k = a + b \geq 3$. Let $A = (i, x)$, $B = (i + a, y)$ and $C = (i + k, z)$ be the points p_i, p_{i+a} and p_{i+k} , respectively. Let m be the distance between B and the intersection point of the segment AC with the interval I_{i+a} . Since each of the points p_j for $i < j < i + k$ are chosen randomly according to a uniform distribution on I_j , it follows that the probability that the triangle determined by A, B and C is empty is precisely

$$\begin{aligned} & \left(1 - \frac{m}{a}\right) \left(1 - 2\frac{m}{a}\right) \cdots \left(1 - (a-1)\frac{m}{a}\right) \left(1 - (b-1)\frac{m}{b}\right) \cdots \left(1 - \frac{m}{b}\right) \\ & \leq \exp\left(-\frac{m}{a} - 2\frac{m}{a} \cdots - (a-1)\frac{m}{a} - (b-1)\frac{m}{b} \cdots - \frac{m}{b}\right) \\ & = \exp\left(-\binom{a}{2}\frac{m}{a} - \binom{b}{2}\frac{m}{b}\right) = \exp\left(-\binom{k-2}{2}\frac{m}{2}\right) \end{aligned}$$

For every fixed choice of A and C , when the point $p_{i+a} = B$ is chosen randomly, the probability that its distance m from the intersection of the segment AC with the interval I_{i+a} is at most d is clearly at most $2d$, for all $d \geq 0$. Therefore the probability that the triangle determined by p_i, p_{i+a} and p_{i+k} is empty is at most

$$2 \int_0^\infty e^{-\frac{1}{2}(k-2)m} dm = \frac{4}{k-2}$$

It follows that the expected value of the total number of empty triangles is at most

$$\begin{aligned}
 & n - 2 + \sum_{i=1}^{n-3} \sum_{k=3}^{n-i} \sum_{a=1}^{k-1} \frac{4}{k-2} \\
 &= n - 2 + \sum_{k=3}^{n-1} (n-k)(k-1) \frac{4}{k-2} \\
 &= n - 2 + \sum_{k=3}^{n-1} (n-k) \frac{4}{k-2} + 4 \sum_{k=3}^{n-1} (n-k) \\
 &= 2n^2 + O(n \log n).
 \end{aligned}$$

This completes the proof.

The result above can be extended to higher dimensions by applying a similar probabilistic construction. A set X of n points in the d -dimensional Euclidean space is called independent if no $d+1$ of the points lie on a hyperplane. A simplex determined by $d+1$ of the points is called empty if it contains no other point of X . Let $f_d(X)$ denote the number of empty simplices of X , and define $f_d(n) = \min f_d(X)$, where X ranges over all independent sets of n points in $\mathbb{R}^d$. Katchalski and Meir (1988) showed that $f_d(n) \geq \binom{n-1}{d}$. The following theorem of Bárány and Füredi

shows that here, again, a probabilistic construction gives a matching upper bound, up to a constant factor (that depends on the dimension). We omit the detailed proof.

Theorem 14.2.2 There exists a constant $K = K(d)$, such that for every convex, bounded set $A \subset \mathbb{R}^d$ with nonempty interior, if X is a random set of n points obtained by n random and independent choices of points of A picked with uniform distribution, then the expected number of empty simplices of X is at most $K \binom{n}{d}$.

18.5.3 14.3 Geometrical Realizations of Sign Matrices

Let $A = (a_{i,j})$ be an m by n matrix with ± 1 entries. We say that A is realizable in $\mathbb{R}^d$ if there are m hyperplanes $H_1, \dots, H_m$ in $\mathbb{R}^d$ passing through the origin and n points $P_1, \dots, P_n$ in $\mathbb{R}^d$, so that for all i and j , P_j lies on the positive side of H_i if $a_{i,j} = +1$, and on the negative side if $a_{i,j} = -1$. Let $d(A)$ denote the minimum dimension d such that A is realizable in $\mathbb{R}^d$, and define $d(m, n) = \max d(A)$, where A ranges over all m by n matrices with ± 1 entries. Since $d(m, n) = d(n, m)$ we can consider only the case $m \geq n$.

The problem of determining or estimating $d(m, n)$ and, in particular, $d(n, n)$ was raised by Paturi and Simon (1984). This problem was motivated by an attempt to estimate the maximum possible "unbounded-error probabilistic communication complexity" of Boolean functions. Alon, Frankl and Rödl (1985) proved that as n grows $\frac{1}{32}n \leq d(n, n) \leq (\frac{1}{2} + o(1))n$. Both the upper and the lower bounds are proved by combining probabilistic arguments with certain other ideas. In the next theorem we prove the upper bound, which is probably closer to the truth.

Theorem 14.3.1 For all $m \geq n$,

$$d(m, n) \leq \frac{n+1}{2} + \sqrt{\frac{n-1}{2} \log m}.$$

For the proof, we need a definition and two lemmas. For a vector $\vec{a} = (a_1, \dots, a_n)$ of ± 1 entries, the number of sign changes in $\vec{a}$ is the number of indices $i, 1 \leq i \leq n-1$ such that $a_i = -a_{i+1}$. For a matrix A of ± 1 entries, denote by $s(A)$ the maximum number of sign changes in a row of A .

Lemma 14.3.2 For any matrix A of ± 1 entries, $d(A) \leq s(A) + 1$.

Proof. Let $A = (a_{i,j})$ be an m by n matrix of ± 1 entries and suppose $s = s(A)$. Let $t_1 < t_2 < \dots < t_n$ be arbitrary reals, and define n points $P_1, P_2, \dots, P_n$ in $\mathbb{R}^{s+1}$ by $P_j = (1, t_j, t_j^2, \dots, t_j^s)$.

These points, whose last s coordinates represent points on the d -dimensional moment curve, will be the points used in the realization of A . To complete the proof we have to show that each row of A can be realized by a suitable hyperplane through the origin. This is proved by applying some of the known properties of the moment curve as follows. Consider the sign vector representing an arbitrary row of A . Suppose this vector has r sign changes, where, of course,

$r \leq s$. Suppose the sign changes in this vector occur between the coordinates i_j and $i_j + 1$, for $1 \leq j \leq r$. Choose arbitrary reals $y_1, \dots, y_r$, where $t_{i_j} < y_j < t_{i_{j+1}}$ for $1 \leq j \leq r$. Consider the polynomial $P(t) = \prod_{j=1}^r (t - y_j)$. Since its degree is at most s there are real numbers a_j such that $P(t) = \sum_{j=0}^s a_j t^j$. Let H be the hyperplane in $\mathbb{R}^{s+1}$ defined by $H = \{(x_0, x_1, \dots, x_s) \in \mathbb{R}^{s+1} : \sum_{j=0}^s a_j x_j = 0\}$. Clearly the point $P_j = (1, t_j, \dots, t_j^s)$ is on the positive side of this hyperplane if $P(t_j) > 0$, and is on its negative side if $P(t_j) < 0$. Since the polynomial P changes sign only in the values y_j , it follows that the hyperplane H separates the points $P_1, \dots, P_n$ according to the sign pattern of the corresponding row of A . Hence, by choosing the orientation of H appropriately, we conclude that A is realizable in $\mathbb{R}^{s+1}$, completing the proof of the lemma.

Lemma 14.3.3 For every m by n matrix A of ± 1 entries there is a matrix B obtained from A by multiplying some of the columns of A by -1 , such that

$$s(B) \leq \frac{n-1}{2} + \sqrt{\frac{n-1}{2} \log m}.$$

Proof. For each column of A , randomly and independently, choose a number $\epsilon \in \{\pm 1\}$, where each of the two choices is equally likely, and multiply this column by ϵ . Let B be the random sign matrix obtained in this way. Consider an arbitrary fixed row of B . One can easily check that the random variable describing the number of sign changes in this row is a binomial random variable with parameters $n-1$ and $p = 1/2$. This is because no matter what the entries of A in this row are, the row of B is a totally random row of ± 1 entries. By the standard estimates for binomial distributions, described in Appendix A, the probability that this number is greater than $\frac{1}{2}(n-1) + \sqrt{\frac{1}{2}(n-1) \log m}$ is smaller than $1/m$. Therefore with positive probability the number of sign changes in each of the m rows is at most that large, completing the proof.

Proof [Theorem 14.3.1]. Let A be an arbitrary m by n matrix of ± 1 entries. By Lemma 14.3.3 there is a matrix B obtained from A by replacing some of its columns by their inverses, such that $s(B) \leq \frac{1}{2}(n-1) + \sqrt{\frac{1}{2}(n-1) \log m}$. Observe that $d(A) = d(B)$, since any realization of one of these matrices by points and hyperplanes through the origin gives a realization of the other one by replacing the points corresponding to the altered columns by their antipodal points. Therefore, by Lemma 14.3.2,

$$d(A) = d(B) \leq s(B) + 1 \leq \frac{n+1}{2} + \sqrt{\frac{n-1}{2} \log m}.$$

This completes the proof.

It is worth noting that by applying the (general) six standard deviations theorem stated at the end of Section 13.2, the estimate in Lemma 14.3.3 (and hence in Theorem 14.3.1) can be improved to $\frac{1}{2}n + O(\sqrt{n \log(m/n)})$. It can also be shown that if n and m grow so that m/n^2 tends to infinity and $(\log_2 m)/n$ tends to 0 then for almost all m by n matrices A of ± 1 entries $d(A) = (\frac{1}{2} + o(1))n$.

18.5.4 14.4 ϵ -nets and Vc-dimensions of Range Spaces

What is the minimum number $f = f(n, \epsilon)$ such that every set X of n points in the plane contains a subset S of at most f points such that every triangle containing at least ϵn points of X contains at least one point of S ? As we shall see in this section, there is an absolute constant c such that $f(n, \epsilon) \leq (c/\epsilon) \log(1/\epsilon)$ and this estimate holds for every n . This somewhat surprising result is a very special case of a general theorem of Vapnik and Chervonenkis (1971), which has been extended by Haussler and Welzl (1987), and which has many interesting applications in computational geometry and in statistics.

In order to describe this result we need a few definitions. A range space S is a pair (X, R) , where X is a (finite or infinite) set and R is a (finite or infinite) family of subsets of X . The members of X are called points and those of R are called ranges. If A is a subset of X then $P_R(A) = \{r \cap A : r \in R\}$ is the projection of R on A . In case this projection contains all subsets of A we say that A is shattered. The Vapnik-Chervonenkis dimension (or VC-dimension) of S , denoted by $VC(S)$, is the maximum cardinality of a shattered subset of X . If there are arbitrarily large shattered subsets then $VC(S) = \infty$.

The number of ranges in any finite range space with a given number of points and a given VC-dimension cannot be too large. For integers $n \geq 0$ and $d \geq 0$, define a function $g(d, n)$ by $g(d, n) = \sum_{i=0}^d \binom{n}{i}$. Observe that for all $n, d \geq 1$, $g(d, n) = g(d, n-1) + g(d-1, n-1)$. The following combinatorial lemma was proved, independently, by Sauer (1972), by Perles and Shelah and, in a slightly weaker form, by Vapnik and Chervonenkis.

Lemma 14.4.1 If (X, R) is a range space of VC-dimension d with $|X| = n$ points then $|R| \leq g(d, n)$.

Proof. We apply induction on $n + d$. The assertion is trivially true for $d = 0$ and $n = 0$. Assuming it holds for n and $d-1$ and for $n-1$ and $d-1$ we prove it for n and d . Let $S = (X, R)$ be a range space of VC-dimension d on n points. Suppose $x \in X$, and consider the two range spaces $S - x$ and $S \setminus x$ defined as follows.

$$\begin{aligned} S - x &= (X - \{x\}, R - x), \text{ where } R - x = \{r \setminus \{x\} : r \in R\}, \\ S \setminus x &= (X - \{x\}, R \setminus x), \text{ where } R \setminus x = \{r \in R : x \notin r, r \cup \{x\} \in R\}. \end{aligned}$$

Clearly the VC-dimension of $S - x$ is at most d . It is also easy to see that the VC-dimension of $S \setminus x$ is at most $d-1$. Therefore, by the induction hypothesis,

$$|R| = |R - x| + |R \setminus x| \leq g(d, n-1) + g(d-1, n-1) = g(d, n),$$

completing the proof.

It is easy to check that the estimate given in the above lemma is sharp for all possible values of n and d . If (X, R) is a range space of VC-dimension d and $A \subset X$, then the VC-dimension of $(A, P_R(A))$ is clearly at most d . Therefore the last lemma implies the following.

Corollary 14.4.2 If (X, R) is a range space of VC-dimension d then for every finite subset A of X , $|P_R(A)| \leq g(d, |A|)$.

There are many range spaces with finite VC-dimension that arise naturally in discrete and computational geometry. One such example is the space $S = (\mathbb{R}^d, H)$, whose points are all the points in the d -dimensional Euclidean space, and whose set of ranges is the set of all (open) half-spaces. Any set of $d+1$ affinely independent points is shattered in this space, and, by Radon's Theorem, no set of $d+2$ points is shattered. Therefore $VC(S) = d+1$. As shown by Dudley (1978), if (X, R) has finite VC-dimension, so does (X, R_k) , where R_k is the set of all Boolean combinations formed from at most k ranges in R . In particular, the following statement is a simple consequence of Corollary 14.4.2.

Corollary 14.4.3 Let (X, R) be a range space of VC-dimension $d \geq 2$ and let (X, R_h) be the range space on X in which $R_h = \{(r_1 \cap \dots \cap r_h) : r_1, \dots, r_h \in R\}$. Then $VC(X, R_h) \leq$

$2dh \log(dh)$.

Proof. Let A be an arbitrary subset of cardinality n of X . By Corollary 14.4.2 $|P_R(A)| \leq g(d, n) \leq n^d$. Since each member of $P_{R_h}(A)$ is an intersection of h members of $P_R(A)$, it follows that $|P_{R_h}(A)| \leq \binom{g(d, n)}{h} \leq n^{dh}$. Therefore, if $n^{dh} < 2^n$, then A cannot be shattered. But this inequality holds for $n \geq 2dh \log(dh)$, since $dh \geq 4$.

As shown above, the range space whose set of points is $\mathbb{R}^d$ and whose set of ranges is the set of all half-spaces has VC-dimension $d + 1$. This and the last corollary imply that the range space $(\mathbb{R}^d, C_h)$, where C_h is the set of all convex d -polytopes with h facets, has a VC-dimension that does not exceed $2(d + 1)h \log((d + 1)h)$.

An interesting property of range spaces with a finite VC-dimension is the fact that each finite subset of such a set contains relatively small good samples in the sense described below. Let (X, R) be a range space and let A be a finite subset of X . For $0 \leq \epsilon \leq 1$, a subset $B \subset A$ is an ϵ -sample for A if for any range $r \in R$ the inequality

$$\left| \frac{|A \cap r|}{|A|} - \frac{|B \cap r|}{|B|} \right| \leq \epsilon$$

holds. Similarly, a subset $N \subset A$ is an ϵ -net for A if any range $r \in R$ satisfying $|r \cap A| > \epsilon|A|$ contains at least one point of N .

Note that every ϵ -sample for A is also an ϵ -net and that the converse is not true. However, both notions define subsets of A that represent approximately some of the behavior of A with respect to the ranges. Our objective is to show the existence of small ϵ -nets or ϵ -samples for finite sets in some range spaces. Observe that if (X, R) is a range space with an infinite VC-dimension then for every n there is a shattered subset A of X of cardinality n . It is obvious that any ϵ -net (and hence certainly any ϵ -sample) for such an A must contain at least $(1 - \epsilon)n$ points; that is, it must contain almost all points of A . Therefore in infinite VC-dimension there are no small nets or samples. However, it turns out that in finite VC-dimension there are always very small nets and samples. The following theorem was proved by Vapnik and Chervonenkis (1971).

Theorem 14.4.4 There is a positive constant c such that if (X, R) is any range space of VC-dimension at most d , $A \subset X$ is a finite subset and $\epsilon, \delta > 0$, then a random subset B of cardinality s of A , where s is at least the minimum between $|A|$ and

$$\frac{c}{\epsilon^2} \left(d \log \frac{d}{\epsilon} + \log \frac{1}{\delta} \right)$$

is an ϵ -sample for A with probability at least $1 - \delta$.

Using similar ideas, Haussler and Welzl (1987) proved the following theorem.

Theorem 14.4.5 Let (X, R) be a range space of VC-dimension d , let A be a finite subset of X and suppose $0 < \epsilon, \delta < 1$. Let N be a set obtained by m random independent draws from A , where

$$m \geq \max \left\{ \frac{4}{\epsilon} \log \frac{2}{\delta}, \frac{8d}{\epsilon} \log \frac{8d}{\epsilon} \right\}. \quad (14.2)$$

Then N is an ϵ -net for A with probability at least $1 - \delta$.

Therefore, if A is a finite subset of a range space of finite VC-dimension d , then for any $\epsilon > 0$, A contains ϵ -nets as well as ϵ -samples whose size is at most some function of ϵ and d , independent of the cardinality of A ! The result about the triangles mentioned in the first paragraph of this section thus follows from Theorem 14.4.5, together with the observation following Corollary 14.4.3 that implies that the range space whose ranges are all triangles in the plane has a finite VC-dimension. We note that, as shown by Pach and Woeginger (1990), there are cases in which, for fixed δ , the dependence of m in $1/\epsilon$ cannot be linear, but there is no known natural

geometric example demonstrating this phenomenon. See also Komlós, Pach and Woeginger (1992) for a tight form of the last theorem.

The proofs of Theorems 14.4.4 and 14.4.5 are very similar. Since the computation in the proof of Theorem 14.4.5 is simpler, we describe here only the proof of this theorem and encourage the reader to try and make the required modifications that yield a proof for Theorem 14.4.4.

Proof [Theorem 14.4.5]. Let (X, R) be a range space with VC-dimension d and let A be a subset of X of cardinality $|A| = n$. Suppose m satisfies (14.2), and let $N = (x_1, \dots, x_m)$ be obtained by m independent random choices of elements of A . (The elements in N are not necessarily distinct, of course.) Let E_1 be the following event:

$$E_1 = \{\exists r \in R : |r \cap A| > \epsilon n, r \cap N = \emptyset\}.$$

To complete the proof we must show that the probability of E_1 is at most δ . To this end, we make an additional random choice and define another event as follows. Independently of our previous choice, we let $T = (y_1, \dots, y_m)$ be obtained by m independent random choices of elements of A . Let E_2 be the event defined by

$$E_2 = \left\{ \exists r \in R : |r \cap A| > \epsilon n, r \cap N = \emptyset, |r \cap T| \geq \frac{\epsilon m}{2} \right\}.$$

[Since the elements of T are not necessarily distinct, the notation $|r \cap T|$ means here $|\{i : 1 \leq i \leq m, y_i \in r\}|$. The quantities $|r \cap N|$ and $|r \cap (N \cup T)|$ are similarly defined.]

Claim 14.4.6 $\Pr[E_2] \geq \frac{1}{2} \Pr[E_1]$.

Proof. It suffices to prove that the conditional probability $\Pr[E_2 \mid E_1]$ is at least $\frac{1}{2}$. Suppose that the event E_1 occurs. Then there is an $r \in R$ such that $|r \cap A| > \epsilon n$ and $r \cap N = \emptyset$. The conditional probability above is clearly at least the probability that for this specific r , $|r \cap T| \geq \frac{1}{2}\epsilon m$. However, $|r \cap T|$ is a binomial random variable with expectation pm and variance $(1-p)pm \leq pm$, where $p = |r \cap A|/|A| \geq \epsilon$. Hence, by Chebyshev's Inequality,

$$\Pr\left[|r \cap T| < \frac{\epsilon m}{2}\right] \leq \Pr\left[|r \cap T| < \frac{pm}{2}\right] \leq \frac{pm}{(pm/2)^2} \leq \frac{4}{\epsilon m} \leq \frac{1}{2},$$

where the last inequality follows from (14.2). Thus the assertion of Claim 14.4.6 is correct.

Claim 14.4.7 $\Pr[E_2] \leq g(d, 2m)2^{-\epsilon m/2}$.

Proof. The random choice of N and T can be described in the following way, which is equivalent to the previous one. First one chooses $N \cup T = (z_1, \dots, z_{2m})$ by making $2m$ random independent choices of elements of A , and then one chooses randomly precisely m of the elements z_i to be the set N (the remaining elements z_j form the set T , of course). For each range $r \in R$ satisfying $|r \cap A| > \epsilon n$, let E_r be the event that $|r \cap T| > \frac{1}{2}\epsilon m$ and $r \cap N = \emptyset$. A crucial fact is that if $r, r' \in R$ are two ranges, $|r \cap A| > \epsilon n$ and $|r' \cap A| > \epsilon n$ and if $r \cap (N \cup T) = r' \cap (N \cup T)$, then the two events E_r and $E_{r'}$, when both are conditioned on the choice of $N \cup T$, are identical. This is because the occurrence of E_r depends only on the intersection $r \cap (N \cup T)$. Therefore, for any fixed choice of $N \cup T$, the number of distinct events E_r does not exceed the number of different sets in the projection $P_{N \cup T}(R)$.

Since the VC-dimension of X is d , Corollary 14.4.2 implies that this number does not exceed $g(d, 2m)$.

Let us now estimate the probability of a fixed event of the form E_r given the choice of $N \cup T$. This probability is at most

$$\Pr\left[r \cap N = \emptyset \mid |r \cap (N \cup T)| > \frac{\epsilon m}{2}\right].$$

Define $s = |r \cap (N \cup T)|$. Since the choice of N among the elements of $N \cup T$ is independent of the choice of $N \cup T$, the last conditional probability is precisely

$$\begin{aligned} & \frac{(2m-s)(2m-s-1)\cdots(m-s+1)}{2m(2m-1)\cdots(m+1)} \\ &= \frac{m(m-1)\cdots(m-s+1)}{2m(2m-1)\cdots(2m-s+1)} \leq 2^{-s} \leq 2^{-\epsilon m/2} \end{aligned}$$

Since there are at most $g(d, 2m)$ potential distinct events E_r , it follows that the probability that at least one of them occurs given the choice of $N \cup T$ is at most $g(d, 2m)2^{-\epsilon m/2}$. Since this estimate holds conditioned on every possible choice of $N \cup T$ it follows that the probability of the event E_2 is at most $g(d, 2m)2^{-\epsilon m/2}$. This establishes Claim 14.4.7

By Claims 14.4.6 and 14.4.7, $\Pr[E_1] \leq 2g(d, 2m)2^{-\epsilon m/2}$. To complete the proof of the theorem it remains to show that if m satisfies inequality (14.2) then

$$2g(d, 2m)2^{-\epsilon m/2} \leq \delta$$

We describe the proof for $d \geq 2$. The computation for $d = 1$ is easier. Since $g(d, 2m) \leq (2m)^d$ it suffices to show that $2(2m)^d \leq \delta 2^{\epsilon m/2}$; that is,

$$\frac{\epsilon m}{2} \geq d \log(2m) + \log \frac{2}{\delta}$$

From (14.2) it follows that

$$\frac{\epsilon m}{4} \geq \log \frac{2}{\delta}$$

and hence it suffices to show that $\frac{1}{4}\epsilon m \geq d \log(2m)$.

The validity of the last inequality for some value of m implies its validity for any larger m and hence it suffices to check that it is satisfied for $m = (8d/\epsilon) \log(8d/\epsilon)$; that is,

$$2d \log \frac{8d}{\epsilon} \geq d \log \left(\frac{16d}{\epsilon} \log \frac{8d}{\epsilon} \right)$$

The last inequality is equivalent to $4d/\epsilon \geq \log(8d/\epsilon)$, which is certainly true. This completes the proof of the theorem.

Theorems 14.4.4 and 14.4.5 have been used for constructing efficient data structures for various problems in computational geometry. A trivial example is just the observation that Theorem 14.4.4 implies the following: For every $\epsilon > 0$ there is a constant $c = c(\epsilon)$ such that for every n and every set A of n points in the plane there is a data structure of size $c(\epsilon)$ that enables us to estimate, given any triangle in the plane, the number of points of A in this triangle up to an additive error of ϵn . This is done simply by storing the coordinates of a set of points that form an ϵ -sample for A considered as a subset of the range space whose ranges are all planar triangles. More sophisticated data structures whose construction relies on the above two theorems can be found in the paper of Haussler and Welzl (1987).

18.5.5 14.5 Dual Shatter Functions and Discrepancy

The dual shatter function h of a range space $S = (X, R)$ is the function h mapping integers to integers, defined by letting $h(g)$ denote the maximum, over all possible choices of g members of R , of the number of atoms in the Venn diagram of these members. It is not too difficult to prove that if the VC-dimension of S is d , then $h(g) \leq O\left(g^{2^{d+1}-1}\right)$, but in geometric applications it is usually better to bound this function directly.

In Matoušek, Welzl and Wernisch (1993) it is proved that if the dual shatter function of a range space $S = (X, R)$ satisfies $h(g) \leq O(g^t)$, A is any set of n points in the range space and $\mathcal{F}$ is the projection $P_R(A)$ of R on A , then the discrepancy of $\mathcal{F}$ satisfies

$$\text{disc}(\mathcal{F}) \leq O\left(n^{1/2-1/2t}\sqrt{\log n}\right) \quad (14.3)$$

This supplies nontrivial estimates in various geometric situations, improving the trivial bound that follows from Theorem 13.1.1 of Chapter 13. In most of these geometric applications it is widely believed that the $\sqrt{\log n}$ factor can be omitted. In the abstract setting, however, this factor cannot be omitted, as proved for $t = 2, 3$ in Matoušek (1997) and later for all t in Alon, Rónyai and Szabó (1999).

The proof of (14.3) is based on a beautiful result of Chazelle and Welzl (1989) and its improvement by Haussler (1995). It is somewhat simpler to prove the result with an extra logarithmic factor, and this is the proof we present here. See Pach and Agarwal (1995) for some additional information.

Let $\mathcal{F}$ be a family of subsets of a finite set A . In what follows we consider graphs whose edges are (unordered) pairs of points of A . For $F \in \mathcal{F}$ and $x, y \in A$, the edge xy stabs F if F contains exactly one of the two points x and y . The following theorem is proved in Chazelle and Welzl (1989). An improvement by a logarithmic factor appears in Haussler (1995).

Theorem 14.5.1 Let $(A, \mathcal{F})$ be a finite range space, where $|A| = n$, and suppose that its dual shatter function h satisfies $h(g) \leq cg^t$ for some fixed $c, t > 0$. Then, there is a $C = C(c, t)$ and a Hamiltonian path on A , such that each member F of $\mathcal{F}$ is stabbed by at most $Cn^{1-1/t} \log n$ edges of the path.

To prove the above theorem, we need the following lemma.

Lemma 14.5.2 Let $(A, \mathcal{F}), n, h, t$ and c be as above, let B be a finite subset of $p > 1$ points of A , and let $\mathcal{G}$ be a collection of m (not necessarily distinct) members of $\mathcal{F}$. Then there are two distinct points x, y in B , such that the edge xy stabs at most $(bm \log p)/p^{1/t}$ members of $\mathcal{G}$, where $b = b(c)$.

Proof. We may and will assume that p is larger than $c + 1$. Let g be the largest integer such that $cg^t \leq p - 1$, that is, $g = \lfloor ((p - 1)/c)^{1/t} \rfloor$. Let L be a random collection of g members of $\mathcal{G}$, each picked, randomly and independently (with possible repetitions), among all m members of $\mathcal{G}$ with uniform distribution. The Venn diagram of all members of L partition B into at most $h(g) \leq cg^t < p$ atoms and hence there are two distinct points x, y of B that lie in the same atom.

To complete the proof it suffices to show that with positive probability, for each pair of points of B that stabs more than $(bm \log p)/p^{1/t}$ members of $\mathcal{G}$, at least one of these members lies in L (and hence the pair does not lie in an atom of the corresponding Venn diagram.) There are $\binom{p}{2}$ such pairs and for each of them the probability that L contains no member of $\mathcal{G}$ it stabs is at most

$$\left(1 - \frac{b \log p}{p^{1/t}}\right)^g \leq \exp\left(-\frac{b \log p}{p^{1/t}} \left\lfloor \left(\frac{p-1}{c}\right)^{1/t} \right\rfloor\right),$$

which is less than $1/p^2$ for an appropriately chosen constant $b = b(c)$. This completes the proof.

Proof [Theorem 14.5.1]. Note first that if d is the VC-dimension of the given space, then there is a shattered set D of size d . It is not difficult to see that there are $g = \lceil \log_2 d \rceil$ sets among those shattering D , so that no two points of D lie in the same atom of their Venn diagram. Therefore $d \leq c(\lceil \log_2 d \rceil)^t$, implying that $d \leq 2^{c't \log t}$, where $c' = c'(c)$. By Lemma 14.4.1 this implies that the total number of ranges in R is at most $n^{2^{c't \log t}}$.

We next prove that there is a spanning tree of A satisfying the assertion of Theorem 14.5.1, and then show how to replace it by a Hamiltonian path. By Lemma 14.5.2 with $B_0 = A, p_0 = n$ and $\mathcal{G}_0 = \mathcal{F}, m_0 = |\mathcal{G}_0| \left(\leq n^{2^{c't \log t}}\right)$, we conclude that there is a pair x_0, y_0 of points in A

such that the edge x_0y_0 does not stab more than $m_0(b \log n)/n^{1/t}$ members of $\mathcal{G}$. Let $\mathcal{G}_1$ be the collection obtained from $\mathcal{G}$ by duplicating all members of $\mathcal{G}$ that are stabbed by x_0y_0 , and define $B_1 = B - x_0, p_1 = n - 1, m_1 = |\mathcal{G}_1| \leq m_0 [1 + (b \log n)/n^{1/t}]$. Applying Lemma 14.5.2 again, this time to B_1 and $\mathcal{G}_1$, we obtain another pair x_1, y_1 , define $B_2 = B_1 - x_1, p_2 = p_1 - 1 = n - 2$, and let $\mathcal{G}_2$ be the collection obtained from $\mathcal{G}_1$ by duplicating all members of $\mathcal{G}_1$ stabbed by $x_1y_1, m_2 = |\mathcal{G}_2|$. By the assertion of the lemma, $m_2 \leq m_1 [1 + (b \log n)/(n - 1)^{1/t}]$. Proceeding in this manner we get a sequence $x_0y_0, x_1y_1, \dots, x_{n-1}y_{n-1}$ of edges of a graph on A , a sequence of subsets $B_0 = A, B_1, \dots, B_{n-1}$, where each B_i is obtained from the previous one by omitting the point x_{i-1} , and a sequence of collections $\mathcal{G}_0, \mathcal{G}_1, \dots, \mathcal{G}_{n-1}$, where

$$\begin{aligned} |\mathcal{G}_{n-1}| &\leq m_0 \prod_{i=0}^{n-1} \left(1 + \frac{b \log n}{(n-i)^{1/t}} \right) \\ &\leq n^{2^{c't \log t}} \exp \left(b \log n \sum_{i=0}^{n-1} (n-i)^{-1/t} \right) \leq 2^{b'n^{1-1/t} \log n}, \end{aligned}$$

for an appropriate $b' = b'(c, t)$.

Note, now, that the edges x_iy_i form a spanning tree on the set A . The crucial observation is the fact that if a member of $\mathcal{F}$ is stabbed by s of the edges, then it is being duplicated s times during the above process that generates $\mathcal{G}_{n-1}$, implying that $2^s \leq |\mathcal{G}_{n-1}|$ and hence that $s \leq b'n^{1-1/t} \log n$. It remains to replace the spanning tree by a Hamiltonian path. To do so, replace each edge of the tree by two parallel edges and take a Euler tour in the resulting graph (in which all degrees are even). This is a sequence $x_0, x_1, x_2, \dots, x_{2n-2} = x_0$ of points of A such that each adjacent pair of elements of the sequence is an edge of the tree, and each edge appears twice this way. The subsequence of the above one obtained by keeping only the first appearance of each point of A is a Hamiltonian path; it is easy to check that each member of $\mathcal{F}$ is stabbed by at most $2b'n^{1-1/t} \log n$ of its edges, completing the proof.

The following result is a simple consequence of Theorem 14.5.1. As mentioned above, its assertion can be improved by a factor of $\sqrt{\log n}$.

Theorem 14.5.3 Let $(A, \mathcal{F})$ be a finite range space, where $|A| = n$, and suppose that its dual shatter function h satisfies $h(g) \leq cg^t$ for some fixed $c, t > 0$. Then, there is a $C' = C'(c, t)$ such that the discrepancy of $\mathcal{F}$ satisfies

$$\text{disc}(\mathcal{F}) \leq C'n^{1/2-1/2t} \log n.$$

Proof. Without loss of generality, assume that the number of points of A is even (otherwise, simply omit a point). By Theorem 14.5.1 there is a Hamiltonian path $x_1x_2 \cdots x_n$ on these points such that each member of $\mathcal{F}$ is stabbed by at most $Cn^{1-1/t} \log n$ edges of the path. Let $f : A \rightarrow \{-1, 1\}$ be a random coloring of A , where for each $i, 1 \leq i \leq n/2$, randomly and independently, either $f(x_{2i-1}) = 1, f(x_{2i}) = -1$ or $f(x_{2i-1}) = -1, f(x_{2i}) = 1$, the two choices being equally likely. Fix a member $F \in \mathcal{F}$, and note that the contribution of each pair $x_{2i-1}x_{2i}$ to the sum $\sum_{x_j \in F} f(x_j)$ is zero, if the edge $x_{2i-1}x_{2i}$ does not stab F , and is either $+1$ or -1 otherwise. It thus follows that this sum has, in the notation of Theorem A.1.1, the distribution S_r for some $r \leq Cn^{1-1/t} \log n$. Thus the probability it is at least α in absolute value can be bounded by $2e^{-\alpha^2/2r}$. As shown in the first paragraph of the proof of Theorem 14.5.1, the total number of members of $\mathcal{F}$ does not exceed $n^{2^{c't \log t}}$, and thus the probability that there exists a member $F \in \mathcal{F}$ for which the sum $\sum_{x_j \in F} f(x_j)$ exceeds $C'n^{1/2-1/2t} \log n$ is less than 1 for an appropriately chosen constant $C' = C'(c, t)$.

The range space whose set of points is an arbitrary set of points in the plane, and whose ranges are all discs in the plane, has dual shatter function $O(g^2)$. The above theorem thus

shows that it is possible to color any set of n points in the plane red and blue, such that the absolute value of the difference between the number of red points and the number of blue points inside any disc would not exceed $n^{1/4+o(1)}$. Similar results can be proved for many other geometric range spaces.

18.5.6 14.6 Exercises

1. Let A be a set of n points in the plane, and let $\mathcal{F}$ be the set of all intersections of A with an open triangle in the plane. Prove that the discrepancy of $\mathcal{F}$ does not exceed $n^{1/4+o(1)}$.
2. Prove that n distinct points in the plane determine at most $O(n^{4/3})$ unit distances.

The Probabilistic Lens: Efficient Packing

Let $C \subset \mathbb{R}^n$ be bounded with Riemann measure $\mu = \mu(C) > 0$. Let $N(C, x)$ denote the maximal number of disjoint translates of C that may be packed in a cube of side x and define the packing constant

$$\delta(C) = \mu(C) \lim_{x \rightarrow \infty} N(C, x) x^{-n},$$

the maximal proportion of space that may be packed by copies of C . The following result improves the one described in Section 3.4.

Theorem 1 Let C be bounded, convex and centrally symmetric about the origin. Then

$$\delta(C) \geq 2^{-(n-1)}.$$

Proof. Fix $\epsilon > 0$. Normalize so $\mu = \mu(C) = 2 - \epsilon$. For any real z let C_z denote the "slab" of $(z_1, \dots, z_{n-1}) \in \mathbb{R}^{n-1}$ such that $(z_1, \dots, z_{n-1}, z) \in C$ and let $\mu(C_z)$ be the usual $(n-1)$ -dimensional measure of C_z . Riemann measurability implies

$$\lim_{\gamma \rightarrow 0} \sum_{m \in \mathbb{Z}} \mu(C_{m\gamma}) \gamma = \mu(C).$$

Let K be an integer sufficiently large so that

$$\sum_{m \in \mathbb{Z}} \mu(C_{mK^{-(n-1)}}) K^{-(n-1)} < 2$$

and further that all points of C have all coordinates less than $K/2$.

For $1 \leq i \leq n-1$ let $v_i \in \mathbb{R}^n$ be that vector with all coordinates zero except K as the i th coordinate. Let

$$v = (z_1, \dots, z_{n-1}, K^{-(n-1)})$$

where $z_1, \dots, z_{n-1}$ are chosen uniformly and independently from the real interval $[0, K)$. Let Λ_v denote the lattice generated by the v 's; that is,

$$\begin{aligned} \Lambda_v &= \{m_1 v_1 + \dots + m_{n-1} v_{n-1} + m v : m_1, \dots, m_{n-1}, m \in \mathbb{Z}\} \\ &= \{(m z_1 + m_1 K, \dots, m z_{n-1} + m_{n-1} K, m K^{-(n-1)}) : \\ &\quad m_1, \dots, m_{n-1}, m \in \mathbb{Z}\}. \end{aligned}$$

Let $\theta(x)$ denote that unique $x' \in (-\frac{1}{2}K, \frac{1}{2}K]$ so that $x - mK = x'$ for some $m \in \mathbb{Z}$. For $m \in \mathbb{Z}$ let A_m be the event that some $m_1 v_1 + \dots + m_{n-1} v_{n-1} + m v \in C$. Since all coordinates of all points of C are less than $K/2$, A_m occurs if and only if

$$(\theta(mz_1), \dots, \theta(mz_{n-1}), mK^{-(n-1)}) \in C$$

which occurs if and only if $(\theta(mz_1), \dots, \theta(mz_{n-1})) \in C_{mK^{-(n-1)}}$. The independence and uniformity of the z_i over $[0, K)$ implies the independence and uniformity of the $\theta(z_i)$ over $(-\frac{1}{2}K, \frac{1}{2}K]$ and so

$$\Pr[A_m] = K^{-(n-1)} \mu(C_{mK^{-(n-1)}})$$

Summing over positive m , and employing the central symmetry,

$$\sum_{m>0} \Pr[A_m] < \frac{1}{2} \sum_{m \in \mathbb{Z}} K^{-(n-1)} \mu(C_{mK^{-(n-1)}}) < \frac{1}{2} \cdot 2 = 1$$

Hence there exists v with all $\{A_m\}_{m>0}$ not holding. By the central symmetry A_m and A_{-m} are the same event so no $\{A_m\}_{m \neq 0}$ holds. When $m = 0$ the points $m_1v_1 + \dots + m_{n-1}v_{n-1} = K(m_1, \dots, m_{n-1}, 0)$ all lie outside C except the origin. For this v

$$\Lambda_v \cap C = \{0\}$$

Consider the set of translates $C + 2w, w \in \Lambda_v$. Suppose

$$z = c_1 + 2w_1 = c_2 + 2w_2 \text{ with } c_1, c_2 \in C, w_1, w_2 \in \Lambda_v$$

Then $\frac{1}{2}(c_1 - c_2) = w_2 - w_1$. From convexity and central symmetry $\frac{1}{2}(c_1 - c_2) \in C$. As $w_2 - w_1 \in \Lambda_v$, it is zero and hence $c_1 = c_2$ and $w_1 = w_2$. That is, the translates form a packing of $\mathbb{R}^n$. As $\det(2\Lambda_v) = 2^n \det(\Lambda_v) = 2^n$ this packing has density $2^{-n}\mu = 2^{-n}(2 - \epsilon)$. As $\epsilon > 0$ was arbitrary, $\delta(C) \geq 2^{-(n-1)}$.

18.6 15 Codes, Games and Entropy

18.6.1 15.1 Codes

Suppose we want to send a message, here considered a string of bits, across a noisy channel. There is a probability p that any bit sent will be received incorrectly. The value p is a parameter of the channel and cannot be changed. We assume that p is both the probability that a sent zero is received as a one and that a sent one is received as a zero. Sent bits are always received, but perhaps incorrectly. We further assume that the events that the bits are received incorrectly are mutually independent. The case $p = 0.1$ will provide a typical example.

How can we improve the reliability of the system? One simple way is to send each bit three times. When the three bits are received we use majority rule to decode. The probability of incorrect decoding is then $3p^2 + p^3 = 0.031$ in our instance. We have sacrificed speed - the rate of transmission of this method

is $\frac{1}{3}$ - and gained accuracy in return. If we send each bit five times and use majority rule to decode, the probability of incorrect decoding drops to 0.01051 but the rate of transmission also drops to $\frac{1}{5}$. Clearly we may make the probability of incorrect decoding as low as needed, but seemingly with the trade-off that the rate of transmission tends to zero. It is the fundamental theorem of Information Theory due to Claude Shannon - that this trade-off is not necessary: there are codes with rate of transmission approaching a positive constant (dependent on p) with probability of incorrect transmission approaching zero.

A coding scheme consists of positive integers m, n , a function $f : \{0, 1\}^m \rightarrow \{0, 1\}^n$ called the encoding function, and a function $g : \{0, 1\}^n \rightarrow \{0, 1\}^m$ called the decoding function. The notion is that a message (or segment of message) $x \in \{0, 1\}^m$ will be encoded and sent as $f(x)$

and a received message $y \in \{0, 1\}^n$ will be decoded as $g(y)$. The rate of transmission of such a scheme is defined as m/n . Let $E = (e_1, \dots, e_n)$ be a random string defined by $\Pr[e_i = 1] = p$, $\Pr[e_i = 0] = 1 - p$, the values e_i mutually independent. We define the probability of correct transmission as $\Pr[g(f(x) + E) = x]$. Here x is assumed to be uniformly distributed over $\{0, 1\}^m$ and independent of E ; '+' here is mod2 vector addition.

A crucial role is played by the entropy function

$$H(p) = -p \log_2 p - (1 - p) \log_2 (1 - p)$$

defined for $0 < p < 1$. For any fixed p the entropy function appears in the asymptotic formula

$$\binom{n}{pn} = \frac{n^n e^{-n}}{(pn)^{pn} e^{-pn} ((1-p)n)^{(1-p)n} e^{-(1-p)n}} (1 + o(1))^n = 2^{n(H(p) + o(1))}.$$

For $0 < p < \frac{1}{2}$ we further bound

$$\sum_{i \leq pn} \binom{n}{i} \leq (1 + pn) \binom{n}{pn} = 2^{n(H(p) + o(1))}.$$

Theorem 15.1.1 [Shannon's Theorem] Let $0 < p < \frac{1}{2}$ be fixed. For $\epsilon > 0$ arbitrarily small there exists a coding scheme with rate of transmission greater than $1 - H(p) - \epsilon$ and probability of incorrect transmission less than ϵ .

Remark. It is not difficult to show that, for every such p , any coding scheme whose rate of transmission exceeds $1 - H(p) + \epsilon$ must have a significant error probability. Indeed, if $f(x)$, the image of x , is transmitted, then with high probability, the obtained output, y , is of distance $(1 + o(1))pn$ from $f(x)$. Hence, if there are 2^m input words, the total size of all typical outputs is about $2^m \cdot \binom{n}{pn} = 2^{m + (1 + o(1))H(p)n}$. If this quantity is much larger than 2^n , then there must be significant overlaps between the output sets of different input words, making the decoding likely to err.

Proof. Let $\delta > 0$ be such that $p + \delta < \frac{1}{2}$ and $H(p + \delta) < H(p) + \epsilon/2$. For n large set $m = n(1 - H(p) - \epsilon)$, guaranteeing the rate of transmission. Let $f : \{0, 1\}^m \rightarrow \{0, 1\}^n$ be a random function - each $f(x)$ uniformly and independently chosen. Given f define the decoding function $g : \{0, 1\}^n \rightarrow \{0, 1\}^m$ by setting $g(y) = x$ if x is the unique vector in $\{0, 1\}^m$ whose image, $f(x)$, is within $n(p + \delta)$ of y . We measure distance by the Hamming metric $\rho : \rho(y, y')$ is the number of coordinates in which y, y' differ. If there is no such x , or more than one such x , then we shall consider decoding to be incorrect.

There are two ways decoding can be incorrect. Possibly $f(x) + E$ is not within $n(p + \delta)$ of $f(x)$. The distance from $f(x) + E$ to $f(x)$ is simply the number of ones in E , which has a binomial distribution $B(n, p)$ and so this occurs with probability $o(1)$ (in fact, with exponentially small probability.) The only other possibility is that there is some $x' \neq x$ with $f(x') \in S$, where S is the set of y' within $n(p + \delta)$ of $f(x) + E$. Conditioning on the values $f(x)$ and E , $f(x')$ is still uniformly distributed over $\{0, 1\}^n$ and hence this occurs with probability $|S|2^{-n}$ for any particular x' and thus with total probability at most

$$2^m |S| 2^{-n} < 2^{-n(\epsilon/2 + o(1))} = o(1).$$

The total probability for incorrect decoding from both sources is thus $o(1)$ and, in fact, exponentially small. For n sufficiently large this is less than ϵ .

The average over all choices of f, x of the probability of incorrect decoding is less than ϵ . Thus there exists a specific f (hence a specific coding scheme) with probability of incorrect coding less than ϵ .

Shannon's Theorem, dealing with the intensely practical subject of communications, puts the shortcomings of the probabilistic approach in sharp contrast. Where is the coding scheme? Supposing that a coding scheme may be found, how can encoding and decoding be rapidly processed?

A group code is a coding scheme in which the map $f : \{0, 1\}^m \rightarrow \{0, 1\}^n$ is linear; that is, $f(0) = 0$ and $f(x + x') = f(x) + f(x')$, all calculations done mod 2. Group codes are of particular interest, in part because of the ease of encoding.

Theorem 15.1.2 Let $0 < p < \frac{1}{2}$ be fixed. For $\epsilon > 0$ arbitrarily small there exists a group code with rate of transmission greater than $1 - H(p) - \epsilon$ and probability of incorrect transmission less than ϵ .

Proof. For $1 \leq i \leq m$ let $u_i \in \{0, 1\}^m$ be that vector with a one in position i , all other entries zero. Let $f(u_1), \dots, f(u_m)$ be chosen randomly and independently and then extend f by setting

$$f(\epsilon_1 u_1 + \dots + \epsilon_m u_m) = \epsilon_1 f(u_1) + \dots + \epsilon_m f(u_m).$$

We follow the proof of Shannon's Theorem until bounding the probability that $f(x) + E$ lies within $n(p + \delta)$ of $f(x)$. Set $z = x - x' = \epsilon_1 u_1 + \dots + \epsilon_m u_m$, again all mod 2. As $x \neq x', z \neq 0$. Reorder for convenience so that $\epsilon_m = 1$.

By linearity $f(z) = f(x) - f(x')$ so we bound $\Pr[f(z) \in S]$, where S is the set of vectors within $n(p + \delta)$ of E . Fixing E and the $f(u_i), i < m, f(z)$ still has an additive term $f(u_m)$ that is uniform and independent. Hence $f(z)$ is distributed uniformly. Thus $\Pr[f(z) \in S] = |S|2^{-n}$ and the remainder of the proof is as in Shannon's Theorem.

18.6.2 15.2 Liar Game

Paul is trying to find a number $x \in \{1, \dots, n\}$ from a recalcitrant and mendacious Carole. He may ask q questions of the form "Is $x \in S$?", where S can be any subset of the possibilities. The questions are asked sequentially and Paul's choice of his i th question can depend on previous responses. Carole is allowed to lie - but she can lie at most k times. For which n, q, k can Paul determine the number?

When $k = 0$ Paul can win exactly when $n \leq 2^q$. The values $n = 100, q = 10, k = 1$ make for an amusing parlor game. Carole is hardly a passive observer; she may play an adversary strategy. By that we mean that she does not select an x in advance but answers consistently with at least one x . At the end of the game if her answers were consistent with more than one x then she has won. The game, called the (n, q, k) -Liar Game, is now a perfect information game with no hidden move and no draws. Hence either Paul or Carole has a perfect winning strategy. But who?

We describe an equivalent game, the Chip-Liar Game. There is a board with positions $0, 1, \dots, k$. There are n chips labeled $1, \dots, n$ which are initially at position k . There are q rounds. On each round Paul selects a set S of the chips. Carole can either move every chip not in S one position to the left or move every chip in S one position to the left. (Here position $i - 1$ is one position to the left of position i . Chips moved one position to the left from position 0 are removed from the board.) At the end of the q rounds Carole wins if there is more than one chip remaining on the board and Paul wins if there is one or zero chip remaining on the board. Basically, chip i at position j represents that the answer $x = i$ has already received $k - j$ lies; Paul selecting S represents his asking if $x \in S$; Carole moving the chips not in S represents a Yes answer, moving the chips in S represents a No answer. (In the Chip-Liar Game Carole can remove all chips from the board while in the Liar Game Carole must play consistently with at least one x . But when Carole removes all chips from the board she automatically has lost and hence this difference does not affect the determination of the winner.)

In the Chip-Liar Game there is no reason to place all chips at position k at the start. More generally, for $x_0, \dots, x_k \geq 0$, we define the $(x_0, \dots, x_k)$, q -Chip-Liar Game to be the above q round game with initial position consisting of x_i chips at position i . This, in turn, corresponds to a Liar Game in which there are x_i possibilities for which Carole is constrained to lie at most i times.

Let us define $B(q, j)$ as the probability that in q flips of a fair coin there are at most j heads. Of course, we have the exact formula

$$B(q, j) = 2^{-q} \sum_{i=0}^j \binom{q}{i}.$$

Theorem 15.2.1 If

$$\sum_{i=0}^k x_i B(q, i) > 1$$

then Carole wins the $(x_0, \dots, x_k)$, q -Chip-Liar Game.

Corollary 15.2.2 If

$$n > \frac{2^q}{\sum_{i=0}^k \binom{q}{i}}$$

then Carole wins the (n, q, k) -Liar Game.

Proof [Theorem 15.2.1]. Fix a strategy for Paul. Now Carole plays randomly! That is, at each round, after Paul has selected a set S of chips Carole flips a coin - if it comes up heads she moves every chip not in S one position to the left and if it comes up tails she moves every chip in S one position to the left. For each chip c let I_c be the indicator random variable for c remaining on the board at the end of the game. Set $X = \sum I_c$, the number of chips remaining on the board at the end of the game. Consider a single chip c . Each round Paul may have chosen $c \in S$ or $c \notin S$ but in either case c is moved to the left with probability $\frac{1}{2}$. Suppose c starts at position i . It remains on the board at the end of the game if and only if in the q rounds it has been moved to the left at most i times. Then $E[I_c]$, the probability of this occurring, is precisely $B(q, i)$. By linearity of expectation $E[X] = \sum_{i=0}^k x_i B(q, i)$. The assumption of the theorem gives $E[X] > 1$. But then $X > 1$ must occur with positive probability. That is, Carole must win with positive probability.

No strategy of Paul allows him to always win. But this is a perfect information game with no draws so someone has a perfect strategy that always wins. That someone isn't Paul, so it must be Carole.

The above proof certainly illustrated the magical element of the probabilistic method. Carole has a winning strategy but what is it? The general notion of moving from a probabilistic existence proof to an explicit construction is called derandomization and will be dealt with in detail in the next chapter. Here we can give an explicit strategy. With l moves remaining in the game and y_i chips on position i define the weight of the position as $\sum_i y_i B(l, i)$ - note this is $E[Y]$, where Y is the number of chips that would remain on the board should Carole play the rest of the game randomly. Carole's explicit strategy is to always move so as to maximize the weight.

Consider any position with weight W and any move S by Paul. Let W^y, W^n be the new weights should Carole move all chips not in S or all chips in S , respectively. We claim $W = \frac{1}{2}(W^y + W^n)$. One argument is that by linearity this identity reduces to the case of one chip and it then follows from the identity $B(l, j) = \frac{1}{2}(B(l-1, j) + B(l-1, j-1))$. But we needn't actually do any calculation. Carole's playing randomly can be thought of as first flipping a coin

to decide on her first move and then playing randomly so that $E[Y]$ is the average of the two conditional expectations.

At the start of the game, by assumption, the weight is larger than one. Carole's explicit strategy assures that the weight does not decrease so at the end of the game the weight is larger than one. But at the end of the game the weight is the number of chips remaining. Being larger than one, Carole has won the game.

The converse of the theorem, and even the corollary, is false. Consider the Liar Game with $n = 5, q = 5$ questions and $k = 1$ possible lie. In the Chip-Liar version this is the $(0, 5)$, 5-Chip-Liar Game. Here $B(5, 1) = \frac{6}{32}$ and $5 \cdot \frac{6}{32} \leq 1$. Still, Carole wins with perfect play. The problem is that Paul has no good first move. Suppose he selects two chips as S (asks "Is $x \leq 2$?" in the Liar Game). Then Carole moves the two chips one to the left (responds Yes) leaving the position $(2, 3)$ with four questions remaining. As $2B(4, 0) + 3B(4, 1) = \frac{17}{16} > 1$, Carole will now win. It is easy to check that all other moves of Paul fail. The difficulty here is that Paul was in a position with weight $W \leq 1$ but was unable to find a move such that $W^y \leq 1$ and $W^n \leq 1$.

18.6.3 15.3 Tenure Game

Paul, Chair of Department, is trying to promote one of his faculty to tenure but standing in his way is a recalcitrant and meanspirited Carole, the Provost. There are k pretenure levels, labeled $1, \dots, k$, level 1 the highest, and a level 0, representing tenure. For our purposes each faculty member is represented by a chip. The $(x_1, \dots, x_k)$ -Tenure Game begins with x_i chips at level i for $1 \leq i \leq k$ and no chips on level zero. Each year Paul presents a set S of chips to Carole. Carole may either:

- Promote all chips in S and fire the others or
- Promote all chips not in S and fire those in S .

Promote, as used above, means to move from level i to level $i - 1$. Fired means just that: removing the chip from the game. If a chip reaches level 0 then Paul is the winner. The draconian promotion or perish provision ensures that the game will end within k years with either Paul winning or Carole having successfully eliminated all chips.

Theorem 15.3.1 If $\sum_i x_i 2^{-i} < 1$ then Carole wins the $(x_1, \dots, x_k)$ -Tenure Game.

Proof. Fix a strategy for Paul. Now Carole plays randomly! That is, at each round, after Paul has selected a set S of chips Carole flips a coin - if it comes up heads she moves every chip not in S one position to the left and if it comes up tails she moves every chip in S one position to the left. For each chip c let I_c be the indicator random variable for c reaching level 0. Set $X = \sum I_c$, the number of chips reaching level 0 at the end of the game. Consider a single chip c . Each round Paul may have chosen $c \in S$ or $c \notin S$ but in either case c is moved to the left with probability $\frac{1}{2}$. Suppose c starts at position i . It remains on the board at the end of the game if and only if the first i coin flips of Carole led to promotions for c . Then $E[I_c]$, the probability

of this occurring, is precisely 2^{-i} . By linearity of expectation $E[X] = \sum_{i=1}^k x_i 2^{-i}$. The assumption of the theorem gives $E[X] < 1$. But then $X < 1$ must occur with positive probability. That is, Carole must win with positive probability.

No strategy of Paul allows him to always win. But this is a perfect information game with no draws so someone has a perfect strategy that always wins. That someone isn't Paul, so it must be Carole.

As with the Liar Game we may derandomize the above argument to give an explicit strategy for Carole. With y_i chips on position i define the weight of the position as $\sum_i y_i 2^{-i}$ - note this is $E[Y]$, where Y is the number of chips that would reach level 0 should Carole play the rest

of the game randomly. Carole's explicit strategy is to always move so as to minimize the weight. Consider any position with weight W and any move S by Paul. Let W^y, W^n be the new weights should Carole move all chips not in S or all chips in S , respectively. As in the Liar Game $W = \frac{1}{2}(W^y + W^n)$. At the start of the game, by assumption, the weight is less than one. Carole's explicit strategy assures that the weight does not increase so at all times the weight is smaller than one. A chip at level 0 would add one to the weight by itself so that this never occurs and hence Carole wins.

In the Liar Game the sufficient condition for Carole to win was not necessary because Paul did not always have an appropriately splitting move. Here, however, we have an amusing lemma.

Lemma 15.3.2 If a set of chips has weight at least one it may be split into two parts, each of weight at least $\frac{1}{2}$.

Proof. There must be two chips at some position i , otherwise the weight is less than one. If there are two chips at position 1 simply split them. If there are two chips at position $i > 1$ glue them together, and consider them as one superchip at position $i - 1$. Then the proof follows by induction on the number of chips.

Theorem 15.3.3 If $\sum x_i 2^{-i} \geq 1$ then Paul wins the $(x_1, \dots, x_k)$ -Tenure Game.

Proof. The initial weight is at least one. Applying the lemma Paul splits the chips into two parts, each of weight at least $\frac{1}{2}$, and sets S equal to one of the parts. Carole moves all chips in one part one position to the left, doubling their weight, leaving a new position of weight at least one. Thus the weight never goes below one. Therefore the game cannot end with all chips having been removed (which would have weight zero) and so it must end with a win for Paul.

18.6.4 15.4 Balancing Vector Game

The balancing vector game is a perfect information game with two players, Pusher and Chooser. There is a parameter $n \geq 1$, and we shall be concerned with asymptotics in n . There are n rounds, each involving vectors in $\mathbb{Z}^n$. There is a position vector $P \in \mathbb{Z}^n$, initially set at 0. Each round has two parts. First Pusher picks $v \in \{-1, +1\}^n$. Then Chooser either resets P to $P + v$ or to $P - v$. At the end of the n th round the payoff to Pusher is $|P|_\infty$, the maximal absolute value of the coordinates of P . Let $\text{VAL}(n)$ denote the value of this game to Pusher, that is, the maximum payoff Pusher can ensure when both players play optimally. Let S_n denote, as usual, the sum of n independent uniform ± 1 random variables.

Theorem 15.4.1 If $\Pr[|S_n| > \alpha] < n^{-1}$ then $\text{VAL}(n) \leq \alpha$.

Proof. Consider the game a win for Pusher if the final $|P|_\infty > \alpha$. Suppose Chooser announces that she will flip a fair coin each round to determine whether to reset P as $P + v$ or $P - v$. Let x_i be the i th coordinate for the final value of the position vector P . Let W_i be the event $|x_i| > \alpha$ and $W = \bigvee_{i=1}^n W_i$ so that W is the event of Pusher winning. Regardless of Pusher's strategy x_i has distribution S_n so that

$$\Pr[W] \leq \sum_{i=1}^n \Pr[|S_n| > \alpha] < 1$$

Pusher cannot always win so Chooser always wins.

Corollary 15.4.2 $\text{VAL}(n) = O(\sqrt{n \ln n})$.

To give a lower bound on $\text{VAL}(n)$ one wants to find a strategy for Pusher that wins against any Chooser. It is not sufficient to find a strategy that does well against a randomly playing Chooser - the Chooser is an adversary. Still, the notion of a randomly playing Chooser motivates the following result.

Theorem 15.4.3 If $\Pr[|S_n| > \alpha] > cn^{-1/2}$, where c is an absolute constant, then $\text{VAL}(n) > \alpha$.

Corollary 15.4.4 $\text{VAL}(n) = \Omega(\sqrt{n \ln n})$ and hence $\text{VAL}(n) = \Theta(\sqrt{n \ln n})$.

Proof [Theorem 15.4.3]. Define, for $x \in \mathbb{Z}, 0 \leq i \leq n$,

$$w_i(x) = \Pr [|x + S_{n-i}| > \alpha].$$

For $P = (x_1, \dots, x_n)$ set $w_i(P) = \sum_{j=1}^n w_i(x_j)$. When P is the position vector at the end of the i th round, $w_i(P)$ may be interpreted as the expected number of coordinates with absolute value greater than α at the end of the game, assuming random play by Chooser. At the beginning of the game $w_0(P) = w_0(0) > c\sqrt{n}$ by assumption. Given position P at the end of round i , Pusher's strategy will be to select $v \in \{-1, +1\}^n$ so that $w_{i+1}(P - v)$ and $w_{i+1}(P + v)$ are close together.

The distribution $x + S_{n-i}$ splits into $x + 1 + S_{n-i-1}$ and $x - 1 + S_{n-i-1}$ depending on the first coin flip so that for any i, x ,

$$w_i(x) = \frac{1}{2} (w_{i+1}(x + 1) + w_{i+1}(x - 1)).$$

Set $P = (x_1, \dots, x_n), v = (v_1, \dots, v_n)$. For $1 \leq j \leq n$ set

$$\Delta_j = w_{i+1}(x_j + 1) - w_{i+1}(x_j - 1)$$

so that

$$w_{i+1}(P + v) - w_{i+1}(P - v) = \sum_{j=1}^n v_j \Delta_j$$

and, for $\epsilon = \pm 1$,

$$w_{i+1}(P + \epsilon v) = w_i(P) + \frac{1}{2} \epsilon \sum_{j=1}^n v_j \Delta_j.$$

Now we bound $|\Delta_j|$. Observe that

$$\Delta_j = \Pr [S_{n-i-1} = y] - \Pr [S_{n-i-1} = z]$$

where y is the unique integer of the same parity as $n - i - 1$ in the interval $(\alpha - (x_j + 1), \alpha - (x_j - 1)]$ and z is the same in $[-\alpha - (x_j + 1), -\alpha - (x_j - 1))$. Let us set

$$g(m) = \max_s \Pr [S_m = s] = \binom{m}{\lfloor m/2 \rfloor} 2^{-m} \sim \sqrt{\frac{2}{\pi m}}$$

so that $|\Delta_j| \leq g(n - i - 1)$ for all j .

A simple strategy for Pusher is then to reorder the coordinates so that $|\Delta_1| \geq \dots \geq |\Delta_n|$ and then select $v_1, \dots, v_n \in \{-1, +1\}$ sequentially, giving $v_i \Delta_i$ the opposite sign of $v_1 \Delta_1 + \dots + v_{i-1} \Delta_{i-1}$. (When $i = 1$ or the sum is zero, choose v_i arbitrarily.) This assures

$$|v_1 \Delta_1 + \dots + v_n \Delta_n| \leq |\Delta_1| \leq g(n - i - 1).$$

Let P^i denote the position vector at the end of the i th round and v Pusher's choice for the $(i + 1)$ st round. Then regardless of Chooser's choice of $\epsilon = \pm 1$,

$$w_{i+1}(P^{i+1}) = w_{i+1}(P^i + \epsilon v) \geq w_i(P^i) - \frac{1}{2} \left| \sum_{j=1}^n v_j \Delta_j \right| \geq w_i(P^i) - \frac{1}{2} g(n - i - 1).$$

Thus

$$w_n(P^n) \geq w_0(P^0) - \frac{1}{2} \sum_{i=0}^{n-1} g(n - i - 1).$$

Simple asymptotics give that the above sum is asymptotic to $(8n/\pi)^{1/2}$. Choosing $c > (2/\pi)^{1/2}$, $w_n(P^n) > 0$. But $w_n(P^n)$ is simply the number of coordinates with absolute value greater than α in the final $P = P^n$. This Pusher strategy assures there is more than zero, hence at least one such coordinate and therefore Pusher wins.

18.6.5 15.5 Nonadaptive Algorithms

Let us modify the balancing game of Section 15.4 by requiring the vectors selected by Pusher to have coordinates zero and one rather than plus and minus one. Let $\text{VAL}^*(n)$ denote the value of the modified game. One can use the bounds on $\text{VAL}(n)$ to show $\text{VAL}^*(n) = \Theta(\sqrt{n \ln n})$.

In Chapter 13 we showed that any family of n sets $S_1, \dots, S_n$ on n points $1, \dots, n$ has discrepancy $O(\sqrt{n})$; that is, there is a coloring $\chi : \{1, \dots, n\} \rightarrow \{-1, +1\}$ so that all $|\chi(S_i)| \leq c\sqrt{n}$. The proof of this result does not yield an effective algorithm for finding such a coloring and indeed it is not known if there is a polynomial time algorithm to do so. Suppose one asks for a nonadaptive or on-line algorithm in the following sense. Instead of being presented the entire data of $S_1, \dots, S_n$ at once, one is presented with the points sequentially. At the j th "round" the algorithm looks at point j - more specifically, at which sets S_i contain j or, equivalently, at the j th column of the incidence matrix. At that stage the algorithm must decide how to color j and, once colored, the coloring cannot be changed. How small can we assure $\max |\chi(S_i)|$ with such an algorithm? We may think of the points as being presented by an adversary. Thinking of the points as their associated column vectors, Pusher as the worst case adversary and Chooser as the algorithm, the best such an algorithm can do is precisely $\text{VAL}^*(n)$.

The requirement that an algorithm be nonadaptive is both stronger and weaker than the requirement that an algorithm take polynomial time. Still, this lends support to the conjecture that there is no polynomial time algorithm for finding a coloring with all $|\chi(S_i)| \leq c\sqrt{n}$.

18.6.6 15.6 Half Liar Game

We modify the Liar Game of Section 15.2 by limiting Carole's mendacity. If the correct answer is Yes then Carole is now required to answer Yes. She may answer Yes when the correct answer is No, and that would count as one of her k lies. Let $A_k(q)$ denote the maximal n for which Paul wins the Half Liar Game with n values, q queries, and a maximum of k lies.

Theorem 15.6.1 [Dumitriu and Spencer (2004)] For each fixed $k \geq 1$,

$$A_k(q) \sim 2^{q+k} / \binom{q}{k}$$

While the methods below extend to arbitrary k , we give the proof only for the case $k = 1$. This case was first given by Cicalese and Mundici (2000).

Proof. Let us fix a winning strategy for Paul with $n = A_1(q)$. This may be described by a binary decision tree of depth q . For each value i , $1 \leq i \leq n$, let $\sigma_i = (x_{i1}, \dots, x_{iq}) \in \{Y, N\}^q$ be the string of truthful responses to Paul's queries with that value. Let T_i be the set of possible response strings given by Carole with that value. For each $x_{ij} = N$ Carole may lie on the j th question, otherwise responding truthfully. Thus $|T_i| = W(\sigma_i)$, where we define the weight $W(\sigma)$ to be one plus the number of N 's in the string σ . We cannot have any common $\sigma \in T_i \cap T_j$ as then Carole could respond with σ and Paul would not be able to distinguish i, j . Thus

$$\sum_{i=1}^n W(\sigma_i) \leq 2^q. \quad (15.1)$$

For a given u , call σ Carole friendly if $W(\sigma) \leq 1 + \frac{1}{2}(q - u)$, otherwise Paul friendly. There are at most $2^q \Pr[S_q \leq -u]$ Carole friendly σ 's, where S_q is the sum of q independent ± 1 random variables. By (15.1) there are at most $2^{q+1}/(q - u)$ Paul friendly σ 's. Thus

$$n \leq 2^q \Pr[S_q \leq -u] + \frac{2^{q+1}}{q - u}$$

The optimization of u is left as an exercise but even taking a suboptimal $u = \lfloor q^{2/3} \rfloor$ gives

$$A_1(q) = n \leq (1 + o(1)) \frac{2^{q+1}}{q}.$$

For larger n Paul cannot have a winning strategy and thus Carole must have a winning adversary strategy. Intriguingly, this argument does not yield an explicit strategy for Carole.

In the other direction let $\epsilon > 0$ be fixed and small and set $n = \lfloor (1 - \epsilon)2^{q+1}/q \rfloor$. We will give a strategy for Paul. For $r \geq 1$ let M_r denote those $\sigma \in \{Y, N\}^r$ with at least $\frac{1}{2}(r - u)N$'s and let $f(r) = |M_r|$. For definiteness, take $u = \lfloor r^{2/3} \rfloor$. Then $f(r) \sim 2^r$. We first message n . Pick r with, say, $10/\epsilon \leq n/f(r) \leq 21/\epsilon$, set $A = \lceil n/f(r) \rceil$, and boost n to $n = Af(r)$. As the boost (which makes things only harder for Paul) was by a factor less than $1 + (\epsilon/10)$ the new n still has $n \leq (1 - \epsilon/2)2^{q+1}/q$.

Paul associates the $n = f(r)A$ values with pairs (σ, j) , $\sigma \in M_r$, $1 \leq j \leq A$. For his first r queries he asks for the coordinates of σ . Carole responds τ , which can differ from the truthful σ in at most one coordinate. Thus τ has at most $\frac{1}{2}(r + u) + 1Y$'s. (Basically, these r queries are nearly even splits and force Carole to answer No nearly half the time.) What does Paul know at this moment? If Carole has not lied the answer must be one of the A values (τ, j) . If Carole has lied the answer must be one of the at most $\frac{1}{2}A(r + u + 2)$ values (τ^+, j) , where τ^+ is derived from τ by shifting a Y to a N .

Set $s = q - r$, the number of remaining queries. As A is bounded and $2^r A \sim f(r)A = n = \Theta(2^q/q)$, we have $r = q - \log_2 q - O(1)$. In particular, $r \sim q$, the first r queries were the preponderance of the queries. Then

$$A \leq (1 + o(1))n2^{-r} \leq \left(1 - \frac{\epsilon}{2} + o(1)\right) 2^{q+1-r}r$$

and

$$A \frac{r + u + 2}{2} \sim A \frac{r}{2} \leq \left(1 - \frac{\epsilon}{2} + o(1)\right) 2^s$$

Paul may now give further ground and allow Carole to lie in either direction for the remaining s questions. This is the (x_0, x_1) , s -Chip-Liar Game with $x_0 = A$ and $x_1 \leq (1 - \epsilon/2 + o(1))2^s$. The endgame strategy required at this point is given in the exercises.

18.6.7 15.7 Entropy

Let X be a random variable taking values in some range S , and let $P(X = x)$ denote the probability that the value of X is x . The binary entropy of X , denoted by $H(X)$, is defined by

$$H(X) = \sum_{x \in S} P(X = x) \log_2 \frac{1}{P(X = x)}$$

If Y is another random variable, taking values in T , and (X, Y) is the random variable taking values in $S \times T$ according to the joint distribution of X and Y , then the conditional entropy of X given Y is

$$H(X | Y) = H(X, Y) - H(Y)$$

In this section we prove some simple properties of the entropy function and describe several surprising combinatorial and geometric applications. Intuitively, the entropy of a random variable measures the amount of information it encodes. This provides an intuitive explanation to the four parts of the next simple lemma. The formal proof, given below, uses the properties of the functions $\log z$ and $z \log z$, where here, and in the rest of this section, all logarithms are in base 2.

Lemma 15.7.1 Let X, Y and Z be three random variables taking values in S, T and U , respectively. Then

- (i) $H(X) \leq \log |S|$.
- (ii) $H(X, Y) \geq H(X)$.
- (iii) $H(X, Y) \leq H(X) + H(Y)$.
- (iv) $H(X | Y, Z) \leq H(X | Y)$.

Proof.

- (i) Since the function $\log z$ is concave it follows, by Jensen's Inequality, that

$$\begin{aligned} H(X) &= \sum_{i \in S} P(X = i) \log \frac{1}{P(X = i)} \\ &\leq \log \left[\sum_{i \in S} P(X = i) \frac{1}{P(X = i)} \right] = \log |S|. \end{aligned}$$

- (ii) By the monotonicity of $\log z$ for all $z > 0$,

$$\begin{aligned} H(X, Y) &= \sum_{i \in S} \sum_{j \in T} P(X = i, Y = j) \log \frac{1}{P(X = i, Y = j)} \\ &\geq \sum_{i \in S} \sum_{j \in T} P(X = i, Y = j) \log \frac{1}{P(X = i)} \\ &= \sum_{i \in S} P(X = i) \log \frac{1}{P(X = i)} = H(X) \end{aligned}$$

- (iii) By definition

$$\begin{aligned} &H(X) + H(Y) - H(X, Y) \\ &= \sum_{i \in S} \sum_{j \in T} P(X = i, Y = j) \log \frac{P(X = i, Y = j)}{P(X = i)P(Y = j)} \\ &= \sum_{i \in S} \sum_{j \in T} P(X = i)P(Y = j)f(z_{ij}) \end{aligned}$$

where $f(z) = z \log z$ and $z_{ij} = P(X = i, Y = j)/[P(X = i)P(Y = j)]$. Since $f(z)$ is convex it follows, by Jensen's Inequality, that the last quantity is at least

$$f \left(\sum_{i \in S} \sum_{j \in T} P(X = i)P(Y = j)z_{ij} \right) = f(1) = 0.$$

- (iv) Note that

$$\begin{aligned}
H(X | Y) &= H(X, Y) - H(Y) \\
&= \sum_{i \in S} \sum_{j \in T} P(X = i, Y = j) \log \frac{P(Y = j)}{P(X = i, Y = j)}.
\end{aligned}$$

Similarly

$$\begin{aligned}
&H(X | Y, Z) \\
&= \sum_{i \in S} \sum_{j \in T} \sum_{k \in U} P(X = i, Y = j, Z = k) \log \frac{P(Y = j, Z = k)}{P(X = i, Y = j, Z = k)}.
\end{aligned}$$

Therefore

$$\begin{aligned}
&H(X | Y) - H(X | Y, Z) \\
&= \sum_{i \in S} \sum_{j \in T} \sum_{k \in U} P(X = i, Y = j, Z = k) \\
&\quad \cdot \log \frac{P(Y = j)P(X = i, Y = j, Z = k)}{P(X = i, Y = j)P(Y = j, Z = k)} \\
&= \sum_{i \in S} \sum_{j \in T} \sum_{k \in U} \frac{P(X = i, Y = j)P(Y = j, Z = k)}{P(Y = j)} f(z_{ijk}),
\end{aligned}$$

where $f(z) = z \log z$ and

$$z_{ijk} = \frac{P(Y = j)P(X = i, Y = j, Z = k)}{P(X = i, Y = j)P(Y = j, Z = k)}$$

By the convexity of $f(z)$ and since

$$\sum_{i \in S} \sum_{j \in T} \sum_{k \in U} \frac{P(X = i, Y = j)P(Y = j, Z = k)}{P(Y = j)} = 1,$$

it follows that the above quantity is at least

$$f \left(\sum_{i \in S} \sum_{j \in T} \sum_{k \in U} \frac{P(X = i, Y = j)P(Y = j, Z = k)}{P(Y = j)} z_{ijk} \right) = f(1) = 0.$$

The following simple but useful fact that the entropy is subadditive has already been applied in Section 13.2.

Proposition 15.7.2 Let $X = (X_1, \dots, X_n)$ be a random variable taking values in the set $S = S_1 \times S_2 \times \dots \times S_n$, where each of the coordinates X_i of X is a random variable taking values in S_i . Then

$$H(X) \leq \sum_{i=1}^n H(X_i).$$

Proof. This follows by induction from Lemma 15.7.1, part (iii).

The above proposition is used in Kleitman, Shearer and Sturtevant (1981) to derive several interesting applications in Extremal Finite Set Theory, including an upper estimate for the maximum possible cardinality of a family of k -sets in which the intersection of no two is contained in a third. The basic idea in Kleitman et al. (1981) can be illustrated by the following very simple corollary of the last proposition.

Corollary 15.7.3 Let $\mathcal{F}$ be a family of subsets of $\{1, 2, \dots, n\}$ and let p_i denote the fraction of sets in $\mathcal{F}$ that contain i . Then

$$|\mathcal{F}| \leq 2^{\sum_{i=1}^n H(p_i)},$$

where $H(p) = -p \log_2 p - (1-p) \log_2 (1-p)$.

Proof. Associate each set $F \in \mathcal{F}$ with its characteristic vector $v(F)$, which is a binary vector of length n . Let $X = (X_1, \dots, X_n)$ be the random variable taking values in $\{0, 1\}^n$, where $P(X = v(F)) = 1/|\mathcal{F}|$ for all $F \in \mathcal{F}$. Clearly $H(X) = |\mathcal{F}| \cdot (1/|\mathcal{F}|) \log |\mathcal{F}| = \log |\mathcal{F}|$ and since here $H(X_i) = H(p_i)$ for all $1 \leq i \leq n$, the result follows from Proposition 15.7.2.

The following interesting extension of Proposition 15.7.2 has been proved by Shearer; see Chung, Frankl, Graham and Shearer (1986). As in that proposition, let $X = (X_1, \dots, X_n)$ be a random variable taking values in the set $S = S_1 \times S_2 \times \dots \times S_n$, where each X_i is a random variable taking values in S_i . For a subset I of $\{1, 2, \dots, n\}$, let $X(I)$ denote the random variable $(X_i)_{i \in I}$.

Proposition 15.7.4 Let $X = (X_1, \dots, X_n)$ and S be as above. If $\mathcal{G}$ is a family of subsets of $\{1, \dots, n\}$ and each $i \in \{1, \dots, n\}$ belongs to at least k members of $\mathcal{G}$ then

$$kH(X) \leq \sum_{G \in \mathcal{G}} H(X(G)).$$

Proof. We apply induction on k . For $k = 1$, replace each set $G \in \mathcal{G}$ by a subset of it to obtain a family $\mathcal{G}'$ whose members form a partition of $\{1, \dots, n\}$. By Lemma 15.7.1, part (ii), $\sum_{G \in \mathcal{G}} H(X(G)) \geq \sum_{G' \in \mathcal{G}'} H(X(G'))$ and by Lemma 15.7.1, part (iii), $\sum_{G' \in \mathcal{G}'} H(X(G')) \geq H(X)$, supplying the desired result for $k = 1$.

Assuming the result holds for $k - 1$, we prove it for $k (\geq 2)$. If there is a $G \in \mathcal{G}$ with $G = \{1, \dots, n\}$, the result follows from the induction hypothesis. Otherwise, let G, G' be two members of $\mathcal{G}$. By applying Lemma 15.7.1, part (iv), we conclude that

$$H(X(G \setminus G') \mid X(G \cap G'), X(G' \setminus G)) \leq H(X(G \setminus G') \mid X(G \cap G')),$$

implying that

$$H(X(G \cup G')) - H(X(G')) \leq H(X(G)) - H(X(G \cap G')).$$

Therefore $H(X(G \cup G')) + H(X(G \cap G')) \leq H(X(G)) + H(X(G'))$. It follows that if we modify $\mathcal{G}$ by replacing G and G' by their union and intersection, then the sum $\sum_{G \in \mathcal{G}} H(X(G))$ can only decrease. After a finite number of such modifications we can reach the case in which one of the sets in $\mathcal{G}$ is $\{1, \dots, n\}$, and as this case has already been proved, this completes the proof.

Corollary 15.7.5 Let $\mathcal{F}$ be a family of vectors in $S_1 \times S_2 \times \dots \times S_n$. Let $\mathcal{G} = \{G_1, G_2, \dots, G_m\}$ be a collection of subsets of $N = \{1, 2, \dots, n\}$, and suppose that each element $i \in N$ belongs to at least k members of $\mathcal{G}$. For each $1 \leq i \leq m$ let $\mathcal{F}_i$ be the set of all projections of the members of $\mathcal{F}$ on G_i . Then

$$|\mathcal{F}|^k \leq \prod_{i=1}^m |\mathcal{F}_i|$$

Proof. Let $X = (X_1, \dots, X_n)$ be the random variable taking values in $\mathcal{F}$, where $P(X = F) = 1/|\mathcal{F}|$ for all $F \in \mathcal{F}$. By Proposition 15.7.4,

$$kH(X) \leq \sum_{i=1}^m H(X(G_i))$$

But $H(X) = \log_2 |\mathcal{F}|$, whereas by Lemma 15.7.1, part (i), $H(X(G_i)) \leq \log_2 |\mathcal{F}_i|$, implying the desired result.

Since the volume of every d -dimensional measurable set in $\mathbb{R}^n$ can be approximated by the volume of an appropriate approximation of it by standard aligned boxes in a fine enough grid, the last result has the following geometric application, proved in Loomis and Whitney (1949) in a different manner.

Corollary 15.7.6 Let B be a measurable body in the n -dimensional Euclidean space, let $\text{Vol}(B)$ denote its (n -dimensional) volume, and let $\text{Vol}(B_i)$ denote the ($n-1$)-dimensional volume of the projection of B on the hyperplane spanned by all coordinates besides the i th one. Then

$$(\text{Vol}(B))^{n-1} \leq \prod_{i=1}^n \text{Vol}(B_i)$$

If $S_i = \{0, 1\}$ for all i in Corollary 15.7.5, we get the following statement about set systems.

Corollary 15.7.7 [Chung et al. (1986)] Let N be a finite set, and let $\mathcal{F}$ be a family of subsets of N . Let $\mathcal{G} = \{G_1, \dots, G_m\}$ be a collection of subsets of N , and suppose that each element of S belongs to at least k members of $\mathcal{G}$. For each $1 \leq i \leq m$ define $\mathcal{F}_i = \{F \cap G_i : F \in \mathcal{F}\}$. Then

$$|\mathcal{F}|^k \leq \prod_{i=1}^m |\mathcal{F}_i|$$

We close the section with the following application of the last result, given in Chung et al. (1986).

Corollary 15.7.8 Let $\mathcal{F}$ be a family of graphs on the labeled set of vertices $\{1, \dots, t\}$, and suppose that for any two members of $\mathcal{F}$ there is a triangle contained in both of them. Then

$$|\mathcal{F}| < \frac{1}{4} 2^{\binom{t}{2}}.$$

Proof. Let N be the set of all $\binom{t}{2}$ unordered pairs of vertices in $T = \{1, 2, \dots, t\}$, and consider $\mathcal{F}$ as a family of subsets of N . Let $\mathcal{G}$ be the family of all subsets of N consisting of the edge sets of unions of two vertex disjoint nearly equal complete graphs in T . Let

$$s = \binom{\lceil t/2 \rceil}{2} + \binom{\lfloor t/2 \rfloor}{2}$$

denote the number of edges of such a union, and let m denote the total number of members in $\mathcal{G}$. By symmetry, each edge in N lies in precisely $k = sm / \binom{t}{2}$ members of $\mathcal{G}$. The crucial point is that every two graphs in $\mathcal{F}$ must have at least one common edge in each $G \in \mathcal{G}$, since their intersection contains a triangle (and there are no triangles in the complement of G .) Therefore, in the notation of Corollary 15.7.7, the cardinality of each $\mathcal{F}_i$ is at most 2^{s-1} . We thus conclude that

$$|\mathcal{F}|^{sm / \binom{t}{2}} \leq (2^{s-1})^m,$$

implying that

$$|\mathcal{F}| \leq 2^{\binom{t}{2} - \binom{t}{2}/s},$$

and the desired result follows, as $s < \frac{1}{2} \binom{t}{2}$.

Simonovits and Sós conjectured that if $\mathcal{F}$ satisfies the assumptions of the last corollary, then, in fact,

$$|\mathcal{F}| \leq \frac{1}{8} 2^{\binom{t}{2}},$$

which, if true, is tight. This remains open. It seems plausible to conjecture that there is some absolute constant $\epsilon > 0$, such that for any fixed graph H that is not a star-forest (i.e., a forest each connected component of which is a star), the following holds. Let $\mathcal{F}$ be a family of graphs on the labeled set of vertices $\{1, 2, \dots, t\}$, and suppose that for any two members of $\mathcal{F}$ there is a copy of H contained in both of them. Then

$$|\mathcal{F}| < \left(\frac{1}{2} - \epsilon\right) 2^{\binom{t}{2}}.$$

This is also open, though it is not difficult to show that it is true for every H of chromatic number at least 3, and that the conclusion fails for every star-forest H .

18.6.8 15.8 Exercises

1. Suppose that in the $(x_1, \dots, x_k)$ -Tenure Game of Section 15.3 the object of Paul is to maximize the number of faculty receiving tenure while the object of Carole is to minimize that number. Let v be that number with perfect play. Prove $v = \left\lfloor \sum_{i=1}^k x_i 2^{-i} \right\rfloor$.
2. Let $A_1, \dots, A_n \subseteq \{1, \dots, m\}$ with $\sum_{i=1}^n 2^{-|A_i|} < 1$. Paul and Carole alternately select distinct vertices from $\{1, \dots, m\}$, Paul having the first move, until all vertices have been selected. Carole wins if she has selected all the vertices of some A_i . Paul wins if Carole does not win. Give a winning strategy for Paul.
3. Let $\mathcal{F}$ be a family of graphs on the labeled set of vertices $\{1, 2, \dots, 2t\}$, and suppose that for any two members of $\mathcal{F}$ there is a perfect matching of t edges contained in both of them. Prove that

$$|\mathcal{F}| \leq 2^{\binom{2t}{2} - t}$$

4. (Han's Inequality.) Let $X = (X_1, \dots, X_m)$ be a random variable and let $H(X)$ denote its entropy. For a subset I of $\{1, 2, \dots, m\}$, let $X(I)$ denote the random variable $(X_i)_{i \in I}$. For $1 \leq q \leq m$, define

$$H_q(X) = \frac{1}{\binom{m-1}{q-1}} \sum_{Q \subset \{1, \dots, m\}, |Q|=q} H(X(Q))$$

Prove that

$$H_1(X) \geq H_2(X) \geq \dots \geq H_m(X) = H(X)$$

5. Let $X_i = \pm 1, 1 \leq i \leq n$, be uniform and independent and let $S_n = \sum_{i=1}^n X_i$. Let $0 \leq p \leq \frac{1}{2}$. Prove

$$\Pr[S_n \geq (1 - 2p)n] \leq 2^{H(p)n} 2^{-n}$$

by computing precisely the Chernoff bound $\min_{\lambda \geq 0} \mathbb{E}[e^{\lambda S_n}] e^{-\lambda(1-2p)n}$. (The case $p = 0$ will require a slight adjustment in the method though the end result is the same.)

6. (Parameter optimization in the Half Liar Game.) Find, asymptotically, the $u = u(q)$ that

minimizes $2^q \Pr[S_q \leq -u] + 2^{q+1}/(q-u)$ and express the minimal value in the form $2^{q+1}/q + (1+o(1))g(q)$ for some function g .

7. Show that for A fixed and r sufficiently large Paul wins the $(2^r - (r+1)A, A), r$ Chip Liar Game.

The Probabilistic Lens: an Extremal Graph

Let T (top) and B (bottom) be disjoint sets of size m and let G be a bipartite graph, all edges between T and B . Suppose G contains no 4-cycle. How many edges can G have? This is a question from Extremal Graph Theory. Surprisingly, for some m , we may give the precise answer.

Suppose $m = n^2 + n + 1$ and that a projective plane P of order n (and hence containing m points) exists. Identify T with the points of P and B with the lines of P and define $G = G_P$ by letting $t \in T$ be adjacent to $b \in B$ if and only if point t is on line b in P . As two points cannot lie on two lines, G_P contains no 4-cycle. We claim that such a G_P has the largest number of edges of any G containing no 4-cycle and further that any G containing no 4-cycle and having that many edges can be written in the form $G = G_P$.

Suppose G contains no 4-cycle. Let $b_1, b_2 \in B$ be a uniformly selected pair of distinct elements. For $t \in T$ let $D(t)$ be the set of $b \in B$ adjacent to t and $d(t) = |D(t)|$, the degree of t . Let I_t be the indicator random variable for t being adjacent to b_1, b_2 . Then

$$\mathbb{E}[I_t] = \Pr[b_1, b_2 \in D(t)] = \binom{d(t)}{2} / \binom{m}{2}.$$

Now set

$$X = \sum_{t \in T} I_t,$$

the number of $t \in T$ adjacent to b_1, b_2 . Then $X \leq 1$; that is, all b_1, b_2 have at most one common neighbor. ($X \leq 1$ is actually equivalent to G containing no 4-cycle.)

Linearity of expectation gives

$$\mathbb{E}[X] = \sum_{t \in T} \mathbb{E}[I_t] = \sum_{t \in T} \binom{d(t)}{2} / \binom{m}{2}.$$

Let $\hat{d} = m^{-1} \sum_{t \in T} d(t)$ be the average degree. Convexity of the function $\binom{y}{2}$ gives

$$\sum_{t \in T} \binom{d(t)}{2} / \binom{m}{2} \geq m \binom{\hat{d}}{2} / \binom{m}{2}$$

with equality if and only if all $t \in T$ have the same degree. Now

$$1 \geq \max X \geq \mathbb{E}[X] \geq m \binom{\hat{d}}{2} / \binom{m}{2}$$

When $G = G_P$ all $d(x) = \hat{d}$ (every line has $n+1$ points) and X is identically 1 (two points determine precisely one line) so that the above inequalities are all equalities and

$$1 = m \binom{\hat{d}}{2} / \binom{m}{2}.$$

Any graph with more edges would have a strictly larger $\hat{d}$ so that $1 \geq m \binom{\hat{d}}{2} / \binom{m}{2}$ would fail and the graph would contain a 4-cycle.

Suppose further G has the same number of edges as G_P and contains no 4-cycle. The inequalities then must be equalities and so $X = 1$ always. Define a geometry with points T and lines given by the neighbor sets of $b \in B$. As $X = 1$ any two points determine a unique line. Reversing the roles of T, B one also has that any two lines must determine a unique point. Thus G is generated from a projective plane.

18.7 16 Derandomization

Math is natural. Nobody could have invented the mathematical universe. It was there, waiting to be discovered, and it's crazy; it's bizarre.

- John Conway

As mentioned in Chapter 1, the probabilistic method supplies, in many cases, effective randomized algorithms for various algorithmic problems. In some cases, these algorithms can be derandomized and converted into deterministic ones. In this chapter we discuss some examples.

18.7.1 16.1 The Method of Conditional Probabilities

An easy application of the basic probabilistic method implies the following statement, which is a special case of Theorem 2.3.1.

Proposition 16.1.1 For every integer n there exists a coloring of the edges of the complete graph K_n by two colors so that the total number of monochromatic copies of K_4 is at most $\binom{n}{4} \cdot 2^{-5}$.

Indeed, $\binom{n}{4} \cdot 2^{-5}$ is the expected number of monochromatic copies of K_4 in a random 2-edge-coloring of K_n , and hence a coloring as above exists.

Can we actually find deterministically such a coloring in time which is polynomial in n ? Let us describe a procedure that does, which is a special case of a general technique called the method of conditional probabilities.

We first need to define a weight function for any partially colored K_n . Given a coloring of some of the edges of K_n by red and blue, we define, for each copy K of K_4 in K_n , a weight $w(K)$ as follows. If at least one edge of K is colored red and at least one edge is colored blue then $w(K) = 0$. If no edge of K is colored, then $w(K) = 2^{-5}$, and if $r \geq 1$ edges of K are colored, all with the same color, then $w(K) = 2^{r-6}$. Also define the total weight W of the partially colored K_n as the sum $\sum w(K)$, as K ranges over all copies of K_4 in K_n . Observe that the weight of each copy K of K_4 is precisely the probability that it will be monochromatic, if all the presently uncolored edges of K_n will be assigned randomly and independently one of the two colors red and blue. Hence, by linearity of expectation, the total weight W is simply the expected number of monochromatic copies of K_4 in such a random extension of the partial coloring of K_n to a full coloring.

We can now describe the procedure for finding a coloring as in Proposition 16.1.1. Order the $\binom{n}{2}$ edges of K_n arbitrarily, and construct the desired two-coloring by coloring each edge either red or blue in its turn. Suppose $e_1, \dots, e_{i-1}$ have already been colored, and we now have to color e_i . Let W be the weight of K_n , as defined above, with respect to the given partial coloring c of $e_1, \dots, e_{i-1}$. Similarly, let W_{red} be the weight of K_n with respect to the partial coloring obtained from c by coloring e_i red, and let W_{blue} be the weight of K_n with respect to the partial coloring obtained from c by coloring e_i blue. By the definition of W (and as follows from its interpretation as an expected value),

$$W = \frac{W_{\text{red}} + W_{\text{blue}}}{2}$$

The color of e_i is now chosen so as to minimize the resulting weight; that is, if $W_{\text{red}} \leq W_{\text{blue}}$ then we color e_i red, otherwise, we color it blue. By the above inequality, the weight function never increases during the algorithm. Since at the beginning its value is exactly $\binom{n}{4}2^{-5}$, its value at the end is at most this quantity. However, at the end all edges are colored, and the weight is precisely the number of monochromatic copies of K_4 . Thus the procedure above produces, deterministically and in polynomial time, a 2-edge-coloring of K_n satisfying the conclusion of Proposition 16.1.1.

Let us describe, now, the method of conditional probabilities in a more general setting. An instance of this method is due, implicitly, to Erdős and Selfridge (1973) and more explicit examples appear in Spencer (1987) and in Raghavan (1988). Suppose we have a probability space, and assume, for simplicity, that it is symmetric and contains 2^l points, denoted by the binary vectors of length l . Let $A_1, \dots, A_s$ be a collection of events and suppose that $\sum_{i=1}^s \Pr[A_i] = k$. Thus k is the expected value of the number of events A_i that hold, and hence there is a point $(\epsilon_1, \dots, \epsilon_l)$ in the space in which at most k events hold. Our objective is to find such a point deterministically.

For each choice of $(\epsilon_1, \dots, \epsilon_{j-1})$ and for each event A_i , the conditional probability

$$\Pr[A_i \mid \epsilon_1, \dots, \epsilon_{j-1}]$$

of the event A_i given the values of $\epsilon_1, \dots, \epsilon_{j-1}$ is clearly the average of the two conditional probabilities corresponding to the two possible choices for ϵ_j . That is,

$$\Pr[A_i \mid \epsilon_1, \dots, \epsilon_{j-1}] = \frac{\Pr[A_i \mid \epsilon_1, \dots, \epsilon_{j-1}, 0] + \Pr[A_i \mid \epsilon_1, \dots, \epsilon_{j-1}, 1]}{2}.$$

Consequently,

$$\begin{aligned} & \sum_{i=1}^s \Pr[A_i \mid \epsilon_1, \dots, \epsilon_{j-1}] \\ &= \frac{1}{2} \sum_{i=1}^s \Pr[A_i \mid \epsilon_1, \dots, \epsilon_{j-1}, 0] + \frac{1}{2} \sum_{i=1}^s \Pr[A_i \mid \epsilon_1, \dots, \epsilon_{j-1}, 1] \\ &\geq \min \left\{ \sum_{i=1}^s \Pr[A_i \mid \epsilon_1, \dots, \epsilon_{j-1}, 0], \sum_{i=1}^s \Pr[A_i \mid \epsilon_1, \dots, \epsilon_{j-1}, 1] \right\}. \end{aligned}$$

Therefore, if the values of ϵ_j are chosen, each one in its turn, so as to minimize the value of $\sum_{i=1}^s \Pr[A_i \mid \epsilon_1, \dots, \epsilon_j]$, then the value of this sum cannot increase. Since this sum is k at the beginning, it follows that it is at most k at the end. But at the end each ϵ_j is fixed, and hence the value of this sum is precisely the number of events A_i that hold at the point $(\epsilon_1, \dots, \epsilon_l)$, showing that our procedure works.

Note that the assumptions that the probability space is symmetric and that it has 2^l points can be relaxed. The procedure above is efficient provided l is not too large (as is usually the case in combinatorial examples), and, more importantly, provided the conditional probabilities $\Pr[A_i \mid \epsilon_1, \dots, \epsilon_j]$ can be computed efficiently for each of the events A_i and for each possible value of $\epsilon_1, \dots, \epsilon_j$. This is, indeed, the case in the example considered in Proposition 16.1.1. However, there are many interesting examples where this is not the case. A trick that can be useful in such cases is the introduction of pessimistic estimators, introduced by Raghavan (1988). Consider, again, the symmetric probability space with 2^l points described above, and the events $A_1, \dots, A_s$ in it. Suppose that for each event A_i and for each $0 \leq j \leq l$ we have a function $f_j^i(\epsilon_1, \dots, \epsilon_j)$ that can be efficiently computed. Assume also that

$$f_{j-1}^i(\epsilon_1, \dots, \epsilon_{j-1}) \geq \frac{f_j^i(\epsilon_1, \dots, \epsilon_{j-1}, 0) + f_j^i(\epsilon_1, \dots, \epsilon_{j-1}, 1)}{2}, \quad (16.1)$$

and that f_j^i is an upper bound on the conditional probabilities for the event A_i ; that is,

$$f_j^i(\epsilon_1, \dots, \epsilon_j) \geq \Pr[A_i \mid \epsilon_1, \dots, \epsilon_j]. \quad (16.2)$$

Clearly the same inequalities hold for the sums over i . In this case, if in the beginning $\sum_{i=1}^s f_0^i \leq t$, and we choose the values of the ϵ_j so as to minimize the sum $\sum_{i=1}^s f_j^i(\epsilon_1, \dots, \epsilon_j)$ in each step, we get in the end a point $(\epsilon_1, \dots, \epsilon_l)$ for which the sum $\sum_{i=1}^s f_l^i(\epsilon_1, \dots, \epsilon_l) \leq t$. The number of events A_i that hold in this point is at most t . The functions f_j^i in the argument above are called pessimistic estimators.

This enables us to obtain efficient algorithms in some cases where there is no known efficient way of computing the required conditional probabilities. The following theorem is an example; it is related to some of the results in Chapters 13 and 15.

Theorem 16.1.2 Let $(a_{ij})_{i,j=1}^n$ be an n by n matrix of reals, where $-1 \leq a_{ij} \leq 1$ for all i, j . Then one can find, in polynomial time, $\epsilon_1, \dots, \epsilon_n \in \{-1, 1\}$ such that for every i , $1 \leq i \leq n$,

$$\left| \sum_{j=1}^n \epsilon_j a_{ij} \right| \leq \sqrt{2n \ln(2n)}$$

Proof. Consider the symmetric probability space on the 2^n points corresponding to the 2^n possible vectors $(\epsilon_1, \dots, \epsilon_n) \in \{-1, 1\}^n$. Define $\beta = \sqrt{2n \ln(2n)}$ and let A_i be the event $\left| \sum_{j=1}^n \epsilon_j a_{ij} \right| > \beta$. We next show that the method of conditional probabilities with appropriate pessimistic estimators enables us to find efficiently a point of the space in which no event A_i holds.

Define $\alpha = \beta/n$ and let $G(x)$ be the function

$$G(x) = \cosh(\alpha x) = \frac{e^{\alpha x} + e^{-\alpha x}}{2}$$

By comparing the terms of the corresponding Taylor series it is easy to see that, for every real x ,

$$G(x) \leq e^{\alpha^2 x^2 / 2}$$

with strict inequality if neither x nor α are 0. It is also simple to check that for every real x and y ,

$$G(x)G(y) = \frac{G(x+y) + G(x-y)}{2}.$$

We can now define the functions f_p^i that will form our pessimistic estimators. For each $1 \leq i \leq n$ and for each $\epsilon_1, \dots, \epsilon_p \in \{-1, 1\}$ we define

$$f_p^i(\epsilon_1, \dots, \epsilon_p) = 2e^{-\alpha\beta} G\left(\sum_{j=1}^p \epsilon_j a_{ij}\right) \prod_{j=p+1}^n G(a_{ij}).$$

Obviously, these functions can be efficiently computed. It remains to check that they satisfy the conditions described in equations (16.1) and (16.2), and that the sum $\sum_{i=1}^n f_0^i$ is less than 1. This is proved in the following claims.

Claim 16.1.3 For every $1 \leq i \leq n$ and every $\epsilon_1, \dots, \epsilon_{p-1} \in \{-1, 1\}$,

$$f_{p-1}^i(\epsilon_1, \dots, \epsilon_{p-1}) = \frac{f_p^i(\epsilon_1, \dots, \epsilon_{p-1}, -1) + f_p^i(\epsilon_1, \dots, \epsilon_{p-1}, 1)}{2}.$$

Proof. Put $v = \sum_{j=1}^{p-1} \epsilon_j a_{ij}$. By the definition of f_p^i and by the properties of G ,

$$\begin{aligned}
 f_{p-1}^i(\epsilon_1, \dots, \epsilon_{p-1}) &= 2e^{-\alpha\beta} G(v) G(a_{ip}) \prod_{j=p+1}^n G(a_{ij}) \\
 &= 2e^{-\alpha\beta} \frac{G(v - a_{ip}) + G(v + a_{ip})}{2} \prod_{j=p+1}^n G(a_{ij}) \\
 &= \frac{f_p^i(\epsilon_1, \dots, \epsilon_{p-1}, -1) + f_p^i(\epsilon_1, \dots, \epsilon_{p-1}, 1)}{2}
 \end{aligned}$$

completing the proof of the claim.

Claim 16.1.4 For every $1 \leq i \leq n$ and every $\epsilon_1, \dots, \epsilon_{p-1} \in \{-1, 1\}$,

$$f_{p-1}^i(\epsilon_1, \dots, \epsilon_{p-1}) \geq \Pr[A_i \mid \epsilon_1, \dots, \epsilon_{p-1}].$$

Proof. Define v as in the proof of Claim 16.1.3. Then

$$\begin{aligned}
 &\Pr[A_i \mid \epsilon_1, \dots, \epsilon_{p-1}] \\
 &= \Pr\left[v + \sum_{j \geq p} \epsilon_j a_{ij} > \beta\right] + \Pr\left[-v - \sum_{j \geq p} \epsilon_j a_{ij} > \beta\right] \\
 &= \Pr\left[e^{\alpha(v + \sum_{j \geq p} \epsilon_j a_{ij})} > e^{\alpha\beta}\right] + \Pr\left[e^{-\alpha(v + \sum_{j \geq p} \epsilon_j a_{ij})} > e^{\alpha\beta}\right] \\
 &\leq e^{\alpha v} e^{-\alpha\beta} \mathbb{E}\left[\exp\left(\alpha \sum_{j \geq p} \epsilon_j a_{ij}\right)\right] + e^{-\alpha v} e^{-\alpha\beta} \mathbb{E}\left[\exp\left(-\alpha \sum_{j \geq p} \epsilon_j a_{ij}\right)\right] \\
 &= 2e^{-\alpha\beta} G(v) \prod_{j \geq p} G(a_{ij}) = f_{p-1}^i(\epsilon_1, \dots, \epsilon_{p-1})
 \end{aligned}$$

This completes the proof of Claim 16.1.4.

To establish the theorem it remains to show that $\sum_{i=1}^n f_0^i < 1$. Indeed, by the properties of G and by the choice of α and β :

$$\begin{aligned}
 \sum_{i=1}^n f_0^i &= \sum_{i=1}^n 2e^{-\alpha\beta} \prod_{j=1}^n G(a_{ij}) \leq \sum_{i=1}^n 2e^{-\alpha\beta} \prod_{j=1}^n e^{\alpha^2 a_{ij}^2 / 2} \\
 &\leq \sum_{i=1}^n 2e^{-\alpha\beta} e^{\alpha^2 n / 2} = 2ne^{\alpha^2 n / 2 - \alpha\beta} = 2ne^{-\alpha^2 n / 2} = 1.
 \end{aligned}$$

Moreover, the first inequality is strict unless $a_{ij} = 0$ for all i, j , whereas the second is strict unless $a_{ij}^2 = 1$ for all i, j . This completes the proof of the theorem.

18.7.2 16.2D-wise Independent Random Variables in Small Sample Spaces

The complexity class **NC** is, roughly speaking, the class of all problems that can be solved in time that is polylogarithmic (in the size of the input) using a polynomial number of parallel processors. Several models of computation, which are a theoretical abstraction of the parallel computer, have been used in considering this class. The most common one is the EREW (Exclusive Read, Exclusive Write) PRAM, in which different processors are not allowed to read from or write into the same memory cell simultaneously. See Karp and Ramachandran (1990) for more details.

Let n denote the size of the input. There are several simple tasks that can easily be performed in NC. For example, it is possible to copy the content of a cell c into $m = n^{O(1)}$ cells in

time $O(\log n)$, using, say, m processors. To do so, consider a complete binary tree with m leaves and associate each of its internal vertices with a processor. At first, the processor corresponding to the root of the tree reads from c and writes its content in two cells, corresponding to its two children. Next, each of these two, in parallel, reads from its cell and writes its content in two cells corresponding to its two children. In general, at the i th step all the processors whose distance from the root of the tree is $i - 1$, in parallel, read the content of c previously stored in their cells and write it twice. The procedure clearly ends in time $O(\log m)$, as claimed. [In fact, it can be shown that $O(m/\log m)$ processors suffice for this task but we do not try to optimize this number here.]

A similar technique can be used for computing the sum of m numbers with m processors in time $O(\log m)$. We consider the numbers as if they lie on the leaves of a complete binary tree with m leaves, and in the i th step each one of the processors whose distance from the leaves is i computes, in parallel, the sum of the two numbers previously computed by its children. The root will clearly have, in such a way, the desired sum in time $O(\log m)$.

Let us now return to the edge-coloring problem of the complete graph K_n discussed in Proposition 16.1.1. By the remarks above, the problem of checking if in a given edge-coloring there are at most $\binom{n}{4}2^{-5}$ monochromatic copies of K_4 is in NC; that is, this checking can be done in time $(\log n)^{O(1)}$ - in fact, in time $O(\log n)$ - using $n^{O(1)}$ processors. Indeed, we can first copy the given coloring $\binom{n}{4}$ times. Then we assign a processor for each copy of K_4 in K_n , and this processor checks if its copy is monochromatic or not (all these checkings can be done in parallel, since we have enough copies of the coloring). Finally, we sum the number of processors whose copies are monochromatic. Clearly we can complete the work in time $O(\log n)$ using $n^{O(1)}$ parallel processors.

Thus we can check, in NC, if a given coloring of K_n satisfies the assertion of Proposition 16.1.1. Can we find such a coloring deterministically in NC? The method described in the previous section does not suffice, as the edges have been colored one by one, so the procedure is sequential and requires time $\Omega(n^2)$. However, it turns out that in fact we can find, in NC, a coloring with the desired properties by applying a method that relies on a technique first suggested by Joffe (1974), and later

developed by many researchers. This method is a general technique for converting randomized algorithms whose analysis only depends on d -wise rather than fully independent random choices (for some constant d) into deterministic (and in many cases also parallel) ones. Our approach here follows the one of Alon, Babai and Itai (1986), but for simplicity we only consider here the case of random variables that take the two values 0,1 with equal probability.

The basic idea is to replace an exponentially large sample space by one of polynomial size. If a random variable on such a space takes a certain value with positive probability, then we can find a point in the sample space in which this happens simply by deterministically checking all the points. This can be done with no loss of time by using a polynomial number of parallel processors. Note that for the edge-coloring problem considered in Proposition 16.1.1, 6-wise independence of the random variables corresponding to the colors of the edges suffice, since this already gives a probability of 2^{-5} for each copy of K_4 to be monochromatic, and hence gives the required expected value of monochromatic copies. Therefore, for this specific example, it suffices to construct a sample space of size $n^{O(1)}$ and $\binom{n}{2}$ random variables in it, each taking the values 0 and 1 with probability $\frac{1}{2}$, such that each 6 of the random variables are independent.

Small sample spaces with many d -wise independent 0,1 random variables in them can be constructed from any linear error correcting code with appropriate parameters. The construction we describe here is based on the binary BCH codes [see, e.g., MacWilliams and Sloane (1977)].

Theorem 16.2.1 Suppose $n = 2^k - 1$ and $d = 2t + 1$. Then there exists a symmetric probability space Ω of size $2(n + 1)^t$ and d -wise independent random variables $y_1, \dots, y_n$ over Ω

each of which takes the values 0 and 1 with probability $\frac{1}{2}$.

The space and the variables are explicitly constructed, given a representation of the field $F = GF(2^k)$ as a k -dimensional algebra over $GF(2)$.

Proof. Let $x_1, \dots, x_n$ be the n nonzero elements of F , represented as column vectors of length k over $GF(2)$. Let H be the following $1 + kt$ by n matrix over $GF(2)$.

$$H = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ x_1 & x_2 & \cdots & x_n \\ x_1^3 & x_2^3 & \cdots & x_n^3 \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{2t-1} & x_2^{2t-1} & \cdots & x_n^{2t-1} \end{pmatrix}$$

This is the parity check matrix of the extended binary BCH code of length n and designed distance $2t + 2$. It is well known that any $d = 2t + 1$ columns of H are linearly independent over $GF(2)$. For completeness, we present the proof in the next lemma.

Lemma 16.2.2 Any set of $d = 2t + 1$ columns of H is linearly independent over $GF(2)$.

Proof. Let $J \subset \{1, 2, \dots, n\}$ be a subset of cardinality $|J| = 2t + 1$ of the set of indices of the columns of H . Suppose that $\sum_{j \in J} z_j H_j = 0$, where H_j denotes the j th column of H and $z_j \in GF(2)$. To complete the proof we must show that $z_j = 0$ for all $j \in J$. By the assumption,

$$\sum_{j \in J} z_j x_j^i = 0 \quad (16.3)$$

for $i = 0$ and for every odd i satisfying $1 \leq i \leq 2t - 1$. Suppose, now, that $a = 2^b \cdot l$, where $l \leq 2t - 1$ is an odd number. By squaring the equation (16.3) b times, where $i = l$, using the fact that $(u + v)^2 = u^2 + v^2 \pmod{2}$ and the fact that since each z_j is either 0 or 1, the equality $z_j = z_j^2$ holds for all j , we conclude that equation (16.3) holds for $i = a$. Consequently, (16.3) holds for all $i, 0 \leq i \leq 2t$. This is a homogeneous system of $2t + 1$ linear equations in $2t + 1$ variables. The matrix of the coefficients is a Vandermonde matrix, which is nonsingular. Thus the only solution is the trivial one $z_j = 0$ for all $j \in J$, completing the proof of the lemma.

Returning to the proof of the theorem, we define $\Omega = \{1, 2, \dots, 2(n + 1)^t\}$, and let $A = (a_{ij}), i \in \Omega, 1 \leq j \leq n$ be the $(0, 1)$ -matrix whose $2(n + 1)^t = 2^{kt+1}$ rows are all the linear combinations (over $GF(2)$) of the rows of H . The sample space Ω is now endowed with the uniform probability measure, and the random variable y_j is defined by the formula $y_j(i) = a_{ij}$ for all $i \in \Omega, 1 \leq j \leq n$.

It remains to show that the variables y_j are d -wise independent, and that each of them takes the values 0 and 1 with equal probability. For this we have to show that for every set J of up to d columns of A , the rows of the $|\Omega|$ by $|J|$ submatrix $A_J = (a_{ij}), i \in \Omega, j \in J$ take on each of the $2^{|J|}$ $(0, 1)$ -vectors of length $|J|$ equally often. However, by Lemma 16.2.2 the columns of the corresponding submatrix H_J of H are linearly independent. The number of rows of A_J that are equal to any given vector is precisely the number of linear combinations of the rows of H_J that are equal to this vector. This number is the number of solutions of a system of $|J|$ linearly independent linear equations in $kt + 1$ variables, which is, of course, $2^{kt+1-|J|}$, independent of the vector of free coefficients. This completes the proof of the theorem.

Theorem 16.2.1 supplies an efficient way of constructing, for every fixed d and every n , a sample space of size $O(n^{\lceil d/2 \rceil})$ and nd -wise independent random variables in it, each taking the values 0 and 1 with equal probability. In particular, we can use such a space of size $O\left(\binom{n}{2}^3\right) = O(n^6)$ for finding a coloring as in Proposition 16.1.1 in NC. Several other applications of Theorem 16.2.1 appear in the paper of Alon et al. (1986).

It is natural to ask if the size $O(n^{\lceil d/2 \rceil})$ can be improved. We next show that this size is optimal, up to a constant factor (depending on d).

Let us call a random variable almost constant if it attains a single value with probability 1. Let $m(n, d)$ denote the function defined by

$$m(n, d) = \sum_{j=0}^{d/2} \binom{n}{j} \text{ if } d \text{ is even}$$

and

$$m(n, d) = \sum_{j=0}^{(d-1)/2} \binom{n}{j} + \binom{n-1}{(d-1)/2} \text{ if } d \text{ is odd.}$$

Observe that for every fixed d , $m(n, d) = \Omega(n^{\lfloor d/2 \rfloor})$.

Proposition 16.2.3 If the random variables $y_1, \dots, y_n$ over the sample space Ω are d -wise independent and none of them is almost constant then $|\Omega| \geq m(n, d)$.

Note that we assume here neither that Ω is a symmetric space nor that the variables y_j are $(0, 1)$ -variables.

Proof. Clearly we may assume that the expected value of each y_j is 0 [since otherwise we can replace y_j by $y_j - E[y_j]$.] For each subset S of $\{1, \dots, n\}$, define $\alpha_S = \prod_{j \in S} y_j$. Observe that since no y_j is almost constant and since the variables are d -wise independent,

$$E[\alpha_S \alpha_S] = \prod_{j \in S} \text{Var}[y_j] > 0 \quad (16.4)$$

for all S satisfying $|S| \leq d$. Similarly, for all S and T satisfying $|S \cup T| \leq d$ and $S \neq T$ we have

$$E[\alpha_S \alpha_T] = \prod_{j \in S \cap T} \text{Var}[y_j] \prod_{j \in S \cup T \setminus (S \cap T)} E[y_j] = 0. \quad (16.5)$$

Let $S_1, \dots, S_m$, where $m = m(n, d)$, be subsets of $\{1, \dots, n\}$ such that the union of each two is of size at most d . [Take all subsets of size at most $d/2$, and if d is odd add all the subsets of size $(d+1)/2$ containing 1.]

To complete the proof, we show that the m functions α_{S_j} (considered as real vectors of length $|\Omega|$) are linearly independent. This implies that $|\Omega| \geq m = m(n, d)$, as stated in the proposition.

To prove linear independence, suppose $\sum_{j=1}^m c_j \alpha_{S_j} = 0$. Multiplying by α_{S_i} and computing expected values we obtain, by (16.5),

$$0 = \sum_{j=1}^m c_j E[\alpha_{S_j} \alpha_{S_i}] = c_i E[\alpha_{S_i} \alpha_{S_i}]$$

This implies, by (16.4), that $c_i = 0$ for all i . The required linear independence follows, completing the proof.

The last proposition shows that the size of a sample space with nd -wise independent non-trivial random variables can be polynomial in n only when d is fixed. However, as shown by Naor and Naor (1990), if we only require the random variables to be almost d -wise independent, the size can be polynomial even when $d = \Omega(\log n)$. Such sample spaces and random variables, which can be constructed explicitly in several ways, have various interesting applications in which almost d -wise independence suffices. More details appear in Naor and Naor (1990) and in Alon, Goldreich, Håstad and Peralta (1990).

18.7.3 16.3 Exercises

1. Let $A_1, \dots, A_n \subseteq \{1, \dots, m\}$ with $\sum_{i=1}^n 2^{1-|A_i|} < 1$. Prove there exists a two-coloring $\chi : \{1, \dots, m\} \rightarrow \{0, 1\}$ with no A_i monochromatic. With $m = n$ give a deterministic algorithm to find such a χ in polynomial time.
2. Describe a deterministic algorithm that, given n , constructs, in time polynomial in n , a family $\mathcal{F}$ of n^{10} subsets of the set $N = \{1, 2, \dots, n\}$, where each $F \in \mathcal{F}$ is of size at most $10 \log n$ and for every family $\mathcal{G}$ of n subsets each of cardinality $n/2$ of N , there is an $F \in \mathcal{F}$ that intersects all members of $\mathcal{G}$.

The Probabilistic Lens: Crossing Numbers, Incidences, Sums and Products

In this lens we start with a simple result in graph theory, whose proof is probabilistic, and then describe some of its fascinating consequences in combinatorial geometry and combinatorial number theory. Some versions of most of these seemingly unrelated consequences have been proved before, in a far more complicated manner. Before the discovery of the new proofs shown here, the only clue that there might be a connection between all of them has been the fact that Endre Szemerédi is one of the coauthors of each of the papers providing the first proofs.

An embedding of a graph $G = (V, E)$ in the plane is a planar representation of it, where each vertex is represented by a point in the plane, and each edge uv is represented by a curve connecting the points corresponding to the vertices u and v . The crossing number of such an embedding is the number of pairs of intersecting curves that correspond to pairs of edges with no common endpoints. The crossing number $\text{cr}(G)$ of G is the minimum possible crossing number in an embedding of it in the plane. The following theorem was proved by Ajtai, Chvátal, Newborn and Szemerédi (1982) and, independently, by Leighton. Here we describe a very short probabilistic proof.

Theorem 1 The crossing number of any simple graph $G = (V, E)$ with $|E| \geq 4|V|$ is at least $|E|^3/64|V|^2$.

Proof. By Euler's formula any simple planar graph with $n \geq 3$ vertices has at most $3n - 6$ edges, implying that any simple planar graph with n vertices has at most $3n$ edges. Therefore the crossing number of any simple graph with n vertices and m edges is at least $m - 3n$. Let $G = (V, E)$ be a graph with $|E| \geq 4|V|$ embedded in the plane with $t = \text{cr}(G)$ crossings. Let H be the random induced subgraph of G obtained by picking each vertex of G , randomly and independently, to be a vertex of H with probability p (where p will be chosen later). The expected number of vertices of H is $p|V|$, the expected number of its edges is $p^2|E|$, and the expected number of crossings in its given embedding is p^4t , implying that the expected value of its crossing number is at most p^4t . Therefore $p^4t \geq p^2|E| - 3p|V|$, implying that

$$\text{cr}(G) = t \geq \frac{|E|}{p^2} - 3\frac{|V|}{p^3}.$$

Without trying to optimize the constant factor, substitute $p = 4|V|/|E| (\leq 1)$, to get the desired result.

Székely (1997) noticed that this result can be applied to obtain a surprisingly simple proof of a result of Szemerédi and Trotter in combinatorial geometry. The original proof is far more complicated.

Theorem 2 Let P be a set of n distinct points in the plane, and let L be a set of m distinct lines. Then the number of incidences between the members of P and those of L (i.e., the number of pairs (p, l) with $p \in P, l \in L$ and $p \in l$) is at most $c(m^{2/3}n^{2/3} + m + n)$, for some absolute constant c .

Proof. We may and will assume that every line in L is incident with at least one of the points of P . Denote the number of incidences by I . Let $G = (V, E)$ be the graph whose vertices are all members of P , where two are adjacent if and only if they are consecutive points of P on some line in L . Clearly $|V| = n$ and $|E| = I - m$. Note that G is already given embedded in the plane, where the edges are represented by segments of the corresponding lines in L . In this embedding, every crossing is an intersection point of two members of L , implying that $\text{cr}(G) \leq \binom{m}{2} \leq \frac{1}{2}m^2$. By Theorem 1, either $I - m = |E| < 4|V| = 4n$, that is, $I \leq m + 4n$, or

$$\frac{m^2}{2} \geq \text{cr}(G) \geq \frac{(I - m)^3}{64n^2}$$

implying that $I \leq (32)^{1/3}m^{2/3}n^{2/3} + m$. In both cases $I \leq 4(m^{2/3}n^{2/3} + m + n)$, completing the proof.

An analogous argument shows that the maximum possible number of incidences between a set of n points and a set of m unit circles in the plane does not exceed $O(m^{2/3}n^{2/3} + m + n)$ and this implies that the number of unit distances determined by a set of n points in the plane is at most $O(n^{4/3})$. While the above upper bound for the number of incidences of points and lines is sharp, up to a constant factor, an old conjecture of Erdős asserts that the maximum possible number of unit distances determined by a set of n points in the plane is at most $c_\epsilon n^{1+\epsilon}$ for any $\epsilon > 0$. The

$O(n^{4/3})$ estimate is, however, the best known upper bound and was first proved by Spencer, Szemerédi and Trotter in a far more complicated way.

Elekes (1997) found several applications of Theorem 2 in additive number theory. Here, too, the proofs are amazingly simple. Here is a representative result.

Theorem 3 For any three sets A, B and C of s real numbers each,

$$|A \cdot B + C| = |\{ab + c : a \in A, b \in B, c \in C\}| \geq \Omega(s^{3/2}).$$

Proof. Put $R = A \cdot B + C$, $|R| = r$ and define

$$P = \{(a, t) : a \in A, t \in R\}, \quad L = \{y = bx + c : b \in B, c \in C\}$$

Thus P is a set of $n = sr$ points in the plane, L is a set of $m = s^2$ lines in the plane, and each line $y = bx + c$ in L is incident with s points of P , that is, with all the points $\{(a, ab + c) : a \in A\}$. Therefore, by Theorem 2, $s^3 \leq 4(s^{4/3}(sr)^{2/3} + sr + s^2)$, implying that $r \geq \Omega(s^{3/2})$, as needed.

The same method implies that for every set A of n reals, either $|A + A| \geq \Omega(n^{5/4})$ or $|A \cdot A| \geq n^{5/4}$, greatly improving and simplifying a result of Erdős and Szemerédi.

18.8 17 Graph Property Testing

'Call the first witness,' said the King; and the White Rabbit blew three blasts on the trumpet, and called out, 'First witness!'

- from Alice in Wonderland, by Lewis Carroll

18.8.1 17.1 Property Testing

Property testers are fast randomized algorithms for distinguishing between combinatorial structures that satisfy a certain property, and ones that are far from satisfying it. The basic algorithmic task in this area is to design a randomized algorithm that, given a combinatorial structure S , can distinguish with high probability between the case that S satisfies a prescribed property $\mathcal{P}$ and the case that S is ϵ -far from satisfying $\mathcal{P}$. Here S is said to be ϵ -far from

satisfying $\mathcal{P}$ if an ε -fraction of its representation should be modified in order to turn it to a structure that satisfies $\mathcal{P}$. The main objective is to design randomized algorithms, which look at a very small portion of the input, and using this information distinguish with high probability between the above two cases. Such algorithms are called testers for the property $\mathcal{P}$.

Preferably, a tester should look at a portion of the input whose size is a function of ε only. The general notion of property testing was first formulated by Rubinfeld and Sudan (1996), who were motivated by the study of various algebraic properties such as linearity of functions. Property testing is also motivated by questions in program checking, computational learning, approximation algorithms and probabilistically checkable proofs, as well as by the need to access large data sets, like the graph of the Internet. The investigation of the subject relies heavily on probabilistic methods.

The main focus of this chapter is in testing properties of graphs. In this case a graph G on n vertices is said to be ε -far from satisfying a property $\mathcal{P}$, if one needs to add to or delete from G at least εn^2 edges in order to turn it into a graph satisfying $\mathcal{P}$. Here we assume that the tester can query an oracle whether a pair of vertices, i and j , are adjacent in the input graph G . If the graph satisfies the property, then the tester has to accept with probability at least, say, $\frac{2}{3}$, and if it is ε -far from satisfying it, then the algorithm has to reject with probability at least $\frac{2}{3}$.

The study of the notion of testability for combinatorial structures, and mainly for labeled graphs, was introduced by Goldreich, Goldwasser and Ron (1998). They showed that many natural graph properties such as k -colorability, having a large clique or having a large cut, admit a tester, whose query complexity [i.e., the number of oracle queries of type "does (i, j) belong to $E(G)$ "] as well as their total running time can be upper bounded by a function of ε that is independent of the size of the input. We call properties having such efficient testers, that is, testers whose query complexity is a function of ε only, testable. In general, a property tester may have a small probability of accepting graphs that are ε -far from satisfying the tested property, as well as a small probability of rejecting graphs satisfying the property. In this case the tester is said to have two-sided error. If the tester accepts graphs satisfying the property with probability 1, then the tester is said to have one-sided error.

It is worth noting that the model of graph property testing described here is often referred to as the dense graph model. Other models of graph property testing have also been investigated; see, for example, Goldreich and Ron (2002). For further reading and pointers on testing properties of graphs and other combinatorial structures the reader is referred to the surveys by Goldreich (1999), Fischer (2001), Ron (2001), Alon and Shapira (2006) and their references.

18.8.2 17.2 Testing Colorability

Although the computational problem of deciding whether a given graph is k -colorable is NP-complete for every fixed $k \geq 3$, it turns out that, somewhat surprisingly, for every fixed $\varepsilon > 0$ there is an efficient algorithm for distinguishing between graphs on n vertices that are k -colorable, and graphs from which one has to delete at least εn^2 edges to make them k -colorable. This result, mentioned already in Alon, Duke, Lefmann, Rödl and Yuster (1994), follows from the fact that the property of being k -colorable is testable, as proved implicitly in Rödl and Duke (1985) and explicitly (with a far better dependence on the parameter ε) in Goldreich et al. (1998). Indeed, as we show in this subsection, if a graph $G = (V, E)$ is ε -far from being k -colorable, then an induced subgraph of it on a randomly chosen set of $c(k)/\varepsilon^2$ vertices is not k -colorable with high probability. This is proved in Alon and Krivelevich (2002), with $c(k) = 36k \ln k$, building on the work of Goldreich et al. (1998), who showed that a random set of $O(k^2 \ln k / \varepsilon^3)$ vertices suffices. Note that the above supplies a very simple tester with one-sided error for testing k -colorability; consider the induced subgraph

on a randomly chosen set of $36k \ln k / \varepsilon^2$ vertices, and accept iff this subgraph is k -colorable. Obviously, every k -colorable graph is accepted by this procedure, and graphs that are ε -far from being k -colorable are likely to be rejected. Note also that the validity of this statement implies the nontrivial fact that every graph that is ε -far from being k -colorable contains a small witness (for being non k colorable), that is, a subgraph on only $c(\varepsilon, k) \leq O(k \ln k / \varepsilon^2)$ vertices which is not k -colorable. The existence of some such function $c(\varepsilon, k)$ has been conjectured by Erdős and first proved by Rödl and Duke [for some extremely fast growing function $c(\varepsilon, k)$ of ε and k - see Rödl and Duke (1985).] In this section we describe the improved $c(k)/\varepsilon^2$ bound. For simplicity, we present the proof only for $k = 3$; the proof for the general case is essentially identical. Throughout the proof we omit all floor and ceiling signs whenever these are not crucial.

Theorem 17.2.1 Suppose $0 < \varepsilon < \frac{1}{10}$, let $G = (V, E)$ be a graph on $n > 400/\varepsilon^3$ vertices, and suppose that one has to delete from G at least εn^2 edges to make it 3-colorable. Then the probability that an induced subgraph of G on a randomly chosen set of $s = 40/\varepsilon^2$ vertices is 3-colorable does not exceed $\frac{1}{10}$.

We start with an outline of the proof. Given $G = (V, E)$ as in the theorem, pick a random subset $R \subset V$ of size $|R| = s = 40/\varepsilon^2$ in s rounds, each time choosing uniformly at random a single vertex r_j among the vertices not selected so far.

Suppose that some subset $S \subset R$ has already been 3-colored by $\phi : S \rightarrow C$, where $C = \{1, 2, 3\}$. The objective is to show that with high probability there is a witness showing that this partial coloring cannot be extended to a proper coloring of the induced subgraph on R . If a proper 3-coloring $c : V \rightarrow C$ of G is to coincide with ϕ on S , then for every vertex $v \in V \setminus S$, the colors of the neighbors of v in S under ϕ are forbidden for v in c . The rest of the colors are still feasible for v . It could be that v has no feasible colors left at all. Such a vertex will be called colorless with respect to S and ϕ . If the number of colorless vertices is large, then there is a decent chance that among the next few randomly chosen vertices of R there will be one such colorless vertex v^* . Obviously, adding v^* to S provides the desired witness for non extendibility of ϕ .

If the set of colorless vertices is small, then one can show that, as G is ε -far from being 3-colorable, there is a relatively large subset W of vertices (which will be called restricting) such that adding any vertex $v \in W$ to S and coloring it by any feasible color excludes this color from the lists of feasible colors of at least εn neighbors of v . If such a vertex v is found among the next few vertices of the random sample R , then adding v to S and coloring it by any of its feasible colors reduces substantially the total size of the lists of feasible colors for the remaining vertices of V , which helps to approach the first situation, that is, the case when there are many colorless vertices. This process can be represented by a tree in which every internal node corresponds to a restricting vertex v , and every edge from v to a child corresponds to a feasible color for v . The tree will not be very large. Indeed, each of its internal vertices has at most 3 children, and its depth cannot exceed $3/\varepsilon$, as the total size of the lists of feasible colors at the beginning is $3n$, and this size is reduced by at least εn in each step. It thus suffices to show that with high probability the construction of the whole tree (until no feasible colors are left to any of its leaves) can be completed using the vertices in our random set R .

We proceed with the formal proof. For a subset $S \subseteq V$, a 3-coloring of it $\phi : S \rightarrow C$, and a vertex $v \in V \setminus S$ let $L_\phi(v)$ be the set of all colors in C besides those that appear already on some neighbor of v . This is the set of feasible colors for v . Clearly, for $S = \emptyset$, $L_\phi(v) = C$ for every $v \in V$. A vertex $v \in V \setminus S$ is called colorless if $L_\phi(v) = \emptyset$. Let U denote the set of all colorless vertices under (S, ϕ) .

For every vertex $v \in V \setminus (S \cup U)$ define

$$\delta_\phi(v) = \min_{i \in L_\phi(v)} |\{u \in N(v) \setminus (S \cup U) : i \in L(u)\}|.$$

Therefore coloring v by any one of the colors from $L_\phi(v)$ and then adding it to S will result in deleting this color and thus shortening the lists of feasible colors of at least $\delta_\phi(v)$ neighbors of v outside S .

Claim 17.2.2 For every set $S \subset V$ and every 3-coloring ϕ of S , the graph G is at most

$$(n-1)|S \cup U| + \frac{1}{2} \sum_{v \in V \setminus (S \cup U)} \delta_\phi(v)$$

edges far from being k -colorable.

Proof. Consider the following coloring of G : Every $v \in S$ is colored by $\phi(v)$, every $v \in U$ is colored by an arbitrary color and every $v \in V \setminus (S \cup U)$ is colored by a color $i \in L_\phi(v)$ for which $\delta_\phi(v) = |\{u \in N(v) \setminus (S \cup U) : i \in L(u)\}|$. The number of monochromatic edges incident with $S \cup U$ is at most $(n-1)|S \cup U|$. Every vertex $v \in V \setminus (S \cup U)$ has exactly $\delta_\phi(v)$ neighbors $u \in V \setminus (S \cup U)$, whose color list $L_\phi(v)$ contains the color chosen for v . Therefore v will have at most $\delta_\phi(v)$ neighbors in $V \setminus (S \cup U)$ colored in the same color as v itself. Hence the total number of monochromatic edges is at most

$$(n-1)|S \cup U| + \frac{1}{2} \sum_{v \in V \setminus (S \cup U)} \delta_\phi(v),$$

as claimed.

Given a pair (S, ϕ) , a vertex $v \in V \setminus (S \cup U)$ is called restricting if $\delta_\phi(v) \geq \varepsilon n$. We denote by W the set of all restricting vertices.

Claim 17.2.3 For every pair (S, ϕ) , where $S \subset V$ and $\phi : S \rightarrow C$,

$$|U \cup S \cup W| \geq \varepsilon n / 2.$$

Proof. By the previous claim, and since G is ε -far from being 3-colorable,

$$\begin{aligned} \varepsilon n^2 &< n(|S| + |U|) + \frac{1}{2} \sum_{v \in V \setminus (S \cup U)} \delta_\phi(v) \\ &\leq n(|S| + |U|) + \frac{1}{2}|W|(n-1) + \frac{1}{2} \sum_{v \in V \setminus (S \cup U \cup W)} \delta_\phi(v) \\ &< n(|S| + |U| + |W|) + \frac{1}{2}\varepsilon n^2. \end{aligned}$$

□

Returning to our randomly chosen vertices $r_1, \dots, r_s$ of R , construct an auxiliary ternary tree T . To distinguish between the vertices of G and those of T we call the latter nodes. Each node of T is labeled either by a vertex of G or by the special symbol $\#$, whose meaning will be explained in what follows. If a node t of T is labeled by $\#$, then t is called a terminal node. The edges of T are labeled by integers from C .

Let t be a node of T . Consider the path from the root of T to t , not including t itself. The labels of the nodes along this path form a subset $S(t)$ of V . The labels of the edges along the path define a 3-coloring $\phi(t)$ of $S(t)$ in a natural way: the label of the edge following a node t' in the path determines the color of its label $v(t')$. The labeling of the nodes and edges of T will have the following property: If t is labeled by v and v has a neighbor in $S(t)$ whose color in $\phi(t)$ is i , then the son of v along the edge labeled by i is labeled by $\#$. This label indicates the fact that in this case color i is infeasible for v , given $(S(t), \phi(t))$.

At each step of the construction of T we will maintain the following: All leafs of T are either unlabeled or are labeled by $\#$. Also, only leafs of T can be labeled by $\#$. We start the construction of T from an unlabeled single node, the root of T .

Suppose that $j - 1$ vertices of T have already been chosen, and we are about to choose vertex r_j of R . Consider a leaf t of T . If t is labeled by $\#$, we do nothing for this leaf. (That is the reason such a t is called a terminal node; nothing will ever grow out of it.) Assume now that t is unlabeled. Define the pair $(S(t), \phi(t))$ as described above. Now, for the pair $(S(t), \phi(t))$ we define the set $U(t)$ of colorless vertices and the set $W(t)$ of restricting vertices as described before. Round j is called successful for the node t if the random vertex r_j satisfies $r_j \in U(t) \cup W(t)$. If round j is indeed successful for t , then we label t by r_j , create 3 sons of t and label the corresponding edges by 1, 2, 3. Now, if color i is infeasible for r_j , given $(S(t), \phi(t))$, we label the son of t along the edge with label i by $\#$, otherwise we leave this son unlabeled. Note that if $r_j \in U(t)$, then none of the colors from C is feasible for r_j , and thus all the sons of t will be labeled by $\#$. This completes the description of the process of constructing T . As each edge along a path from a root to a leaf of the tree corresponds to a restricting vertex, and the total size of all lists starts with $3n$ and is reduced by at least εn with each coloring of a restricting vertex, we have the following.

Claim 17.2.4 The depth of T is at most $3/\varepsilon$.

Our construction also implies that if a leaf t^* of T is labeled by $\#$, then $\phi(t^*)$ is not a proper 3-coloring of $S(t^*)$. We thus have the following.

Claim 17.2.5 If after round j all leafs of the tree T are terminal nodes, then the induced subgraph of G on $\{r_1, \dots, r_j\}$ is not 3-colorable.

To complete the proof it thus suffices to show the following.

Claim 17.2.6 After $s = 40/\varepsilon^2$ rounds, with probability at least $\frac{9}{10}$ all leaves of T are terminal nodes.

Proof. As every non leaf node of T has at most 3 sons and by Claim 17.2.4 the depth of T is at most $3/\varepsilon$, it can be embedded naturally in the ternary tree $T_{3,3/\varepsilon}$ of depth $3/\varepsilon$. Moreover, this embedding can be prefixed even before exposing R and T . Note that the number of vertices of $T_{3,3/\varepsilon}$ is $1 + 3 + \dots + 3^{3/\varepsilon} < 3^{1+3/\varepsilon}$.

Recall that during the construction of the random sample R and the tree T , a successful round for a leaf t of T results in creating 3 sons of T . Fix a node t of $T_{3,3/\varepsilon}$. If after $40/\varepsilon^2$ rounds t is a leaf of T , then the total number of successful rounds for the path from the root of T to t is equal to the depth of t . As $S(t) \subseteq R$ and thus $|S(t)| \leq 40/\varepsilon^2 \leq \varepsilon n/10$, by Claim 17.2.3 each round has probability of success at least 0.4ε . Therefore, the probability that t is a non terminal leaf of T after $40/\varepsilon^2$ steps can be bounded from above by the probability that the binomial random variable $B(40/\varepsilon^2, 0.4\varepsilon)$ is at most $3/\varepsilon$. The latter probability is at most

$$\exp\left(-\frac{(16/\varepsilon - 3/\varepsilon)^2}{2 \cdot 16/\varepsilon}\right) = \exp\left(-\frac{169}{32\varepsilon}\right).$$

Thus by the union bound we conclude that the probability that some node of $T(3, 3/\varepsilon)$ is a leaf of T , not labeled by $\#$, is at most

$$|V(T_{3,3/\varepsilon})| \exp\left(-\frac{169}{32\varepsilon}\right) < \frac{1}{10}$$

The assertion of Theorem 17.2.1 follows from Claims 17.2.5 and 17.2.6.

18.8.3 17.3 Szemerédi's Regularity Lemma

In this section we describe a fundamental result, the Regularity Lemma, proved by Endre Szemerédi in the 1970s. The original motivation for proving it was an application in combinato-

rial number theory, leading, together with several additional deep ideas, to a complete solution of the Erdős-Turán conjecture discussed in Appendix B.2: Every set of integers of positive upper density contains arbitrarily long arithmetic progressions. It took some time to realize that the lemma is an extremely powerful

tool in extremal graph theory, combinatorics and theoretical computer science. Stated informally, the Regularity Lemma asserts that the vertices of every large graph can be decomposed into a finite number of parts, so that the edges between almost every pair of parts form a random-looking graph. The power of the lemma is in the fact it deals with an arbitrary graph, making no assumptions, and yet it supplies much useful information about its structure. It should be stressed that the impact of the Regularity Lemma goes far beyond its applications in property testing, which is our focus in this chapter. A detailed survey of the lemma and some of its many variants and fascinating consequences can be found in Komlós and Simonovits (1996).

Let $G = (V, E)$ be a graph. For two disjoint nonempty subsets of vertices $A, B \subset V$, let $e(A, B)$ denote the number of edges of G with one end in A and one in B , and let

$$d(A, B) = \frac{e(A, B)}{|A||B|}$$

denote the density of the pair (A, B) . For a real $\varepsilon > 0$, a pair (A, B) as above is called ε -regular if for every $X \subset A$ and $Y \subset B$ that satisfy $|X| \geq \varepsilon|A|, |Y| \geq \varepsilon|B|$ the inequality $|d(A, B) - d(X, Y)| \leq \varepsilon$ holds. It is not difficult to see that for every fixed positive ε, p a fixed pair of two sufficiently large disjoint subsets A and B of a random graph $G = G(n, p)$ are very likely to be ε -regular of density roughly p . (This is stated in one of the exercises at the end of the chapter.) Conversely, an ε -regular pair A, B with a sufficiently small positive ε is random-looking in the sense that it shares many properties satisfied by random (bipartite) graphs.

A partition $V = V_0 \cup V_1 \cup \dots \cup V_k$ of V into pairwise disjoint sets in which V_0 is called the exceptional set is an equipartition if $|V_1| = |V_2| = \dots = |V_k|$. We view the exceptional set as $|V_0|$ distinct parts, each consisting of a single vertex. For two partitions $\mathcal{P}$ and $\mathcal{P}'$ as above, $\mathcal{P}'$ is a refinement of $\mathcal{P}$, if every part in $\mathcal{P}$ is a union of some of the parts of $\mathcal{P}'$. By the last comment on the exceptional set this means, in particular, that if $\mathcal{P}'$ is obtained from $\mathcal{P}$ by shifting vertices from the other sets in the partition to the exceptional set, then $\mathcal{P}'$ is a refinement of $\mathcal{P}$. An equipartition is called ε -regular if $|V_0| \leq \varepsilon|V|$ and all pairs (V_i, V_j) with $1 \leq i < j \leq k$, except at most εk^2 of them, are ε -regular.

Theorem 17.3.1 [The Regularity Lemma, Szemerédi (1978)] For every $\varepsilon > 0$ and every integer t there exists an integer $T = T(\varepsilon, t)$ so that every graph with at least T vertices has an ε -regular partition $(V_0, V_1, \dots, V_k)$, where $t \leq k \leq T$.

The basic idea in the proof is simple. Start with an arbitrary partition of the set of vertices into t disjoint classes of equal sizes (with a few vertices in the exceptional set, if needed, to ensure divisibility by t). Proceed by showing that as long as the existing partition is not ε -regular, it can be refined in a way that increases the weighted average of the square of the density between a pair of classes of the partition by at least a constant depending only on ε . As this average cannot exceed 1, the process has to terminate after a bounded number of refinement steps. Since in each step we control the growth in the number of parts as well as the number of extra vertices thrown to the exceptional set, the desired result follows. The precise details require some care and are given in what follows.

Let $G = (V, E)$ be a graph on $|V| = n$ vertices. For two disjoint subsets $U, W \subset V$, define $q(U, W) = (|U||W|/n^2) d^2(U, W)$. For partitions $\mathcal{U}$ of U and $\mathcal{W}$ of W , define

$$q(\mathcal{U}, \mathcal{W}) = \sum_{\substack{U' \in \mathcal{U} \\ W' \in \mathcal{W}}} q(U', W')$$

Finally, for a partition $\mathcal{P}$ of V , with an exceptional set V_0 , define $q(\mathcal{P}) = \sum q(U, W)$, where the sum ranges over all unordered pairs of distinct parts U, W in the partition, with each vertex of the exceptional set V_0 forming a singleton part in its own. Therefore $q(\mathcal{P})$ is a sum of $\binom{k+|V_0|}{2}$ terms of the form $q(U, W)$. The quantity $q(\mathcal{P})$ is called the index of the partition $\mathcal{P}$. Since $d^2(U, W) \leq 1$ for all U, W , and since the sum $\sum |U||W|$ over all unordered pairs of distinct parts U, W is at most the number of unordered pairs of vertices, it follows that the index of any partition is smaller than $\frac{1}{2}$.

Lemma 17.3.2

- (i) Let U, W be disjoint nonempty subsets of V ; let $\mathcal{U}$ be a partition of U and $\mathcal{W}$ a partition of W . Then $q(\mathcal{U}, \mathcal{W}) \geq q(U, W)$.
- (ii) If $\mathcal{P}'$ and $\mathcal{P}$ are partitions of V and $\mathcal{P}'$ is a refinement of $\mathcal{P}$, then $q(\mathcal{P}') \geq q(\mathcal{P})$.
- (iii) Suppose $\varepsilon > 0$, and suppose U, W are disjoint nonempty subsets of V and the pair (U, W) is not ε -regular. Then there are partitions $\mathcal{U} = \{U_1, U_2\}$ of U and $\mathcal{W} = \{W_1, W_2\}$ of W so that $q(\mathcal{U}, \mathcal{W}) > q(U, W) + \varepsilon^4 |U||W|/n^2$.

Proof.

(i) Define a random variable Z as follows. Let u be a uniformly chosen random element of U , and let w be a uniformly chosen random element of W . Let $U' \in \mathcal{U}$ and $W' \in \mathcal{W}$ be those members of the partition so that $u \in U', w \in W'$. Then $Z = d(U', W')$. The expectation of Z is

$$\sum_{\substack{U' \in \mathcal{U} \\ W' \in \mathcal{W}}} \frac{|U'| |W'|}{|U||W|} d(U', W') = \sum_{\substack{U' \in \mathcal{U} \\ W' \in \mathcal{W}}} \frac{|U'| |W'|}{|U||W|} \cdot \frac{e(U', W')}{|U'| |W'|} = d(U, W).$$

By Jensen's Inequality, $E[Z^2] \geq (E[Z])^2$ and the desired result follows, as

$$E[Z^2] = \frac{n^2}{|U||W|} q(\mathcal{U}, \mathcal{W}) \text{ and } (E[Z])^2 = d^2(U, W) = \frac{n^2}{|U||W|} q(U, W)$$

(ii) This is an immediate consequence of (i).

(iii) Since the pair (U, W) is not ε -regular, there are subsets $U_1 \subset U, W_1 \subset W$ so that $|U_1| \geq \varepsilon |U|, |W_1| \geq \varepsilon |W|$ and $|d(U_1, W_1) - d(U, W)| > \varepsilon$. Put $U_2 = U \setminus U_1, W_2 = W \setminus W_1$ and define the partitions $\mathcal{U} = \{U_1, U_2\}, \mathcal{W} = \{W_1, W_2\}$. Let Z be the random variable defined in the proof of part (i). Then, as shown in that proof,

$$\text{Var}[Z] = E[Z^2] - (E[Z])^2 = \frac{n^2}{|U||W|} (q(\mathcal{U}, \mathcal{W}) - q(U, W))$$

However, as $E[Z] = d(U, W)$ it follows that Z deviates from $E[Z]$ by more than ε with probability $|U_1| |W_1| / |U||W|$, implying that

$$\text{Var}[Z] > \frac{|U_1| |W_1|}{|U||W|} \varepsilon^2 \geq \varepsilon^4$$

This provides the desired result.

Proposition 17.3.3 Suppose $0 < \varepsilon \leq \frac{1}{4}$, let $\mathcal{P} = \{V_0, V_1, \dots, V_k\}$ be an equipartition of V , where V_0 is the exceptional set, $|V_0| \leq \varepsilon n$, and $|V_i| = c$ for all $1 \leq i \leq k$. If $\mathcal{P}$ is not ε -regular

then there exists a refinement $\mathcal{P}' = \{V'_0, V'_1, \dots, V'_\ell\}$ of $\mathcal{P}$, in which $k \leq \ell \leq k4^k$, $|V'_0| \leq |V_0| + n/2^k$, all other sets V_i are of the same size and $q(\mathcal{P}') \geq q(\mathcal{P}) + \frac{1}{2}\varepsilon^5$.

Proof. For every pair $1 \leq i < j \leq k$ define a partition $\mathcal{V}_{ij}$ of V_i and $\mathcal{V}_{ji}$ of V_j as follows. If the pair (V_i, V_j) is ε -regular, then the two partitions are trivial. Otherwise, each partition consists of two parts, chosen according to Lemma 17.3.2, part (iii). For each $1 \leq i \leq k$, let $\mathcal{V}_i$ be the partition of V_i obtained by the Venn diagram of all $(k-1)$ -partitions $\mathcal{V}_{ij}$. Thus each $\mathcal{V}_i$ has at most 2^{k-1} parts. Let $\mathcal{Q}$ be the partition of V consisting of all parts of the partitions $\mathcal{V}_i$ together with the original exceptional set V_0 . By Lemma 17.3.2, parts (ii) and (iii), and since $\mathcal{P}$ is not ε -regular, we conclude that the index of $\mathcal{Q}$ satisfies

$$q(\mathcal{Q}) \geq q(\mathcal{P}) + \varepsilon k^2 \varepsilon^4 \frac{c^2}{n^2} = q(\mathcal{P}) + \varepsilon^5 \frac{(kc)^2}{n^2} > q(\mathcal{P}) + \frac{\varepsilon^5}{2}$$

where here we used the fact that $kc \geq (1-\varepsilon)n \geq \frac{3}{4}n$. Note that $\mathcal{Q}$ has at most $k2^{k-1}$ parts (besides the exceptional set), but those are not necessarily of equal sizes. Define $b = \lfloor c/4^k \rfloor$ and split every part of $\mathcal{Q}$ arbitrarily into disjoint sets of size b , throwing the remaining vertices in each part, if any, to the exceptional set. This process creates a partition $\mathcal{P}'$ with at most $k4^k$ non exceptional parts of equal size, and a new exceptional set V'_0 of size smaller than $|V_0| + k2^{k-1}b < |V_0| + kc/2^k \leq |V_0| + n/2^k$. Moreover, by Lemma 17.3.2, part (ii), the index $q(\mathcal{P}')$ of $\mathcal{P}'$ is at least $q(\mathcal{Q}) > q(\mathcal{P}) + \frac{1}{2}\varepsilon^5$, completing the proof.

Proof [Theorem 17.3.1]. It suffices to prove the lemma for $\varepsilon \leq \frac{1}{4}$ and t satisfying $2^{t-2} > 1/\varepsilon^6$; hence we assume that these inequalities hold. Put $s = \lceil 1/\varepsilon^5 \rceil$ and note that for this choice $1/2^k \leq \varepsilon/2s$ for all $k \geq t$. Define $k_0 = t$ and $k_{i+1} = k_i 4^{k_i}$ for all $i \geq 0$. We prove the lemma with $T = k_s$.

Let $G = (V, E)$ be a graph with $|V| = n \geq T$ vertices. Start with an arbitrary partition $\mathcal{P} = \mathcal{P}_0$ of its vertices into $k = k_0 = t$ pairwise disjoint parts, each of size $\lfloor n/t \rfloor$, and let the exceptional set consist of the remaining vertices, if any. Note that their number is less than t , which is (much) smaller than $\frac{1}{2}\varepsilon n$. As long as the partition $\mathcal{P}$ we have already defined is not ε -regular, apply Proposition 17.3.3 to refine it to a new equipartition $\mathcal{P}'$ with at most $k4^k$ non exceptional parts, whose index exceeds that of $\mathcal{P}$ by at least $\frac{1}{2}\varepsilon^5$, while the size of the exceptional set increases by at most $n/2^k < \varepsilon n/2s$. As the initial index is nonnegative, and the index never exceeds $\frac{1}{2}$, the process must terminate in at most s steps, yielding an ε -regular partition with at most T non exceptional parts, and an exceptional set of size smaller than εn .

Remark. The proof shows that $T(\varepsilon, \lceil 1/\varepsilon \rceil)$ is bounded by a tower of exponents of height roughly $1/\varepsilon^5$. Surprisingly, as shown by Gowers (1997), this tower-type behavior is indeed necessary.

18.8.4 17.4 Testing Triangle-freeness

The relevance of the Regularity Lemma to property testing is nicely illustrated in the proof that the property of containing no triangle is testable with one-sided error. The required combinatorial lemma here is the (intuitive, yet nontrivial) fact that if one has to delete at least εn^2 edges of an n -vertex graph to destroy all triangles in it, then the graph must contain at least δn^3 triangles, where $\delta = \delta(\varepsilon) > 0$. As shown in the exercises, following Ruzsa and Szemerédi (1978), this fact implies that any set of integers with positive upper density contains a three-term arithmetic progression.

Lemma 17.4.1 For any positive $\varepsilon < 1$ there is a $\delta = \delta(\varepsilon) > 0$ so that if $G = (V, E)$ is a graph on $|V| = n$ vertices that is ε -far from being triangle-free, then it contains at least δn^3 triangles.

Proof. We prove the lemma with

$$\delta = \frac{\varepsilon^3}{2^9 T^3 (\varepsilon/4, \lceil 4/\varepsilon \rceil)},$$

where T is as in Theorem 17.3.1. Let $G = (V, E)$ satisfy the assumption. Note, first, that if $n < T(\varepsilon/4, \lceil 4/\varepsilon \rceil)$ then the assertion is trivial, as in this case δn^3 is less than 1, and it is trivial that if G is ε -far from being triangle-free then it contains a triangle. We thus assume that n is at least $T(\varepsilon, t)$, where $t = \lceil 4/\varepsilon \rceil$. By Theorem 17.3.1 there is an $(\varepsilon/4)$ -regular partition $(V_0, V_1, \dots, V_k)$ of G , where $t \leq k \leq T = T(\varepsilon, t)$. Put $c = |V_1| = |V_2| = \dots = |V_k|$. Let G' be the graph obtained from G by deleting the following edges:

- All edges of G that are incident with a vertex of the exceptional set V_0 (there are less than $\frac{1}{4}\varepsilon n^2$ such edges.)
- All edges of G that lie inside some set V_i (there are less than $\frac{1}{8}\varepsilon n^2$ such edges.)
- All edges of G that lie in irregular pairs [there are at most $\frac{1}{4}\varepsilon k^2 c^2 \leq \varepsilon \frac{1}{4}n^2$ such edges.]
- All edges of G that lie in regular pairs (V_i, V_j) , where the density $d(V_i, V_j)$ is smaller than $\frac{1}{2}\varepsilon$ [there are less than $\binom{k}{2} \frac{1}{2}\varepsilon c^2 < \frac{1}{4}\varepsilon n^2$ such edges.]

Since G' is obtained from G by deleting less than εn^2 edges, it contains a triangle, as G is ε -far from being triangle-free. By the definition of G' , the vertices of this triangle must lie in three distinct sets V_i , any two of which form a regular pair of density at least $\frac{1}{2}\varepsilon$. Without loss of generality assume that these sets are V_1, V_2, V_3 . Call a vertex $v_1 \in V_1$ typical if it has at least $\frac{1}{4}\varepsilon c$ neighbors in V_2 and at least $\frac{1}{4}\varepsilon c$ neighbors in V_3 . We claim that all vertices of V_1 but at most $2 \cdot \frac{1}{4}\varepsilon c < \frac{1}{2}c$ are typical. Indeed, if X_1 is the set of all vertices of V_1 that have less than $\frac{1}{4}\varepsilon c$ neighbors in V_2 then its cardinality must be smaller than $\frac{1}{4}\varepsilon c$, since otherwise the pair X_1 and $X_2 = V_2$, together with the fact that $d(V_1, V_2) \geq \frac{1}{2}\varepsilon$, would violate the $(\varepsilon/4)$ -regularity of this pair. Similarly, there are less than $\frac{1}{4}\varepsilon c$ vertices of V_1 that have less than $\frac{1}{4}\varepsilon c$ neighbors in V_3 , proving the claim.

Fix a typical vertex $v_1 \in V_1$, and let N_2, N_3 denote the sets of all its neighbors in V_2 and V_3 , respectively. Thus $|N_2|, |N_3| \geq \frac{1}{4}\varepsilon c$, and hence, by the $(\varepsilon/4)$ -regularity of the pair (V_2, V_3) and the fact that its density is at least $\frac{1}{2}\varepsilon$, there are at least $\frac{1}{4}\varepsilon |N_2| |N_3| \geq (\varepsilon/4)^3 c^2$ edges between N_2 and N_3 . We conclude that v_1 lies in at least $(\varepsilon/4)^3 c^2$ triangles. As there are at least $\frac{1}{2}c$ typical vertices in V_1 , and since

$$c^3 \geq \left(1 - \frac{\varepsilon}{4}\right)^3 n^3 / T^3 > \frac{n^3}{4T^3},$$

the desired result follows. $\square$

Corollary 17.4.2 The property of being triangle-free is testable with one-sided error.

Proof. For $\varepsilon > 0$, let $\delta = \delta(\varepsilon)$ be as in Lemma 17.4.1. Given a graph $G = (V, E)$ on n vertices, consider the following randomized algorithm for testing if G is triangle-free. Let s be a confidence parameter. Pick randomly and independently s/δ triples of vertices of the graph, and check if at least one of them forms a triangle. If so, then report that the graph is not triangle-free, otherwise, report that the graph is triangle-free. Clearly, if G is triangle-free, the algorithm will decide so. If it is ε -far from being triangle-free, then by Lemma 17.4.1, the probability that the algorithm will err and report that G is triangle-free does not exceed $(1 - 6\delta)^{s/\delta} \leq e^{-6s}$. This completes the proof. $\square$

18.8.5 17.5 Characterizing the Testable Graph Properties

In this section we describe several recent results on graph property testing. The proofs of these results apply a strong variant of the Regularity Lemma, proved in Alon, Fischer, Krivelevich and Szegedy (2000). The detailed proofs are somewhat technical and will not be given here.

A graph property is monotone if it is closed under removing vertices and edges. Thus being k -colorable or triangle-free is a monotone property. A property is hereditary if it is closed under removal of vertices (and not necessarily under removal of edges). Clearly every monotone graph property is also hereditary, but there are also many well-studied hereditary properties that are not monotone. Examples are being a perfect graph, a chordal graph, an interval graph and many more. The results discussed in the previous subsections deal with two special cases of hereditary properties that are also monotone, namely, being triangle-free and being k -colorable. Handling hereditary non-monotone graph properties, such as being perfect or not containing an induced cycle of length 4, is more involved than handling monotone properties.

For a (possibly infinite) family of graphs $\mathcal{F}$, a graph G is said to be induced $\mathcal{F}$ -free if it contains no $F \in \mathcal{F}$ as an induced subgraph. The following lemma is not difficult.

Lemma 17.5.1 Let $\mathcal{F}$ be a (possibly infinite) family of graphs, and suppose there are functions $f_{\mathcal{F}}(\varepsilon)$ and $\delta_{\mathcal{F}}(\varepsilon)$ such that the following holds for every $\varepsilon > 0$. Every graph G on n vertices that is ε -far from being induced $\mathcal{F}$ -free contains at least $\delta_{\mathcal{F}}(\varepsilon)n^{f_{\mathcal{F}}}$ induced copies of a graph $F \in \mathcal{F}$ of size $f \leq f_{\mathcal{F}}(\varepsilon)$. Then, being induced $\mathcal{F}$ -free is testable with one-sided error.

The following general result is proved in Alon and Shapira (2005). A subsequent different, elegant, but non effective proof can be found in Lovász and Szegedy (to appear).

Theorem 17.5.2 [Alon and Shapira (2005)] For any (possibly infinite) family of graphs $\mathcal{F}$ there are functions $f_{\mathcal{F}}(\varepsilon)$ and $\delta_{\mathcal{F}}(\varepsilon)$ satisfying the conditions of Lemma 17.5.1.

It is easy to see that one can define for any hereditary property $\mathcal{P}$, a (possibly infinite) family of graphs $\mathcal{F}_{\mathcal{P}}$ such that satisfying $\mathcal{P}$ is equivalent to being induced $\mathcal{F}_{\mathcal{P}}$ -free. Indeed, we simply put a graph F in $\mathcal{F}_{\mathcal{P}}$ if and only if F does not satisfy $\mathcal{P}$. It thus follows that Theorem 17.5.2, combined with Lemma 17.5.1, implies the following.

Theorem 17.5.3 Every hereditary graph property is testable with one-sided error.

An easy consequence of Theorem 17.5.2 is the following.

Corollary 17.5.4 For every hereditary graph property $\mathcal{P}$, there is a function $W_{\mathcal{P}}(\varepsilon)$ with the following property: If G is ε -far from satisfying $\mathcal{P}$, then G contains an induced subgraph of size at most $W_{\mathcal{P}}(\varepsilon)$ that does not satisfy $\mathcal{P}$.

Using Theorem 17.5.3 one can obtain a characterization of the "natural" graph properties, which are testable with one-sided error.

A tester (one-sided or two-sided) is said to be oblivious if it works as follows: Given ε the tester computes an integer $Q = Q(\varepsilon)$ and asks an oracle for a subgraph induced by a set of vertices S of size Q , where the oracle chooses S randomly and uniformly from the vertices of the input graph. If Q is larger than the size of the input graph then the oracle returns the entire graph. The tester then accepts or rejects according to the graph induced by S .

In some sense, oblivious testers capture the essence of property testing as essentially all the testers that have been analyzed in the literature are in fact oblivious, or could easily be turned into oblivious testers. Clearly some properties cannot have oblivious testers; however, these properties are not natural. An example is the property of not containing an induced cycle of length 4 if the number of vertices is even, and not containing an induced cycle of length 5 if the number of vertices is odd.

Using Theorem 17.5.3 it can be shown that if one considers only oblivious testers, then it is possible to precisely characterize the graph properties, which are testable with one-sided error. To state this characterization we need the following definition.

A graph property $\mathcal{P}$ is called semi-hereditary if there exists a hereditary graph property $\mathcal{H}$ such that the following holds.

1. Any graph satisfying $\mathcal{P}$ also satisfies $\mathcal{H}$.
2. For any $\varepsilon > 0$ there is an $M(\varepsilon)$ such that any graph of size at least $M(\varepsilon)$ that is ε -far from satisfying $\mathcal{P}$ does not satisfy $\mathcal{H}$.

Clearly any hereditary graph property $\mathcal{P}$ is also semi-hereditary because we can take $\mathcal{H}$ in the above definition to be $\mathcal{P}$ itself. In simple words, a semi-hereditary $\mathcal{P}$ is obtained by taking a hereditary graph property $\mathcal{H}$ and removing from it a (possibly infinite, carefully chosen) set of graphs. This means that the first item in the definition above is satisfied. The only restriction, which is needed to get item 2 in the definition, is that $\mathcal{P}$ will be such that for any $\varepsilon > 0$ there will be only finitely many graphs that are ε -far from satisfying it, and yet satisfy $\mathcal{H}$. We are now ready to state the characterization.

Theorem 17.5.5 A graph property $\mathcal{P}$ has an oblivious one-sided tester if and only if $\mathcal{P}$ is semi-hereditary.

The proof can be found in Alon and Shapira (2005). The Regularity Lemma and its strong variant mentioned in the beginning of this subsection play a crucial role in this proof. This is not a coincidence. In Alon, Fischer, Newman and Shapira (2006) it is shown that the property defined by having any given Szemerédi-partition is testable with a constant number of queries. This leads to a combinatorial characterization of the graph properties that are testable with a constant number of queries. This characterization (roughly) says that a graph property $\mathcal{P}$ can be tested by a two-sided error tester with a constant number of queries if and only if testing $\mathcal{P}$ can be reduced to testing the property of satisfying one of finitely many Szemerédi-partitions. See Alon et al. (2006) for the precise formulation and detailed proof.

18.8.6 17.6 Exercises

1. Show that for every fixed $\varepsilon > 0$ and $0 < p < 1$ there is an $m_0 = m_0(\varepsilon, p)$ so that for every $n > 2m > m_0$, the probability that two fixed disjoint sets A and B , each of size m , of the random graph $G(n, p)$ do not form an ε -regular pair is smaller than ε .
2. (Removal Lemma.) Show that for any fixed graph H on h vertices and for any $\varepsilon > 0$ there is a $\delta = \delta(\varepsilon, H) > 0$ so that if one has to delete at least εn^2 edges from an n -vertex graph G to destroy all copies of H , then G contains at least δn^h copies of H .
3. (*) Using Lemma 17.4.1 prove that for any $\varepsilon > 0$ there is an n_0 so that if $n > n_0$ then every subset $A \subset \{1, 2, \dots, n\}$ of size $|A| \geq \varepsilon n$ contains a three-term arithmetic progression.
4. Combine Turán's Theorem with the Regularity Lemma to prove the following result, due to Erdős, Simonovits and Stone: For every fixed graph H of chromatic number $r > 1$ and every $\varepsilon > 0$, there is an $n_0 = n_0(H, \varepsilon)$ so that if $n > n_0$ then any simple graph with n vertices and at least

$$\left(1 - \frac{1}{r-1} + \varepsilon\right) \binom{n}{2}$$

edges contains a copy of H .

5. A graph is chordal if any cycle of length at least 4 in it has a chord. Apply Corollary 17.5.4 to show that for every $\varepsilon > 0$ there is a $k = k(\varepsilon)$ so that every graph on n vertices in which every cycle of length at least 4 and at most k has a chord can be transformed into a chordal graph by adding and/or deleting at most εn^2 edges.
6. (*) A construction of Behrend (1946) gives a subset X of $\{1, 2, \dots, m\}$ of size $|X| \geq m/e^{c\sqrt{\log m}}$ with no three-term arithmetic progression. Show how to construct from such an X a graph on n vertices, which is ε -far from being triangle-free, and yet contains only $\varepsilon^{b \log(1/\varepsilon)} n^3$ triangles.
7. Prove that the property of being triangle-free is not testable with a one-sided error tester whose query complexity is polynomial in $1/\varepsilon$.
8. A graph G is H -free if it contains no copy of H . Prove that for every bipartite graph H

with h vertices, there is a $c = c(h) > 0$ so that any graph G on n vertices that is ε -far from being H -free contains at least $\varepsilon^c n^h$ copies of H .

The Probabilistic Lens: Turán Numbers and Dependent Random Choice

For a graph H and an integer n , the Turán number $\text{ex}(n, H)$ is the maximum possible number of edges in a simple graph on n vertices that contains no copy of H . The asymptotic behavior of these numbers for graphs of chromatic number at least 3 is well known; see, for example, Exercise 4 in Chapter 17. For bipartite graphs H , however, the situation is considerably more complicated, and there are relatively few nontrivial bipartite graphs H for which the order of magnitude of $\text{ex}(n, H)$ is known. Here we prove that for every fixed bipartite graph H in which the degrees of all vertices in one color class are at most r , there is a constant $c = c(H)$ so that $\text{ex}(n, H) \leq cn^{2-1/r}$. This is tight for all values of r , as it is known that for every r and $t > (r-1)!$, there is a simple graph with n vertices and at least $c_{r,t} n^{2-1/r}$ edges, containing no copy of the complete bipartite graph $K_{r,t}$.

The basic tool in the proof is a simple and yet surprisingly powerful method, whose probabilistic proof may be called "dependent random choice," as it involves a random selection of a set of vertices, where the choices are dependent in a way that increases the probability that r -tuples of selected vertices will have many common neighbors. An early variant of this lemma was first proved in Kostochka and Rödl (2004) and Gowers (1998). The proof given here is from Alon, Krivelevich and Sudakov (2003).

Lemma 1 Let a, b, n, r be positive integers. Let $G = (V, E)$ be a graph on $|V| = n$ vertices with average degree $d = 2|E|/n$. If

$$\frac{d^r}{n^{r-1}} - \binom{n}{r} \left(\frac{b-1}{n} \right)^r > a - 1, \quad (1)$$

then G contains a subset A_0 of at least a vertices so that every r vertices of A_0 have at least b common neighbors.

Proof. Let T be a (multi)-set of r random vertices of G , chosen uniformly with repetitions. Set

$$A = \{v \in V : T \subseteq N(v)\},$$

where $N(v)$ denotes the set of all neighbors of v . Denote by X the cardinality of A . By linearity of expectation,

$$\begin{aligned} \mathbb{E}[X] &= \sum_{v \in V} \left(\frac{|N(v)|}{n} \right)^r = \frac{1}{n^r} \sum_{v \in V} |N(v)|^r \\ &\geq \frac{1}{n^r} n \left(\frac{\sum_{v \in V} |N(v)|}{n} \right)^r = \frac{1}{n^{r-1}} \left(\frac{2|E|}{n} \right)^r = \frac{d^r}{n^{r-1}}, \end{aligned}$$

where the inequality follows from the convexity of $f(x) = x^r$.

Let Y denote the random variable counting the number of r -tuples in A with fewer than b common neighbors. For a given r -tuple $R \subseteq V$, the probability that R will be a subset of A is precisely $(|N^*(R)|/n)^r$, where $N^*(R)$ denotes the set of all common neighbors of the vertices in R . As there are at most $\binom{n}{r}$ subsets R of cardinality $|R| = r$ for which $|N^*(R)| \leq b-1$, it follows that

$$\mathbb{E}[Y] \leq \binom{n}{r} \left(\frac{b-1}{n} \right)^r$$

Applying linearity of expectation once again, we conclude by (1) that

$$\mathbb{E}[X - Y] \geq \frac{d^r}{n^{r-1}} - \binom{n}{r} \left(\frac{b-1}{n} \right)^r > a - 1$$

Hence there exists a choice for T so that for the corresponding set A we get $X - Y \geq a$. Pick such a set, and omit a point from every r -tuple in it with fewer than b common neighbors. This gives a set A_0 of at least a vertices so that every r vertices in it have at least b common neighbors.

Theorem 2 Let H be a bipartite graph with maximum degree r on one side. Then there exists a constant $c = c(H) > 0$ such that

$$ex(n, H) < cn^{2-1/r}.$$

Proof. Let A and B be the vertex classes of H , and suppose $|A| = a$, $|A| + |B| = b$, where the degree of every vertex $b \in B$ in H does not exceed r . Let $G = (V, E)$ be a graph on $|V| = n$ vertices with average degree $d = 2|E|/n \geq cn^{1-1/r}$, where c satisfies $c^r > (b-1)^r/r! + a - 1$. It is easy to check that (1) holds. To complete the proof, it suffices to show that G must contain a copy of H . By Lemma 1 there is

a subset $A_0 \subset V$ of cardinality $|A_0| = a$ so that every r -subset of A_0 has at least b common neighbors in G . It is now an easy matter to embed H in G . To do so, start by embedding the vertices of A by an arbitrary injective function from A to A_0 . Proceed by embedding the vertices of B one by one in an arbitrary order, making sure that in each step the image of the new embedded vertex is connected to the images of its neighbors in H and is different from the images of all previously embedded vertices. Since every set of (at most) r vertices of A_0 has at least b common neighbors in G , this process can be performed until the images of all b vertices of H are found. This completes the proof.

18.8.7 Appendix A: Bounding of Large Deviations

A. 1 Chernoff Bounds

We give here some basic bounds on large deviations that are useful when employing the probabilistic method. Our treatment is self-contained. Most of the results may be found in, or immediately derived from, the seminal paper of Chernoff (1952). While we are guided by asymptotic considerations the inequalities are proved for all values of the parameters in the specified region. The first result, while specialized, contains basic ideas found throughout the appendix.

Theorem A.1.1 Let $X_i, 1 \leq i \leq n$, be mutually independent random variables with

$$\Pr[X_i = +1] = \Pr[X_i = -1] = \frac{1}{2}$$

and set, following the usual convention,

$$S_n = X_1 + \cdots + X_n.$$

Let $a > 0$. Then

$$\Pr[S_n > a] < e^{-a^2/2n}$$

We require Markov's Inequality, which states: Suppose that Y is an arbitrary nonnegative random variable, $\alpha > 0$. Then

$$\Pr[Y > \alpha \mathbb{E}[Y]] < \frac{1}{\alpha}$$

Proof. Fix n, a and let, for the moment, $\lambda > 0$ be arbitrary. For $1 \leq i \leq n$,

$$\mathbb{E} [e^{\lambda X_i}] = \frac{e^\lambda + e^{-\lambda}}{2} = \cosh(\lambda)$$

We require the inequality $\cosh(\lambda) \leq e^{\lambda^2/2}$, valid for all $\lambda > 0$, the special case $\alpha = 0$ of Lemma A.1.5 below. (The inequality may be shown more easily by comparing the Taylor series of the two functions termwise.)

$$e^{\lambda S_n} = \prod_{i=1}^n e^{\lambda X_i}$$

Since the X_i are mutually independent so are the $e^{\lambda X_i}$; expectations multiply and

$$\mathbb{E} [e^{\lambda S_n}] = \prod_{i=1}^n \mathbb{E} [e^{\lambda X_i}] = \cosh^n(\lambda) < e^{\lambda^2 n/2}$$

We note that $S_n > a$ if and only if $e^{\lambda S_n} > e^{\lambda a}$ and apply Markov's Inequality so that

$$\Pr[S_n > a] = \Pr[e^{\lambda S_n} > e^{\lambda a}] < \mathbb{E} [e^{\lambda S_n}] / e^{\lambda a} \leq e^{\lambda^2 n/2 - \lambda a}.$$

We set $\lambda = a/n$ to optimize the inequality, $\Pr[S_n > a] < e^{-a^2/2n}$ as claimed.

By symmetry we immediately have the following.

Corollary A.1.2 Under the assumptions of Theorem A.1.1,

$$\Pr[|S_n| > a] < 2e^{-a^2/2n}.$$

The proof of Theorem A.1.1 illustrates the basic idea of the Chernoff bounds. We wish to bound $\Pr[X > a]$ for some random variable X . For any positive λ we bound

$$\Pr[X > a] = \Pr[e^{\lambda X} > e^{\lambda a}] \leq \mathbb{E} [e^{\lambda X}] e^{-\lambda a} \quad (\text{A.1})$$

The core idea of the Chernoff bounds is to select that λ that minimizes $\mathbb{E} [e^{\lambda X}] e^{-\lambda a}$. The art to the Chernoff bounds is to select a λ that is reasonably close to optimal and easy to work with, yielding upper bounds on $\Pr[X > a]$ that are, one hopes, good enough for our purposes. Bounds on $\Pr[X < a]$ are similar. For any positive λ we bound

$$\Pr[X < a] = \Pr[e^{-\lambda X} > e^{-\lambda a}] \leq \mathbb{E} [e^{-\lambda X}] e^{\lambda a}.$$

Chernoff bound arguments tend to be cleaner when $\mathbb{E}[X] = 0$. A simple translation, replacing X by $X - \mu$ where $\mu = \mathbb{E}[X]$, is often quite helpful.

It is instructive to examine the case when N is the standard normal distribution and a is positive. In this instance $\mathbb{E} [e^{\lambda N}] = e^{\lambda^2/2}$ and so

$$\Pr[N > a] = \Pr[e^{\lambda N} > e^{\lambda a}] \leq \mathbb{E} [e^{\lambda N}] e^{-\lambda a} = e^{\lambda^2/2 - \lambda a}.$$

Elementary calculus leads to the optimal choice $\lambda = a$ so that

$$\Pr[N > a] < e^{-a^2/2}$$

This compares well, as $a \rightarrow \infty$, with the actual asymptotics

$$\Pr[N > a] = \frac{1}{\sqrt{2\pi}} \int_a^\infty e^{-t^2/2} dt \sim \frac{1}{\sqrt{2\pi}a} e^{-a^2/2}$$

Results with N being normal with mean μ and variance σ^2 are similarly good. This explains, to some extent, the efficacy of the Chernoff bounds. When a random variable X is "roughly" normal the Chernoff bounds on $\Pr[X > a]$ should be quite close to the actual values for a large. In practice, however, precise calculations of $E[e^{\lambda X}]$ can be difficult or impossible to achieve and there can be considerable art in finding approximations for $E[e^{\lambda X}]$ that will allow for good bounds on $\Pr[X > a]$.

Many of our remaining results will deal with distributions X of the following prescribed type.

Assumptions A.1.3

- $p_1, \dots, p_n \in [0, 1]$,
- $X_1, \dots, X_n$ are mutually independent with
 $\Pr[X_i = 1 - p_i] = p_i$ and $\Pr[X_i = -p_i] = 1 - p_i$,
- $p = (p_1 + \dots + p_n)/n$ and $X = X_1 + \dots + X_n$.

Remark. Clearly $E[X] = E[X_i] = 0$. When all $p_i = 1/2$, X has distribution $S_n/2$. When all $p_i = p$, X has distribution $B(n, p) - np$, where $B(n, p)$ is the usual binomial distribution.

Theorem A.1.4 Under Assumptions A.1.3 and with $a > 0$,

$$\Pr[X > a] < e^{-2a^2/n}$$

Lemma A.1.5 For all reals α, β with $|\alpha| \leq 1$,

$$\cosh(\beta) + \alpha \sinh(\beta) \leq e^{\beta^2/2 + \alpha\beta}$$

Proof. This is immediate if $\alpha = \pm 1$ or $|\beta| \geq 100$. If the lemma were false the function $f(\alpha, \beta) = \cosh(\beta) + \alpha \sinh(\beta) - e^{\beta^2/2 + \alpha\beta}$ would assume a positive global maximum in the interior of the rectangle $R = \{(\alpha, \beta) : |\alpha| \leq 1, |\beta| \leq 100\}$.

Setting partial derivatives equal to zero we find

$$\begin{aligned} \sinh(\beta) + \alpha \cosh(\beta) &= (\alpha + \beta)e^{\beta^2/2 + \alpha\beta}, \\ \sinh(\beta) &= \beta e^{\beta^2/2 + \alpha\beta}, \end{aligned}$$

and thus $\tanh(\beta) = \beta$, which implies $\beta = 0$. But $f(\alpha, 0) = 0$ for all α , a contradiction.

Lemma A.1.6 For all $\theta \in [0, 1]$ and all $\lambda, \theta e^{\lambda(1-\theta)} + (1-\theta)e^{-\lambda\theta} \leq e^{\lambda^2/8}$.

Proof. Setting $\theta = \frac{1}{2}(1 + \alpha)$ and $\lambda = 2\beta$, this lemma reduces to Lemma A.1.5.

Proof [Theorem A.1.4]. Let, for the moment, $\lambda > 0$ be arbitrary.

$$E[e^{\lambda X_i}] = p_i e^{\lambda(1-p_i)} + (1-p_i)e^{-\lambda p_i} \leq e^{\lambda^2/8}$$

by Lemma A.1.6. Then

$$E[e^{\lambda X}] = \prod_{i=1}^n E[e^{\lambda X_i}] \leq e^{\lambda^2 n/8}$$

Applying Markov's Inequality,

$$\Pr[X > a] = \Pr[e^{\lambda X} > e^{\lambda a}] < E[e^{\lambda X}] / e^{\lambda a} \leq e^{\lambda^2 n/8 - \lambda a}.$$

We set $\lambda = 4a/n$ to optimize the inequality: $\Pr[X > a] < e^{-2a^2/n}$, as claimed.

Again by symmetry we immediately have the following.

Corollary A.1.7 Under Assumptions A.1.3 and with $a > 0$,

$$\Pr[|X| > a] < 2e^{-2a^2/n}$$

Under Assumptions A.1.3 with λ arbitrary,

$$\begin{aligned} \mathbb{E}[e^{\lambda X}] &= \prod_{i=1}^n \mathbb{E}[e^{\lambda X_i}] = \prod_{i=1}^n (p_i e^{\lambda(1-p_i)} + (1-p_i) e^{-\lambda p_i}) \\ &= e^{-\lambda p n} \prod_{i=1}^n (p_i e^{\lambda} + (1-p_i)) \end{aligned}$$

With λ fixed, the function

$$f(x) = \ln(xe^{\lambda} + 1 - x) = \ln(Bx + 1) \text{ with } B = e^{\lambda} - 1$$

is concave and hence $\sum_{i=1}^n f(p_i) \leq n f(p)$ (Jensen's Inequality). Exponentiating both sides,

$$\prod_{i=1}^n (p_i e^{\lambda} + (1-p_i)) \leq (p e^{\lambda} + (1-p))^n,$$

so that we have the following.

Lemma A.1.8 Under the Assumptions A.1.3,

$$\mathbb{E}[e^{\lambda X}] < e^{-\lambda p n} (p e^{\lambda} + (1-p))^n.$$

Applying this lemma with inequality (A.1) yields the following.

Theorem A.1.9 Under the Assumptions A.1.3 and with $a > 0$,

$$\Pr[X > a] < e^{-\lambda p n} (p e^{\lambda} + (1-p))^n e^{-\lambda a}$$

for all $\lambda > 0$.

Remark. For given p, n, a , an optimal assignment of λ in Theorem A.1.9 is found by elementary calculus to be

$$\lambda = \ln \left[\left(\frac{1-p}{p} \right) \left(\frac{a + np}{n - (a + np)} \right) \right].$$

This value is oftentimes too cumbersome to be useful. We employ suboptimal λ to achieve more convenient results.

Setting $\lambda = \ln(1 + a/pn)$ and using the fact that $(1 + u/n)^n \leq e^u$, Theorem A.1.9 implies the following.

$$\text{Corollary A.1.10 } \Pr[X > a] < e^{a - pn \ln(1 + a/pn) - a \ln(1 + a/pn)}.$$

To simplify further, apply the inequality $\ln(1 + u) \geq u - u^2/2$, valid for all $u \geq 0$, to Corollary A. 1.10 with $u = a/pn$. This gives the following.

$$\text{Theorem A.1.11 } \Pr[X > a] < e^{-a^2/2pn + a^3/2(pn)^2}.$$

When all $p_i = p$, X has variance $np(1-p)$. With $p = o(1)$ and $a = o(pn)$ this bound reflects the approximation of X by a normal distribution with variance $\sim np$. The bound of Theorem A.1.11 hits a minimum at $a = 2pn/3$. For $a > 2pn/3$ we have the simple bound

$$\Pr[X > a] \leq \Pr[X > 2pn/3] < e^{-2pn/27}.$$

This is improved by the following.

Theorem A.1.12 For $\beta > 1$,

$$\Pr[X > (\beta - 1)pn] < [e^{\beta-1} \beta^{-\beta}]^{pn}$$

Proof. Direct "plug in" to Corollary A.1.10.

$X + pn$ may be interpreted as the number of successes in n independent trials when the probability of success in the i th trial is p_i .

Theorem A.1.13 Under Assumptions A.1.3 and with $a > 0$,

$$\Pr[X < -a] < e^{-a^2/2pn}$$

Note that one cannot simply employ "symmetry" as then the roles of p and $1 - p$ are interchanged.

Proof. Let $\lambda > 0$ be, for the moment, arbitrary. Then by the argument preceding Lemma A.1.8,

$$\mathbb{E}[e^{-\lambda X}] \leq e^{\lambda pn} (pe^{-\lambda} + (1 - p))^n$$

Thus

$$\Pr[X < -a] = \Pr[e^{-\lambda X} > e^{\lambda a}] < e^{\lambda pn} (pe^{-\lambda} + (1 - p))^n e^{-\lambda a},$$

analogous to Theorem A.1.9. We employ the inequality $1 + u \leq e^u$, valid for all u , so that

$$pe^{-\lambda} + (1 - p) = 1 + (e^{-\lambda} - 1)p < e^{p(e^{-\lambda} - 1)}$$

and

$$\Pr[X < -a] \leq e^{\lambda pn + np(e^{-\lambda} - 1) - \lambda a} = e^{np(e^{-\lambda} - 1 + \lambda) - \lambda a}$$

We employ the inequality

$$e^{-\lambda} \leq 1 - \lambda + \lambda^2/2$$

valid for all $\lambda > 0$. (Note: The analogous inequality $e^\lambda \leq 1 + \lambda + \lambda^2/2$ is not valid for $\lambda > 0$ and so this method, when applied to $\Pr[X > a]$, requires an "error" term as the one found in Theorem A.1.11.) Now

$$\Pr[X < -a] \leq e^{np\lambda^2/2 - \lambda a}$$

Set $\lambda = a/np$ to optimize the inequality: $\Pr[X < -a] < e^{-a^2/2pn}$, as claimed.

For clarity the following result is often useful.

Corollary A.1.14 Let Y be the sum of mutually independent indicator random variables, $\mu = \mathbb{E}[Y]$. For all $\epsilon > 0$,

$$\Pr[|Y - \mu| > \epsilon\mu] < 2e^{-c_\epsilon\mu},$$

where $c_\epsilon > 0$ depends only on ϵ .

Proof. Apply Theorems A.1.12 and A.1.13 with $Y = X + pn$ and

$$c_\epsilon = \min \left\{ -\ln(e^\epsilon(1 + \epsilon)^{-(1+\epsilon)}), \epsilon^2/2 \right\}.$$

The asymmetry between $\Pr[X < a]$ and $\Pr[X > a]$ given by Theorems A.1.12 and A.1.13 is real. The estimation of X by a normal distribution with zero mean and variance np is roughly valid for estimating $\Pr[X < a]$ for any a and for estimating $\Pr[X > a]$ while $a = o(np)$. But when a and np are comparable or when $a \gg np$ the Poisson behavior "takes over" and $\Pr[X > a]$ cannot be accurately estimated by using the normal distribution.

We conclude with several large deviation results involving distributions other than sums of indicator random variables.

Theorem A.1.15 Let P have Poisson distribution with mean μ . For $\epsilon > 0$

$$\begin{aligned}\Pr[P \leq \mu(1 - \epsilon)] &\leq e^{-\epsilon^2 \mu/2} \\ \Pr[P \geq \mu(1 + \epsilon)] &\leq [e^\epsilon(1 + \epsilon)^{-(1+\epsilon)}]^\mu\end{aligned}$$

Proof. For any s

$$\Pr[P = s] = \lim_{n \rightarrow \infty} \Pr\left[B\left(n, \frac{\mu}{n}\right) = s\right]$$

Apply Theorems A.1.12 and A.1.13.

Theorem A.1.16 Let $X_i, 1 \leq i \leq n$, be mutually independent with all $E[X_i] = 0$ and all $|X_i| \leq 1$. Set $S = X_1 + \cdots + X_n$. Then

$$\Pr[S > a] < e^{-a^2/2n}$$

Proof. Set, as in the proof of Theorem A.1.1, $\lambda = a/n$. Set

$$h(x) = \frac{e^\lambda + e^{-\lambda}}{2} + \frac{e^\lambda - e^{-\lambda}}{2}x$$

For $x \in [-1, 1]$, $e^{\lambda x} \leq h(x)$. [$y = h(x)$ is the chord through the points $x = \pm 1$ of the convex curve $y = e^{\lambda x}$.] Thus

$$E[e^{\lambda X_i}] \leq E[h(X_i)] = h(E[X_i]) = h(0) = \cosh \lambda.$$

The remainder of the proof follows as in Theorem A.1.1.

Theorem A.1.17 Suppose $E[X] = 0$ and no two values of X are ever more than one apart. Then for all $\lambda \geq 0$,

$$E[e^{\lambda X}] \leq e^{\lambda^2/8}.$$

Proof. Fix $b \in [-\frac{1}{2}, \frac{1}{2}]$ with $X \in [\frac{1}{2}(-1 + b), \frac{1}{2}(+1 + b)]$. Let $y = h(x)$ be the straight line intersecting the curve $y = e^{\lambda x}$ at the points $\frac{1}{2}(\pm 1 + b)$. As $e^{\lambda x}$ is a convex function, $e^{\lambda x} \leq h(x)$ for all $x \in [\frac{1}{2}(-1 + b), \frac{1}{2}(+1 + b)]$. Thus

$$E[e^{\lambda X}] \leq E[h(X)] = h(E[X]) = h(0).$$

We calculate $h(0) = e^{\lambda b/2}[\cosh(\lambda/2) - b \sinh(\lambda/2)]$, which is at most $e^{\lambda^2/8}$ by Lemma A.1.5.

Theorem A.1.18 Let $X_i, 1 \leq i \leq n$, be independent random variables with each $E[X_i] = 0$ and no two values of any X_i ever more than one apart. (We allow, however, values of different X_i, X_j to be further apart.) Set $S = X_1 + \cdots + X_n$. Then

$$\Pr[S > a] < e^{-2a^2/n}$$

Proof. $E[e^{\lambda S}] = \prod_{i=1}^n E[e^{\lambda X_i}] \leq e^{n\lambda^2/8}$ by Theorem A.1.17. Then for $\lambda \geq 0$,

$$\Pr[S > a] = \Pr[e^{\lambda S} \geq e^{\lambda a}] \leq \exp\left[\frac{n\lambda^2}{8} - \lambda a\right]$$

and we set $\lambda = 4a/n$.

We have been roughly guided by the notion that if X has mean zero and variance σ^2 then $\Pr[X \geq a\sigma]$ should go like $e^{-a^2/2}$. There are times when this idea is very wrong. Consider Assumptions A.1.3 with all $p_i = 1/n$ so that $X = P_n - 1$, where P_n has the binomial distribution $B(n, 1/n)$, which is asymptotically P , the Poisson distribution with mean one. Then $E[X] = 0$ and $\text{Var}[X] \sim 1$. For a fixed $\Pr[X = a] \rightarrow 1/e(a+1)!$, which is far bigger than $e^{-a^2/2}$. With

this cautionary preamble, we give a general situation for which the notion is asymptotically correct when a is not too large.

Theorem A.1.19 For every $C > 0$ and $\varepsilon > 0$ there exists $\delta > 0$ so that the following holds: Let $\{X_i\}_{i=1}^n$, n arbitrary, be independent random variables with $E[X_i] = 0$, $|X_i| \leq C$ and $\text{Var}[X_i] = \sigma_i^2$. Set $X = \sum_{i=1}^n X_i$ and $\sigma^2 = \sum_{i=1}^n \sigma_i^2$ so that $\text{Var}[X] = \sigma^2$. Then for $0 < a \leq \delta\sigma$,

$$\Pr[X > a\sigma] < e^{-\frac{1}{2}a^2(1-\varepsilon)}.$$

Proof. We set $\lambda = a/\sigma$ so that $0 \leq \lambda \leq \delta$. Then

$$E[e^{\lambda X_i}] = \sum_{k=0}^{\infty} E\left[\frac{\lambda^k}{k!} X_i^k\right] = 1 + \frac{\lambda^2}{2} \sigma_i^2 + \sum_{k=3}^{\infty} \frac{\lambda^k}{k!} E[X_i^k].$$

As $|X_i^k| \leq C^{k-2} X_i^2$ we bound

$$E[X_i^k] \leq E[|X_i^k|] \leq C^{k-2} E[X_i^2] = C^{k-2} \sigma_i^2$$

For $k \geq 3$ we bound $2/k! \leq 1/(k-2)!$ so that

$$E[e^{\lambda X_i}] \leq 1 + \frac{\lambda^2}{2} \sigma_i^2 \left[1 + \sum_{k=3}^{\infty} \frac{(C\lambda)^{k-2}}{(k-2)!} \right] = 1 + \frac{\lambda^2}{2} \sigma_i^2 e^{\lambda C}.$$

We choose δ to satisfy $e^{C\delta} \leq 1 + \varepsilon$. As $\lambda \leq \delta$,

$$E[e^{\lambda X_i}] \leq 1 + \frac{\lambda^2}{2} \sigma_i^2 (1 + \varepsilon) < \exp\left[\frac{\lambda^2}{2} \sigma_i^2 (1 + \varepsilon)\right].$$

This inequality has held for all X_i so

$$E[e^{\lambda X}] = \prod_{i=1}^n E[e^{\lambda X_i}] < \exp\left[\frac{\lambda^2}{2} \sigma^2 (1 + \varepsilon)\right]$$

and

$$\Pr[X > a\sigma] \leq E[e^{\lambda X}] e^{-\lambda a\sigma} < e^{-\frac{1}{2}a^2(1-\varepsilon)}$$

A. 2 LOWER BOUNDS

The Chernoff bounds of the previous section give upper bounds for $\Pr[X > a]$ by examining one value (albeit, the right one!) of the Laplace transform $E[e^{\lambda X}]$. Here we use three values of the Laplace transform to give lower bounds for $\Pr[X > a]$. We shall set

$$\begin{aligned} f(\lambda) &= E[e^{\lambda X}] \\ g_a(\lambda) &= f(\lambda) e^{-\lambda a} \end{aligned}$$

With this notation $\Pr[X > a] \leq g_a(\lambda)$ and the Chernoff bound is achieved by taking that λ minimizing $g_a(\lambda)$. For any positive u and ε ,

$$\begin{aligned} X \geq a + u &\Rightarrow \lambda X \leq (\lambda + \varepsilon)X - \varepsilon a - \varepsilon u, \\ X \leq a - u &\Rightarrow \lambda X \leq (\lambda - \varepsilon)X + \varepsilon a - \varepsilon u, \end{aligned}$$

so that

$$\begin{aligned} E[e^{\lambda X} \chi(X \geq a + u)] &\leq f(\lambda + \varepsilon) e^{-\varepsilon a} e^{-\varepsilon u} \\ E[e^{\lambda X} \chi(X \leq a - u)] &\leq f(\lambda - \varepsilon) e^{\varepsilon a} e^{-\varepsilon u} \end{aligned}$$

Subtracting these from $E[e^{\lambda X}]$ yields

$$E[e^{\lambda X} \chi(|X - a| < u)] \geq f(\lambda) - e^{-\varepsilon u} [f(\lambda + \varepsilon)e^{-\varepsilon a} + f(\lambda - \varepsilon)e^{\varepsilon a}]$$

When $|X - a| < u$, $e^{\lambda X} \leq e^{\lambda u} e^{\lambda a}$ so

$$\Pr[|X - a| < u] \geq e^{-\lambda u} e^{-\lambda a} E[e^{\lambda X} \chi(|X - a| < u)]$$

But $\Pr[X > a - u] \geq \Pr[|X - a| < u]$, giving our general result as follows.

Theorem A.2.1 For any $a, u, \lambda, \varepsilon$ with $u, \lambda, \varepsilon, \lambda - \varepsilon$ all positive,

$$\Pr[X > a - u] \geq e^{-\lambda u} [g_a(\lambda) - e^{-\varepsilon u} [g_a(\lambda + \varepsilon) + g_a(\lambda - \varepsilon)]] .$$

We note that this bound has used only three values of the Laplace transform: $f(\lambda), f(\lambda - \varepsilon), f(\lambda + \varepsilon)$.

It is instructive to examine the case when N is the standard normal distribution. We assume a is positive and are interested in the asymptotics as $a \rightarrow +\infty$. We set $\lambda = a$ so that $g_a(\lambda) = e^{-a^2/2}$. Now

$$g_a(\lambda \pm \varepsilon) = e^{(\lambda \pm \varepsilon)^2/2 - a(\lambda \pm \varepsilon)} = g_a(\lambda) e^{\varepsilon^2/2}$$

The cancellation of the linear (in ε) terms was not serendipity, but rather reflected the critical choice of λ to minimize $\ln(g_a(\lambda))$. Now

$$\Pr[N > a - u] \geq g_a(a) e^{-au} [1 - 2e^{-\varepsilon u} e^{\varepsilon^2/2}]$$

Suppose we take $\varepsilon = u = 2$. This gives

$$\Pr[N > a - 2] \geq e^{-a^2/2} e^{-2a} [1 - 2e^{-2}]$$

Rescaling: $\Pr[N > a] = \Omega(e^{-a^2/2} e^{-4a})$. In contrast we have the upper bound $\Pr[N > a] \leq e^{-a^2/2}$.

In many applications one does not have the precise values of the Laplace transform $f(\lambda)$. Suppose, however, that we have reasonably good estimates in both directions on $f(\lambda)$. Then Theorem A.2.1 will give a lower bound for $\Pr[X > a - u]$ by using a lower bound for $g_a(\lambda)$ and upper bounds for $g_a(\lambda \pm \varepsilon)$. Our goal will be less ambitious than the estimate achieved for the standard normal N . We shall be content to find the asymptotics of the logarithm of $\Pr[X > a]$. In the next result, the X_n may be imagined to be near the normal distribution. The interval for λ could easily be replaced by $[(1 - \gamma)a_n, (1 + \gamma)a_n]$ for any fixed positive γ .

Theorem A.2.2 Let X_n be a sequence of random variables and a_n a sequence of positive reals with $\lim_{n \rightarrow \infty} a_n = \infty$. Assume

$$E[e^{\lambda X_n}] = e^{\frac{1}{2}\lambda^2(1+o(1))}$$

uniformly for $\frac{1}{2}a_n \leq \lambda \leq \frac{3}{2}a_n$. Then

$$\ln \Pr[X_n > a_n] \sim -\frac{a_n^2}{2}$$

Remark. For $X_n = S_n n^{-1/2}$, $E[e^{\lambda X_n}] = \cosh^n(\lambda n^{-1/2})$. When $u \rightarrow 0$, $\ln \cosh(u) \sim \frac{1}{2}u^2$. The conditions of Theorem A.2.2 therefore hold when $a_n = o(\sqrt{n})$ and $a_n \rightarrow +\infty$. That is, $\ln \Pr[S_n > b_n] \sim -b_n^2/2n$ when $\sqrt{n} \ll b_n \ll n$.

Proof. The upper bound is the Chernoff bound with $\lambda = a_n$.

$$\Pr[X_n > a_n] \leq E[e^{\lambda X_n}] e^{-a_n \lambda} = e^{-\frac{1}{2}a_n^2(1+o(1))}$$

For the lower bound we first let $\delta \in (0, 0.01)$ be fixed. We set $\lambda = a = a_n(1 + \delta)$, $u = a_n\delta$, $\varepsilon = \lambda\delta/10$. Applying Theorem A.2.1

$$\Pr[X > a_n] \geq e^{-\lambda u} B$$

with

$$B = g_a(a) - e^{-\varepsilon u} [g_a(a + \varepsilon) + g_a(a - \varepsilon)]$$

But $\ln[g_a(a)] \sim -\frac{1}{2}a^2$ and, analogous to our result for the standard normal,

$$\ln[g_a(a \pm \varepsilon)] \sim \frac{a^2}{2} \left(1 \pm \frac{\delta}{10}\right)^2 - a^2 \left(1 \pm \frac{\delta}{10}\right) = \frac{a^2}{2} \left(-1 + \frac{\delta^2}{100}\right).$$

As $\varepsilon u = a^2\delta^2/10(1 + \delta)$ we have $e^{-\varepsilon u}g_a(a \pm \varepsilon) \ll g_a(a)$. Now B is dominated by its initial term and

$$\Pr[X > a_n] \geq e^{-\lambda u} g_a(a)(1 - o(1))$$

Taking logarithms:

$$\ln[\Pr[X > a_n]] \geq -a_n^2\delta(1 + \delta) - \frac{a_n^2}{2}(1 + \delta)^2(1 + o(1)) - o(1)$$

As this holds for any fixed $\delta \in (0, 0.01)$

$$\ln[\Pr[X > a_n]] \geq -\frac{a_n^2}{2}(1 + o(1))$$

We have seen that $\Pr[S_n > b_n]$ can be well approximated by $\Pr[\sqrt{n}N > b_n]$ as long as $\sqrt{n} \ll b_n \ll n$. For $b_n = \Theta(n)$ this approximation by the normal distribution is no longer valid. Still, we shall see that the Chernoff bounds continue to give the right asymptotic value for $\ln \Pr[S_n > b_n]$. We place this in a somewhat wider context. Ellis (1984) has given far more general results.

Theorem A.2.3 Let Z_n be a sequence of random variables. Let a be a fixed positive real. Set

$$F(\lambda) = \lim_{n \rightarrow \infty} \frac{1}{n} \ln \mathbf{E}[e^{\lambda Z_n}]$$

Suppose that there exists $\lambda > 0$ and an open interval I containing λ such that

1. $F(s)$ exists and has a first and second derivative for all $s \in I$.
2. $F'(\lambda) = a$.
3. F' is a strictly increasing function in I .
4. There is a K so that $|F''(s)| \leq K$ for all $s \in I$.

Then

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln \Pr[Z_n > an] = F(\lambda) - a\lambda$$

Remark. Let X be a random variable whose Laplace transform is well defined. Let Z_n denote the sum of n independent copies of X . Then $F(\lambda) = \ln \mathbf{E}[e^{\lambda X}]$. In particular, suppose $\Pr[X = 1] = \Pr[X = -1] = \frac{1}{2}$ so that $Z_n = S_n$. Then $F(\lambda) = \ln \cosh(\lambda)$. For any $a \in (0, 1)$

there is a positive λ for which $a = F'(\lambda) = \tanh(\lambda)$. The conditions of Theorem A.2.3 hold and give the asymptotics of $\ln(S_n > an)$.

Proof. The upper bound is the Chernoff bound as

$$\Pr[Z_n > an] \leq \mathbb{E}[e^{\lambda Z_n}] e^{-a\lambda n} = e^{n(F(\lambda) - a\lambda + o(1))}$$

For the lower bound we will apply Theorem A.2.1. First note that since F' is continuous and monotone over I it has a continuous inverse H defined over some interval J containing a . Note $H(a) = \lambda$. Let u be positive and sufficiently small so that $H(a + u) \pm u/k \in I$. As

$$\lim_{u \rightarrow 0} H(a + u) \pm \frac{u}{K} = H(a) = \lambda$$

all sufficiently small u satisfy this criterion.

Set $a^* = a + u$ and $\lambda^* = H(a^*)$ so that $F'(\lambda^*) = a^*$. We define

$$g_n(s) = \mathbb{E}[e^{sZ_n}] e^{-sa^*}$$

Theorem A.2.1 (noting that $an = a^*n - un$) states

$$\Pr[Z_n > an] \geq e^{-\lambda^* a^* n} [g_n(\lambda^*) - e^{-\varepsilon un} [g_n(\lambda^* + \varepsilon) + g_n(\lambda^* - \varepsilon)]]$$

We select $\varepsilon = u/K$. Our selection of u assures us that $\lambda^* \pm \varepsilon$ belong to I . We have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln \left(\frac{e^{-\varepsilon un} g_n(\lambda^* + \varepsilon)}{g_n(\lambda^*)} \right) = -\varepsilon u + F(\lambda^* + \varepsilon) - F(\lambda^*) - \varepsilon a^*.$$

We have selected λ^* so that $F'(\lambda^*) = a^*$. Since $|F''(s)| \leq K$ in the interval I Taylor series bounds

$$|F(\lambda^* + \varepsilon) - F(\lambda^*) - \varepsilon a^*| \leq \frac{K}{2} \varepsilon^2.$$

Our choice of ε (chosen to minimize the quadratic though any sufficiently small ε would do) gives that

$$-\varepsilon u + F(\lambda^* + \varepsilon) - F(\lambda^*) - \varepsilon a^* \leq -\frac{u^2}{2K}.$$

Thus $e^{-\varepsilon n} g_n(\lambda^* + \varepsilon) / g_n(\lambda^*)$ drops exponentially quickly. We only use that for n sufficiently large the ratio is less than 0.25. The same argument shows that for n sufficiently large $e^{-\varepsilon n} g_n(\lambda^* - \varepsilon) / g_n(\lambda^*) < 0.25$. For such n we then have

$$\Pr[Z_n > an] \geq \frac{1}{2} e^{-\lambda^* a^* n} g_n(\lambda^*).$$

This lower bound is $\exp[n(F(\lambda^*) - \lambda^* a^* + o(1))]$. Now consider $F(\lambda^*) - \lambda^* a^*$ as a function of u . As $u \rightarrow 0$, $\lambda^* = H(a + u) \rightarrow H(a) = \lambda$. As F is continuous $F(\lambda^*) \rightarrow F(\lambda)$. Clearly $a^* = a + u \rightarrow a$ and therefore $\lambda^* a^* \rightarrow \lambda a$. Thus

$$F(\lambda^*) - \lambda^* a^* \rightarrow F(\lambda) - \lambda a$$

so

$$\Pr[Z_n > an] \geq e^{n(F(\lambda) - \lambda a + o(1))}$$

Remark. Let Z_n be a sequence of random variables with mean and variance μ_n and σ_n^2 , respectively. The analysis of $\Pr[Z_n > \mu_n + \lambda_n \sigma_n]$ frequently (S_n being the premier example) splits into three parts:

1. Small Deviations. $\lambda_n \rightarrow \lambda$, a positive constant. One hopes to prove asymptotic normality so that $\Pr[Z_n > \mu_n + \lambda_n \sigma_n] \rightarrow \Pr[N > \lambda]$. There is a huge literature on asymptotic normality but, for the most part, asymptotic normality is not covered in this work.
2. Large Deviations. $\lambda_n \rightarrow +\infty$ and $\lambda_n = o(\sigma_n)$. One hopes to show that Z_n is approximately normal in the sense that $\ln \Pr[Z_n > \mu_n + \lambda_n \sigma_n] \sim -\lambda_n^2/2$.
3. Very Large Deviations. $\lambda_n \rightarrow +\infty$ and $\lambda_n = \Omega(\sigma_n)$. Here the approximation of Z_n by the normal distribution generally fails but one hopes that the asymptotics of $\ln \Pr[Z_n > \mu_n + \lambda_n \sigma_n]$ may still be found by the methods we have given.

A. 3 EXERCISES

1. The Hajós number of a graph G is the maximum number k such that there are k vertices in G with a path between each pair so that all the $\binom{k}{2}$ paths are internally pairwise vertex disjoint (and no vertex is an internal vertex of a path and an endpoint of another). Is there a graph whose chromatic number exceeds twice its Hajós number?
2. For two subsets A and B of the set $\mathbb{Z}_m$ of integers modulo m and for $g \in \mathbb{Z}_m$, denote

$$s(A, B, g) = |\{(a, b) : a \in A, b \in B, a + b = g\}|.$$

For a partition of $\mathbb{Z}_m$ into two disjoint sets $\mathbb{Z}_m = A \cup B, A \cap B = \emptyset$ denote

$$c(A, B) = \max_{x \in \mathbb{Z}_m} |s(A, A, x) + s(B, B, x) - 2s(A, B, x)|.$$

Prove that for every odd m there is a partition of $\mathbb{Z}_m$ into two disjoint sets A and B such that $c(A, B) = O(\sqrt{m \log m})$.

3. For $a \in (0, 1)$ apply Theorem A.2.3 to find $\lim_n (1/n) \ln \Pr[S_n > an]$ explicitly. Express $\Pr[S_n > an]$ combinatorially as 2^{-n} times the sum of binomial coefficients. Use Stirling's formula to asymptotically evaluate this sum and show that you get the same result for $\lim_n (1/n) \ln \Pr[S_n > an]$.
4. More generally, for $p \in (0, 1)$ fixed, apply Theorem A.2.3 to find the asymptotics of $\ln \Pr[B(n, p) > an]$ for $p < a < 1$ and of $\ln \Pr[B(n, p) < an]$ for $0 < a < p$. Show that an application of Stirling's formula gives the same answer.
5. Let $\{X_i\}_{i=1}^n$ be independent random variables, each chosen uniformly from $\{+1, +2, -3\}$. Set $Y_n = \sum_{i=1}^n X_i$. Let $f(n)$ be the minimal value so that $\Pr[Y_n > f(n)] < 1/n$. Find the asymptotics of $f(n)$. Redo with $1/n$ replaced by n^{-50} . (Note that it doesn't change the answer much!)

The Probabilistic Lens: Triangle-free Graphs Have Large Independence Numbers

Let $\alpha(G)$ denote the independence number of a graph G . It is easy and well known that for every graph G on n vertices with maximum degree d , $\alpha(G) \geq n/(d+1)$. Ajtai, Komlós and Szemerédi (1980) showed that in case G is triangle-free, this can be improved by a logarithmic factor and in fact $\alpha(G) \geq (cn \log d)/d$, where c is an absolute positive constant. Shearer (1983) simplified the proof and improved the constant factor to $c = 1 + o(1)$. Here is a very short proof, without any attempt to optimize c , which is based on a different technique of Shearer (1995) and its modification in Alon (1996).

Proposition 1 Let $G = (V, E)$ be a triangle-free graph on n vertices with maximum degree at most $d \geq 1$. Then $\alpha(G) \geq (n \log d)/8d$, where the logarithm here and in what follows is in base 2.

Proof. If, say, $d < 16$ the result follows from the trivial bound $\alpha(G) \geq n/(d+1)$ and hence we may and will assume that $d \geq 16$. Let W be a random independent set of vertices in G , chosen uniformly among all independent sets in G . For each vertex $v \in V$ define a random variable $X_v = d|\{v\} \cap W| + |N(v) \cap W|$, where $N(v)$ denotes the set of all neighbors of v . We claim that the expectation of X_v satisfies $E[X_v] \geq \frac{1}{4} \log d$.

To prove this claim, let H denote the induced subgraph of G on $V - (N(v) \cup \{v\})$, fix an independent set S in H and let X denote the set of all non neighbors of S in the set $N(v)$, $|X| = x$. It suffices to show that the conditional expectation

$$E[X_v \mid W \cap V(H) = S] \geq \frac{\log d}{4} \quad (1)$$

for each possible S . Conditioning on the intersection $W \cap V(H) = S$ there are precisely $2^x + 1$ possibilities for W : one in which $W = S \cup \{v\}$ and 2^x in which $v \notin W$ and W is the union of S with a subset of X . It follows that the conditional expectation considered in (1) is precisely

$$\frac{d}{2^x + 1} + \frac{x2^{x-1}}{2^x + 1}$$

To check that the last quantity is at least $\frac{1}{4} \log d$ observe that the assumption that this is false implies that $x \geq 1$ and $2^x(\log d - 2x) > 4d - \log d$, showing that $\log d > 2x \geq 2$ and hence $4d - \log d < \sqrt{d}(\log d - 2)$, which is false for all $d \geq 16$. Therefore

$$E[X_v \mid W \cap V(H) = S] \geq \frac{\log d}{4}$$

establishing the claim.

By linearity of expectation we conclude that the expected value of the sum $\sum_{v \in V} X_v$ is at least $\frac{1}{4}n \log d$. On the other hand, this sum is clearly at most $2d|W|$, since each vertex $u \in W$ contributes d to the term X_u in this sum, and its degree in G , which is at most d , to the sum of all other terms X_v . It follows that the expected size of W is at least $(n \log d)/8d$, and hence there is an independent set of size at least this expectation, completing the proof.

The Ramsey number $R(3, k)$ is the minimum number r such that any graph with at least r vertices contains either a triangle or an independent set of size k . The asymptotic behavior of this function has been studied for over fifty years. It turns out that $R(3, k) = \Theta(k^2/\log k)$. The lower bound is a recent result of Kim(1995), based on a delicate probabilistic construction together with some thirty pages of computation. There is no known explicit construction of such a graph, and the largest known explicit triangle-free graph with no independent set of size k , described in Alon (1994), has only $\Theta(k^{3/2})$ vertices. The tight upper bound for $R(3, k)$, proved in Ajtai et al. (1980), is a very easy consequence of the above proposition.

Theorem 2 [Ajtai et al. (1980)] There exists an absolute constant b such that $R(3, k) \leq bk^2/\log k$ for every $k > 1$.

Proof. Let $G = (V, E)$ be a triangle-free graph on $8k^2/\log k$ vertices. If G has a vertex of degree at least k then its neighborhood contains an independent set of size k . Otherwise, by Proposition 1 above, G contains an independent set of size at least

$$\frac{8k^2}{\log k} \frac{\log k}{8k} = k$$

Therefore, in any case $\alpha(G) \geq k$, completing the proof.

18.9 Appendix B Paul Erdős

Working with Paul Erdős was like taking a walk in the hills. Every time when I thought that we had achieved our goal and deserved a rest, Paul pointed to the top of another hill and off we would go.

- Fan Chung

B. 1 PAPERS

Paul Erdős was the most prolific mathematician of the twentieth century, with over 1500 written papers and more than 490 collaborators. This highly subjective list gives only some of the papers that created and shaped the subject matter of this volume. MR and Zbl. refer to reviews in Math Reviews and Zentralblatt, respectively. Chapter and section reference are to pertinent areas of this volume.

- A combinatorial problem in geometry, *Compositio Math* 2 (1935), 463-470 (with George Szekeres); Zbl. 12, 270.
Written when Erdős was still a teenager this gem contains a rediscovery of Ramsey's Theorem and the Monotone Subsequence Theorem. Many authors have written that this paper played a key role in moving Erdős toward a more combinatorial view of mathematics.
- Some remarks on the theory of graphs, *Bull. Am. Math. Soc.* 53 (1947), 292-294, MR 8# 479d; Zbl 32, 192.
The three-page paper that "started" the probabilistic method, giving an exponential lower bound on the Ramsey number $R(k, k)$. Section 1.1.
- The Gaussian law of errors in the theory of additive number theoretic functions, *Am. J. Math.* 62 (1940), 738-742 (with Mark Kac); MR 2# 42c; Zbl. 24, 102. Showing that the number of prime factors of x chosen uniformly from 1 to n has an asymptotically normal distribution. A connection between probability and number theory that was extraordinary for its time. Section 4.2.
- Problems and results in additive number theory, *Colloque sur la Théorie des Nombres*, Bruxelles, 1955, 127-137, George Thone, Liège; Masson and Cie, Paris, 1956; MR 18# 18a; Zbl. 73, 31.
Using random subsets to prove the existence of a set of integers such that every n is represented $n = x + y$ at least once but at most $c \ln n$ times. Resolving a problem Sidon posed to Erdős in the 1930s. This problem continued to fascinate Erdős: see, e.g., Erdős and Tetali (1990). Section 8.6.
- On a combinatorial problem, *Nordisk Mat. Tidsskr.* 11 (1963), 220-223; MR 28# 4068; Zbl. 122, 248.
On a combinatorial problem II, *Acta Math. Acad. Sci. Hung.* 15 (1964), 445-447; MR 29# 4700; Zbl. 201, 337.
Property B . Probabilistic proofs that any $m < 2^{n-1}n$ -sets can be two-colored with no set monochromatic yet there exist $cn^2 2^n n$ -sets that cannot be so colored. Section 1.3.
- On the evolution of random graphs, *Magyar. Tud. Akad. Mat. Kutató Int. Közl.* 5 (1960), 17-61 (with Alfred Rényi); MR 23# A2338; Zbl. 103, 163. Rarely in mathematics can an entire subject be traced to one paper. For random graphs this is the paper. Chapter 10.

- Graph theory and probability, Can. J. Math. 11 (1959), 34-38; MR 21# 876; Zbl. 84, 396.

Proving by probabilistic methods the existence of graphs with arbitrarily high girth and chromatic number. This paper convinced many of the power of the methodology, as the problem had received much attention but no construction had been found. The Probabilistic Lens: High Girth and High Chromatic Number, following Chapter 3.

- Graph theory and probability II, Can. J. Math. 13 (1961), 346-352; MR 22# 10925; Zbl. 97, 391.

Showing the existence of a triangle-free graph on n vertices with no independent set of size $cn^{1/2} \ln n$ vertices, and hence that the Ramsey number $R(3, k) = \Omega(k^2 \ln^{-2} k)$. A technical tour de force that uses probabilistic methods in a very subtle way, particularly considering the early date of publication.

- On circuits and subgraphs of chromatic graphs, Mathematika 9 (1962), 170-175; MR 25 # 3035; Zbl. 109, 165.

Destroying the notion that chromatic number is necessarily a local property, Erdős proves the existence of a graph on n vertices that cannot be k -colored but for which every εn vertices can be three-colored. The Probabilistic Lens: Local Coloring, following Chapter 8.

- On a combinatorial game, J. Combin. Theory, Ser. A 14 (1973), 298-301 (with John Selfridge); MR 48# 5655; Zbl. 293, 05004.

Players alternate turns selecting vertices and the second player tries to stop the first from getting a winning set. The weight function method used was basically probabilistic and was an early use of derandomization. Section 16.1.

B. 2 CONJECTURES

Conjectures were always an essential part of the mathematical life of Paul Erdős. Here are some of our favorites.

- Do sets of integers of positive density necessarily contain arithmetic progressions of arbitrary length? In finite form, is there for all k and all $\varepsilon > 0$, an n_0 so that if $n \geq n_0$ and S is a subset of the first n integers of size at least εn then S necessarily contains an arithmetic progression of length k ? This conjecture was first made by Paul Erdős and Paul Turán in the 1930s. It was solved (positively) by Szemerédi in the 1970s. Let $F(k, \varepsilon)$ denote the minimal n_0 that suffices above. The growth rate of F remains an intriguing question with very recent results due to Gowers.
- Call distinct S, T, U a Δ -system if $S \cap T = S \cap U = T \cap U$. Let $F(n)$ be the minimal m such that given any mn -sets some three form a Δ -system. Erdős and Rado showed that $F(n)$ exists and gave the upper bound $F(n) \leq 2^n n!$. Erdős conjectured that $F(n) < C^n$ for some constant C .
- What are the asymptotics of the Ramsey function $R(k, k)$? In particular, what is the value c (if it exists) of $\lim_k R(k, k)^{1/k}$? The classic 1947 paper of Erdős gives $c \geq \sqrt{2}$ and $c \leq 4$ follows from the proof of Ramsey's Theorem but a half-century has seen no further improvements in c , though there have been some results on lower order terms.
- Write $r_S(n)$ for the number of solutions to the equation $n = x + y$ with $x, y \in S$. Does there exist a set S of positive integers such that $r_S(n) > 0$ for all but finitely many n yet $r_S(n)$ is bounded by some constant K ? The 1955 paper of Erdős referenced above gives S with $r_S(n) = \Theta(\ln n)$.

- Let $m(n)$, as defined in Section 1.3, denote the minimal size of a family of n sets that cannot be two-colored without forming a monochromatic set. What are the asymptotics of $m(n)$? In 1963 and 1964 Erdős found the bounds $\Omega(2^n) \leq m(n) = O(2^{n/2})$ and the lower bound of Radhakrishnan and Srinivasan, shown in Section 3.5, is now $\Omega(2^{n/2})$.
- Given $2^{n-2} + 1$ points in the plane, no three on a line, must some n of them form a convex set? This conjecture dates back to the 1935 paper of Erdős and Szekeres referenced above.
- Let $m(n, k, l)$ denote the size of the largest family of k -element subsets of an n -set such that no l -set is contained in more than one of them. Simple counting gives $m(n, k, l) \leq \binom{n}{l} / \binom{k}{l}$. Erdős and Hanani conjectured in 1963 that for fixed $l < k$ this bound is asymptotically correct; that is, the ratio of $m(n, k, l)$ to $\binom{n}{l} / \binom{k}{l}$ goes to one as $n \rightarrow \infty$. Erdős had a remarkable ability to select problems that were very difficult but not impossible. This conjecture was settled affirmatively by Vojtech Rödl in 1985, as discussed in Section 4.7. The asymptotics of the difference $\binom{n}{l} / \binom{k}{l} - m(n, k, l)$ remains open.

B. 3 ON ERDŐS

There have been numerous books and papers written about the life and mathematics of Paul Erdős. Three deserving particular mention are:

- The Mathematics of Paul Erdős (Ron Graham and Jarik Nešetřil, eds.), Springer-Verlag, Berlin, 1996 (Vols. I and II).
- Combinatorics, Paul Erdős Is Eighty (D. Miklós, V. T. Sós, T. Szönyi, eds.), Bolyai Soc. Math. Studies, Vol. I (1990) and Vol. II (1993).
- Erdős on Graphs - His Legacy of Unsolved Problems, Fan Chung and Ron Graham, A.K. Peters, 1998.

Of the many papers by mathematicians we note the following:

- László Babai, In and out of Hungary: Paul Erdős, his friends, and times. In Combinatorics, Paul Erdős Is Eighty (listed above), Vol. II, 7-93.
- Béla Bollobás, Paul Erdős - Life and work, in The Mathematics of Paul Erdős (listed above), Vol. II, 1-42.
- A. Hajnal, Paul Erdős' Set theory, in The Mathematics of Paul Erdős (listed above), Vol. II, 352-393.
- János Pach, Two places at once: a remembrance of Paul Erdős, Math Intelligencer, Vol. 19 (1997), no. 2, 38-48.

Two popular biographies of Erdős have appeared:

- The Man Who Loved Only Numbers, Paul Hoffman, Hyperion, New York, 1998.
- My Brain Is Open - The Mathematical Journeys of Paul Erdős, Bruce Schechter, Simon & Schuster, New York, 1998.

Finally, George Csicsery has made a documentary film, *N Is a Number, A Portrait of Paul Erdős*, available from the publishers A. K. Peters, which allows one to see and hear Erdős in lecture and among friends, proving and conjecturing.

B. 4 UNCLE PAUL

Paul Erdős died in September 1996 at the age of 83. His theorems and conjectures permeate this volume. This tribute,¹ given by Joel Spencer at the National Meeting of the American Mathematical Society in January 1997, attempts to convey some of the special spirit that we and countless others took from this extraordinary man.

Paul Erdős was a searcher, a searcher for mathematical truth.

Paul's place in the mathematical pantheon will be a matter of strong debate for in that rarefied atmosphere he had a unique style. The late Ernst Straus said it best, in a commemoration of Erdős' seventieth birthday.

In our century, in which mathematics is so strongly dominated by "theory constructors" he has remained the prince of problem solvers and the absolute monarch of problem posers. One of my friends - a great mathematician in his own right - complained to me that "Erdős only gives us corollaries of the great metatheorems which remain unformulated in the back of his mind." I think there is much truth to that observation but I don't agree that it would have been either feasible or desirable for Erdős to stop producing corollaries and concentrate on the formulation of his metatheorems. In many ways Paul Erdős is the Euler of our times. Just as the "special" problems that Euler solved pointed the way to analytic and algebraic number theory, topology, combinatorics, function spaces, etc.; so the methods and results of Erdős' work already let us see the outline of great new disciplines, such as combinatorial and probabilistic number theory, combinatorial geometry, probabilistic and transfinite combinatorics and graph theory, as well as many more yet to arise from his ideas.

Straus, who worked as an assistant to Albert Einstein, noted that Einstein chose physics over mathematics because he feared that one would waste one's powers in pursuing the many beautiful and attractive questions of mathematics without finding the central questions. Straus goes on,

Erdős has consistently and successfully violated every one of Einstein's prescriptions. He has succumbed to the seduction of every beautiful problem he has encountered - and a great many have succumbed to him. This just proves to me that in the search for truth there is room for Don Juans like Erdős and Sir Galahads like Einstein.

I believe, and I'm certainly most prejudiced on this score, that Paul's legacy will be strongest in Discrete Math. Paul's interest in this area dates back to a marvellous paper with George Szekeres in 1935 but it was after World War II that it really flourished. The rise of the Discrete over the past half century has, I feel, two main causes. The first was The Computer, how wonderful that this physical object has led to such intriguing mathematical questions. The second, with due respect to the many others, was the constant attention of Paul Erdős with his famous admonition "Prove and Conjecture!" Ramsey Theory, Extremal Graph Theory, Random Graphs, how many turrets in our mathematical castle were built one brick at a time with Paul's theorems and, equally important, his frequent and always penetrating conjectures.

My own research specialty, The Probabilistic Method, could surely be called The Erdős Method. It was begun in 1947 with a three page paper in the Bulletin of the American Math Society. Paul proved the existence of a graph having certain Ramsey property without actually constructing it. In modern language he showed that an appropriately defined random graph would have the property with positive probability and hence there must exist a graph with the property. For the next twenty years Paul was a "voice in the wilderness," his colleagues admired his amazing results but adaption of the methodology was slow. But Paul persevered -

⁰¹ Reprinted with permission from the Bulletin of the American Mathematical Society.

he was always driven by his personal sense of mathematical aesthetics in which he had supreme confidence - and today the method is widely used in both Discrete Math and in Theoretical Computer Science.

There is no dispute over Paul's contribution to the spirit of mathematics. Paul Erdős was the most inspirational man I have ever met. I began working with Paul in the late 1960s, a tumultuous time when "do your own thing" was the admonition that resonated so powerfully. But while others spoke of it, this was Paul's *modus operandi*. He had no job; he worked constantly. He had no home; the world was his home. Possessions were a nuisance, money a bore. He lived on a web of trust, travelling ceaselessly from Center to Center, spreading his mathematical pollen.

What drew so many of us into his circle? What explains the joy we have in speaking of this gentle man? Why do we love to tell Erdős stories? I've thought a great deal about this and I think it comes down to a matter of belief, or faith. We mathematicians know the beauties of our subject and we hold a belief in its transcendent quality. God created the integers, the rest is the work of Man. Mathematical truth is immutable, it lies outside physical reality. When we show, for example, that two n th powers never add to an n th power for $n \geq 3$ we have discovered a Truth. This is our belief, this is our core motivating force. Yet our attempts to describe this belief to our nonmathematical friends are akin to describing the Almighty to an atheist. Paul embodied this belief in mathematical truth. His enormous talents and energies were given entirely to the Temple of Mathematics. He harbored no doubts about the importance, the absoluteness, of his quest. To see his faith was to be given faith. The religious world might better have understood Paul's special personal qualities. We knew him as Uncle Paul.

I do hope that one cornerstone of Paul's, if you will, theology will long survive. I refer to The Book. The Book consists of all the theorems of mathematics. For each theorem there is in The Book just one proof. It is the most aesthetic proof, the most insightful proof, what Paul called The Book Proof. And when one of Paul's myriad conjectures was resolved in an "ugly" way Paul would be very happy in congratulating the prover but would add, "Now, let's look for The Book Proof." This platonic ideal spoke strongly to those of us in his circle. The mathematics was there, we had only to discover it.

The intensity and the selflessness of the search for truth were described by the writer Jorge Luis Borges in his story "The Library of Babel." The narrator is a worker in this library which contains on its infinite shelves all wisdom. He wanders its infinite corridors in search of what Paul Erdős might have called The Book. He cries out,

To me, it does not seem unlikely that on some shelf of the universe there lies a total book. I pray the unknown gods that some man - even if only one man, and though it have been thousands of years ago! - may have examined and read it. If honor and wisdom and happiness are not for me, let them be for others. May heaven exist though my place be in hell. Let me be outraged and annihilated but may Thy enormous Library be justified, for one instant, in one being.

In the summer of 1985 I drove Paul to what many of us fondly remember as Yellow Pig Camp - a mathematics camp for talented high school students at Hampshire College. It was a beautiful day - the students loved Uncle Paul and Paul enjoyed nothing more than the company of eager young minds. In my introduction to his lecture I discussed The Book but I made the mistake of describing it as being "held by God." Paul began his lecture with a gentle correction that I shall never forget. "You don't have to believe in God," he said, "but you should believe in The Book."

18.10 References

- Ahlsvede, R. and Daykin, D. E. (1978). An inequality for the weights of two families of sets, their unions and intersections, *Z. Wahrscheinl. V. Geb* 43: 183-185.
- Aho, A. V., Hopcroft, J. E. and Ullman, J. D. (1974). *The Design and Analysis of Computer Algorithms*, Addison Wesley, Reading, MA.
- Aizenman, M. and Barsky, D. J. (1987). Sharpness of the phase transition in percolation models, *Commun. Math. Phys.* 108(3): 489-526.
- Aizenman, M. and Newman, C. M. (1984). Tree graph inequalities and critical behavior in percolation models, *J. Stat. Phys.* 36(1-2): 107-143.
- Aizenman, M., Kesten, H. and Newman, C. M. (1987). Uniqueness of the infinite cluster and continuity of connectivity functions for short and long range percolation, *Commun. Math. Phys.* 111(4): 505-531.
- Ajtai, M. (1983). Σ_1^1 -formulae on finite structures, *Ann. Pure Appl. Logic* 24: 1-48.
- Ajtai, M., Chvátal, V., Newborn, M. M. and Szemerédi, E. (1982). Crossing-free subgraphs, *Theory and Practice of Combinatorics*, North Holland Math. Stud. 60: 9-12.
- Ajtai, M., Komlós, J. and Szemerédi, E. (1980). A note on Ramsey numbers, *J. Combin. Theory, Ser. A* 29: 354-360.
- Ajtai, M., Komlós, J. and Szemerédi, E. (1983). Sorting in $c \log n$ parallel steps, *Combinatorica* 3: 1-19.
- Ajtai, M., Komlós, J. and Szemerédi, E. (1987). Deterministic simulation in LOGSPACE, *Proceedings of the 19th Annual ACM Symposium on Theory of Computing (STOC)*, ACM Press, New York, pp. 132-140.
- Akiyama, J., Exoo, G. and Harary, F. (1981). Covering and packing in graphs IV: Linear arboricity, *Networks* 11: 69-72.
- Alon, N. (1986a). Eigenvalues and expanders, *Combinatorica* 6: 83-96.
- Alon, N. (1986b). Eigenvalues, geometric expanders, sorting in rounds and Ramsey theory, *Combinatorica* 6: 207-219.
- Alon, N. (1988). The linear arboricity of graphs, *Israel J. Math* 62: 311-325.
- Alon, N. (1990a). The maximum number of Hamiltonian paths in tournaments, *Combinatorica* 10: 319-324.
- Alon, N. (1990b). Transversal numbers of uniform hypergraphs, *Graphs and Combinatorics* 6: 1-4.
- Alon, N. (1994). Explicit Ramsey graphs and orthonormal labelings, *The Electronic J. Combinatorics* 1: 8 pp. R12.
- Alon, N. (1996). Independence numbers of locally sparse graphs and a Ramsey type problem, *Random Structures Algorithms* 9: 271-278.
- Alon, N. and Boppana, R. B. (1987). The monotone circuit complexity of boolean functions, *Combinatorica* 7: 1-22.
- Alon, N. and Chung, F. R. K. (1988). Explicit construction of linear sized tolerant networks, *Discrete Math.* 72: 15-19.
- Alon, N. and Frankl, P. (1985). The maximum number of disjoint pairs in a family of subsets, *Graphs and Combinatorics* 1: 13-21.
- Alon, N. and Kleitman, D. J. (1990). Sum-free subsets, in: *A Tribute to Paul Erdős* (A. Baker, B. Bollobás and A. Hajnal, eds.), Cambridge University Press, Cambridge, England, pp. 13-26.
- Alon, N. and Krivelevich, M. (1997). The concentration of the chromatic number of random graphs, *Combinatorica* 17: 303-313.
- Alon, N. and Krivelevich, M. (2002). Testing k -colorability, *SIAM J. Discrete Math.* 15(2): 211-227 (electronic).

- Alon, N. and Linial, N. (1989). Cycles of length 0 modulo k in directed graphs, *J. Combin. Theory, Ser. B* 47: 114-119.
- Alon, N. and Milman, V. D. (1984). Eigenvalues, expanders and superconcentrators, *Proceedings of the 25th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, New York, pp. 320-322. See also: N. Alon and V. D. Milman, λ_1 isoperimetric inequalities for graphs and superconcentrators, *J. Combin. Theory, Ser. B*, 38: 73-88, 1985.
- Alon, N. and Shapira, A. (2005). A characterization of the (natural) graph properties testable with one-sided error, *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 429-438.
- Alon, N. and Shapira, A. (2006). Homomorphisms in graph property testing, *Topics in Discrete Mathematics*, Vol. 26 of *Algorithms Combin.*, Springer, Berlin, pp. 281-313.
- Alon, N., Babai, L. and Itai, A. (1986). A fast and simple randomized parallel algorithm for the maximal independent set problem, *J. Algorithms* 7: 567-583.
- Alon, N., Duke, R. A., Lefmann, H., Rödl, V. and Yuster, R. (1994). The algorithmic aspects of the regularity lemma, *J. Algorithms* 16(1): 80-109.
- Alon, N., Fischer, E., Krivelevich, M. and Szegedy, M. (2000). Efficient testing of large graphs, *Combinatorica* 20(4): 451-476.
- Alon, N., Fischer, E., Newman, I. and Shapira, A. (2006). A combinatorial characterization of the testable graph properties: it's all about regularity, *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, ACM, New York, pp. 251-260.
- Alon, N., Frankl, P. and Rödl, V. (1985). Geometrical realization of set systems and probabilistic communication complexity, *Proceedings of the 26th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, New York, pp. 277-280.
- Alon, N., Goldreich, O., Hastad, J. and Peralta, R. (1990). Simple constructions of almost k -wise independent random variables, *Proceedings of the 31st Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, New York, pp. 544-553.
- Alon, N., Kim, J. H. and Spencer, J. H. (1997). Nearly perfect matchings in regular simple hypergraphs, *Israel J. Math* 100: 171-187.
- Alon, N., Krivelevich, M. and Sudakov, B. (2003). Turán numbers of bipartite graphs and related Ramsey-type questions, *Combin. Probab. Comput.* 12(5-6): 477-494. Special issue on Ramsey theory.
- Alon, N., Rónyai, L. and Szabó, T. (1999). Norm-graphs: variations and applications, *J. Combin. Theory, Ser. B* 76: 280-290.
- Andreev, A. E. (1985). On a method for obtaining lower bounds for the complexity of individual monotone functions, *Dokl. Akad. Nauk SSSR* 282(5): 1033-1037. (In Russian). English translation in *Sov. Math. Dokl.*, 31(3): 530-534.
- Andreev, A. E. (1987). On a method for obtaining more than quadratic effective lower bounds for the complexity of π -schemes, *Vestnik Moskov. Univ. Ser. I Mat. Mekh* 42(1): 63-66. (In Russian).
- Baik, J., Deift, P. and Johansson, K. (1999). On the distribution of the length of the longest increasing subsequence of random permutations, *J. Am. Math. Soc.* 12: 1119-1178.
- Bárány, I. and Füredi, Z. (1987). Empty simplices in Euclidean spaces, *Can. Math. Bull.* 30: 436-445.
- Beck, J. (1978). On 3-chromatic hypergraphs, *Discrete Math.* 24: 127-137.
- Beck, J. (1981). Roth's estimate of the discrepancy of integer sequences is nearly optimal, *Combinatorica* 1: 319-325.
- Beck, J. (1991). An algorithmic approach to the Lovász local lemma. I., *Random Structures Algorithms* 2(4): 343-365.
- Beck, J. and Fiala, T. (1981). Integer-making theorems, *Discrete Appl. Math.* 3: 1-8.
- Behrend, F. A. (1946). On sets of integers which contain no three terms in arithmetical pro-

gression, *Proc. Natl. Acad. Sci. U.S.A.* 32: 331-332.

Bernstein, S. N. (1912). Démonstration du théorème de Weierstrass fondée sur le calcul des probabilités, *Commun. Soc. Math. Kharkov* 13: 1-2.

Billingsley, P. (1995). *Probability and Measure*, 3rd ed., Wiley Series in Probability and Mathematical Statistics, Wiley, Hoboken, NJ.

Blum, N. (1984). A boolean function requiring $3n$ network size, *Theor. Comput. Sci.* 28: 337-345.

Bollobás, B. (1965). On generalized graphs, *Acta Math. Acad. Sci. Hung.* 16: 447-452.

Bollobás, B. (1984). The evolution of random graphs, *Trans. Am. Math. Soc.* 286(1): 257-274.

Bollobás, B. (1988). The chromatic number of random graphs, *Combinatorica* 8: 49-55.

Bollobás, B. (2001). *Random Graphs*, 2nd ed., Vol. 73 of Cambridge Studies in Advanced Mathematics, Cambridge University Press, Cambridge.

Bollobás, B. and Erdős, P. (1976). Cliques in random graphs, *Math. Proc. Cambridge Philos. Soc.* 80: 419-427.

Boppana, R. B. and Spencer, J. H. (1989). A useful elementary correlation inequality, *J. Combin. Theory, Ser. A* 50: 305-307.

Borgs, C., Chayes, J. T., van der Hofstad, R., Slade, G. and Spencer, J. H. (2005). Random subgraphs of finite graphs. I. The scaling window under the triangle condition, *Random Structures Algorithms* 27(2): 137-184.

Brégman, L. M. (1973). Some properties of nonnegative matrices and their permanents, *Sov. Math. Dokl.* 14: 945-949.

Burton, R. M. and Keane, M. (1989). Density and uniqueness in percolation, *Commun. Math. Phys.* 121(3): 501-505.

Chazelle, B. and Welzl, E. (1989). Quasi-optimal range searching in spaces of finite VC-dimension, *Discrete and Computational Geometry* 4: 467-489.

Chernoff, H. (1952). A measure of the asymptotic efficiency for tests of a hypothesis based on the sum of observations, *Ann. Math. Stat.* 23: 493-509.

Chung, F. R. K., Frankl, P., Graham, R. L. and Shearer, J. B. (1986). Some intersection theorems for ordered sets and graphs, *J. Combin. Theory, Ser. A* 43: 23-37.

Chung, F. R. K., Graham, R. L. and Wilson, R. M. (1989). Quasi-random graphs, *Combinatorica* 9: 345-362.

Cicalese, F. and Mundici, D. (2000). Optimal coding with one asymmetric error: below the sphere packing bound, *Computing and Combinatorics (Sydney, 2000)*, Vol. 1858 of Lecture Notes in Computer Science, Springer, Berlin, pp. 159-169.

Cohen, A. and Wigderson, A. (1989). Dispersers, deterministic amplification, and weak random sources, *Proceedings of the 30th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, New York, pp. 14-19.

Danzer, L. and Grünbaum, B. (1962). Über zwei probleme bezüglich konvexer Körper von P. Erdős und von V. L. Klee, *Math. Z.* 79: 95-99.

de la Vega, W. F. (1983). On the maximal cardinality of a consistent set of arcs in a random tournament, *J. Combin. Theory, Ser. B* 35: 328-332.

Dudley, R. M. (1978). Central limit theorems for empirical measures, *Ann. Probab.* 6: 899-929.

Dumitriu, I. and Spencer, J. H. (2004). A halfliar's game, *Theor. Comput. Sci.* 313(3): 353-369. Algorithmic combinatorial game theory.

Elekes, G. (1997). On the number of sums and products, *Acta Arith.* 81: 365-367.

Ellis, R. S. (1984). Large deviations for a general class of random vectors, *Ann. Probab.* 12(1): 1-12.

- Erdős, P. (1947). Some remarks on the theory of graphs, *Bull. Am. Math. Soc.* 53: 292-294.
- Erdős, P. (1956). Problems and results in additive number theory, *Colloque sur le Théorie des Nombres (CBRM, Bruxelles)*, pp. 127-137.
- Erdős, P. (1959). Graph theory and probability, *Can. J. Math.* 11: 34-38.
- Erdős, P. (1962). On circuits and subgraphs of chromatic graphs, *Mathematika* 9: 170-175.
- Erdős, P. (1963a). On a combinatorial problem, *Nordisk Mat. Tidsskr.* 11: 5-10.
- Erdős, P. (1963b). On a problem of graph theory, *Math. Gaz.* 47: 220-223.
- Erdős, P. (1964). On a combinatorial problem II, *Acta Math. Acad. Sci. Hung.* 15: 445-447.
- Erdős, P. (1965a). Extremal problems in number theory, *Proc. Symp. Pure Math. (AMS) VIII*: 181-189.
- Erdős, P. (1965b). On extremal problems of graphs and generalized graphs, *Israel J. Math.* 2: 189-190.
- Erdős, P. and Füredi, Z. (1983). The greatest angle among n points in the d dimensional Euclidean space, *Ann. Discrete Math.* 17: 275-283.
- Erdős, P. and Hanani, H. (1963). On a limit theorem in combinatorial analysis, *Publ. Math. Debrecen* 10: 10-13.
- Erdős, P. and Kac, M. (1940). The Gaussian law of errors in the theory of additive number theoretic functions, *Am. J. Math.* 62: 738-742.
- Erdős, P. and Lovász, L. (1975). Problems and results on 3-chromatic hypergraphs and some related questions, in: *Infinite and Finite Sets (A. Hajnal, R. Rado and V. T. Sós, eds.)*, North-Holland, Amsterdam, pp. 609-628.
- Erdős, P. and Moon, J. W. (1965). On sets of consistent arcs in a tournament, *Can. Math. Bull.* 8: 269-271.
- Erdős, P. and Rényi, A. (1960). On the evolution of random graphs, *Magyar Tud. Akad. Mat. Kutató Int. Közl.* 5: 17-61.
- Erdős, P. and Selfridge, J. L. (1973). On a combinatorial game, *J. Combin. Theory, Ser. A* 14: 298-301.
- Erdős, P. and Spencer, J. H. (1991). Lopsided Lovász local lemma and latin transversals, *Discrete Appl. Math.* 30: 151-154.
- Erdős, P. and Tetali, P. (1990). Representations of integers as the sum of k terms, *Random Structures Algorithms* 1(3): 245-261.
- Fagin, R. (1976). Probabilities in finite models, *J. Symbolic Logic* 41: 50-58.
- Fischer, E. (2001). The art of uninformed decisions: A primer to property testing, *Bull. Eur. Assoc. Theor. Comput. Sci. EATCS* 75: 97-126.
- Fishburn, P. (1992). Correlation in partially ordered sets, *Discrete Appl. Math.* 39: 173-191.
- Fortuin, C. M., Kasteleyn, P. W. and Ginibre, J. (1971). Correlation inequalities on some partially ordered sets, *Commun. Math. Phys.* 22: 89-103.
- Füredi, Z. (1988). Matchings and covers in hypergraphs, *Graphs Combin.* 4: 115-206.
- Frankl, P. and Wilson, R. M. (1981). Intersection theorems with geometric consequences, *Combinatorica* 1: 357-368.
- Frankl, P., Rödl, V. and Wilson, R. M. (1988). The number of submatrices of given type in a Hadamard matrix and related results, *J. Combin. Theory, Ser. B* 44: 317-328.
- Furst, M., Saxe, J. and Sipser, M. (1984). Parity, circuits and the polynomial hierarchy, *Math. Sys. Theory* 17: 13-27.
- Glebskii, Y. V., Kogan, D. I., Liagonkii, M. I. and Talanov, V. A. (1969). Range and degree of realizability of formulas the restricted predicate calculus, *Cybernetics* 5: 142-154. (Russian original: *Kibernetika* 5: 17-27).

- Goldreich, O. (1999). Combinatorial property testing (a survey), Randomization Methods in Algorithm Design (Princeton, NJ, 1997), Vol. 43 of DIMACS Series in Discrete Mathematics and Theoretical Computer Science, American Mathematics Society, Providence, RI, pp. 45-59.
- Goldreich, O. and Ron, D. (2002). Property testing in bounded degree graphs, *Algorithmica* 32(2): 302-343.
- Goldreich, O., Goldwasser, S. and Ron, D. (1998). Property testing and its connection to learning and approximation, *J. ACM* 45(4): 653-750.
- Gowers, W. T. (1997). Lower bounds of tower type for Szemerédi's uniformity lemma, *Geom. Funct. Anal.* 7(2): 322-337.
- Gowers, W. T. (1998). A new proof of Szemerédi's theorem for arithmetic progressions of length four, *Geom. Funct. Anal.* 8(3): 529-551.
- Graham, R. L. and Spencer, J. H. (1971). A constructive solution to a tournament problem, *Can. Math. Bull.* 14: 45-48.
- Graham, R. L., Rothschild, B. L. and Spencer, J. H. (1990). *Ramsey Theory*, 2nd ed., Wiley, Hoboken, NJ.
- Grimmett, G. (1999). *Percolation*, 2nd ed., Vol. 321 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], Springer-Verlag, Berlin.
- Halberstam, H. and Roth, K. F. (1983). *Sequences*, 2nd ed., Springer Verlag, Berlin.
- Hall, M. (1986). *Combinatorial Theory*, 2nd ed., Wiley, Hoboken, NJ.
- Hara, T. and Slade, G. (1990). Mean-field critical behaviour for percolation in high dimensions, *Commun. Math. Phys.* 128(2): 333-391.
- Hara, T. and Slade, G. (1994). Mean-field behaviour and the lace expansion, *Probability and phase transition* (Cambridge, 1993), Vol. 420 of NATO Advanced Study Institute Series C: Mathematical and Physical Science, Kluwer Academic, Dordrecht, pp. 87-122.
- Harper, L. (1966). Optimal numberings and isoperimetric problems on graphs, *J. Combin. Theory* 1: 385-393.
- Harris, T. E. (1960). Lower bound for the critical probability in a certain percolation process, *Math. Proc. Cambridge Philos. Soc.* 56: 13-20.
- Haussler, D. (1995). Sphere packing numbers for subsets of the boolean n -cube with bounded Vapnik-Chervonenkis dimension, *J. Combin. Theory, Ser. A* 69: 217-232.
- Haussler, D. and Welzl, E. (1987). ϵ -nets and simplex range queries, *Discrete and Computational Geometry* 2: 127-151.
- Håstad, J. (1988). Almost optimal lower bounds for small depth circuits, in: *Randomness and Computation* (S. Micali, ed.), Vol. 5 of *Advances in Computer Research*, JAI Press, Greenwich, CT, pp. 143-170.
- Håstad, J. (1998). The shrinkage exponent of De Morgan formulas is 2, *SIAM J. Comput.* 27: 48-64.
- Janson, S. (1990). Poisson approximation for large deviations, *Random Structures Algorithms* 1: 221-230.
- Janson, S. (1998). New versions of Suen's correlation inequality, *Random Structures Algorithms* 13: 467-483.
- Janson, S., Knuth, D., Łuczak, T. and Pittel, B. (1993). The birth of the giant component, *Random Structures Algorithms* 4: 233-358.
- Janson, S., Łuczak, T. and Ruciński, A. (1990). An exponential bound for the probability of nonexistence of a specified subgraph in a random graph, *Random Graphs '87* (Poznań, 1987), Wiley, Chichester, pp. 73-87.
- Janson, S., Łuczak, T. and Ruciński, A. (2000). *Random Graphs*, Wiley, Hoboken, NJ.
- Joffe, A. (1974). On a set of almost deterministic k -independent random variables, *Ann. Probab.* 2: 161-162.
- Kahn, J. (1996). Asymptotically good list colorings, *J. Combin. Theory, Ser. A* 73: 1-59.

- Karchmer, M. and Wigderson, A. (1990). Monotone circuits for connectivity require super-logarithmic depth, *SIAM J. Discrete. Math.* 3: 255-265.
- Karp, R. M. and Ramachandran, V. (1990). Parallel algorithms for shared memory machines, in: *Handbook of Theoretical Computer Science* (J. Van Leeuwen, ed.), Vol. A, Chapter 17, Elsevier, New York, pp. 871-941.
- Katchalski, M. and Meir, A. (1988). On empty triangles determined by points in the plane, *Acta Math. Hung.* 51: 323-328.
- Katona, G. O. H. (1972). A simple proof of the Erdős-Ko-Rado theorem, *J. Combin. Theory, Ser. B* 13: 183-184.
- Khrapchenko, V. M. (1971). A method of determining lower bounds for the complexity of Π -schemes, *Mat. Zametki* 10(1): 83-92. (In Russian). English translation in *Math. Notes Acad. Sci. USSR*, 11:474-479, 1972.
- Kim, J. H. (1995). The Ramsey number $R(3, t)$ has order of magnitude $t^2/\log t$, *Random Structures Algorithms* 7: 173-207.
- Kim, J. H. and Vu, V. (2000). Concentration of multivariate polynomials and its applications, *Combinatorica* 20(3): 417-434.
- Kleitman, D. J. (1966a). Families of non-disjoint subsets, *J. Combin. Theory* 1: 153-155.
- Kleitman, D. J. (1966b). On a combinatorial conjecture of Erdős, *J. Combin. Theory* 1: 209-214.
- Kleitman, D. J., Shearer, J. B. and Sturtevant, D. (1981). Intersection of k -element sets, *Combinatorica* 1: 381-384.
- Kolountzakis, M. N. (1999). An effective additive basis for the integers, *Discrete Math.* 145: 307-313.
- Komlós, J. and Simonovits, M. (1996). Szemerédi's regularity lemma and its applications in graph theory, *Combinatorics, Paul Erdős Is Eighty* (Keszthely, 1993), Vol. 2 of *Bolyai Soc. Math. Stud.*, János Bolyai Math. Soc., Budapest, pp. 295-352.
- Komlós, J., Pach, J. and Woeginger, G. (1992). Almost tight bounds on ϵ -nets, *Discrete and Computational Geometry* 7: 163-173.
- Komlós, J., Pintz, J. and Szemerédi, E. (1982). A lower bound for Heilbronn's problem, *J. London Math. Soc.* 25(2): 13-24.
- Kostochka, A. V. and Rödl, V. (2004). On graphs with small Ramsey numbers. II, *Combinatorica* 24(3): 389-401.
- Loomis, L. H. and Whitney, H. (1949). An inequality related to the isoperimetric inequality, *Bull. Am. Math. Soc.* 55: 961-962.
- Lovász, L. and Szegedy, B. (to appear). Testing properties of graphs and functions, Preprint.
- Lovász, L., Spencer, J. H. and Vesztergombi, K. (1986). Discrepancy of set systems and matrices, *Eur. J. Comb.* 7: 151-160.
- Lubotzky, A., Phillips, R. and Sarnak, P. (1986). Explicit expanders and the Ramanujan conjectures, *Proceedings of the 18th Annual ACM Symposium on Theory of Computing (STOC)*, ACM Press, pp. 240-246. See also: A. Lubotzky, R. Phillips and P. Sarnak, Ramanujan graphs, *Combinatorica* 8: 261-277, 1988.
- Luczak, T. (1990a). Component behavior near the critical point of the random graph process, *Random Structures Algorithms* 1(3): 287-310.
- Luczak, T. (1990b). Component behavior near the critical point of the random graph process, *Random Structures Algorithms* 1: 287-310.
- Luczak, T. (1991). A note on the sharp concentration of the chromatic number of random graphs, *Combinatorica* 11: 295-297.
- MacWilliams, F. J. and Sloane, N. J. A. (1977). *The Theory of Error Correcting Codes*, North Holland, Amsterdam.

Mani-Levitska, P. and Pach, J. (1988). Decomposition problems for multiple coverings with unit balls, manuscript.

Margulis, G. A. (1973). Explicit constructions of concentrators, *Problemy Peredachi Informatsii* 9: 71-80. (In Russian). English translation in *Problems of Information Transmission* 9: 325-332.

Margulis, G. A. (1988). Explicit group-theoretical constructions of combinatorial schemes and their application to the design of expanders and superconcentrators, *Problemy Peredachi Informatsii* 24: 51-60. (In Russian). English translation in *Problems of Information Transmission* 24: 39-46, 1988.

Marica, J. and Schonheim, J. (1969). Differences of sets and a problem of Graham, *Can. Math. Bull.* 12: 635-637.

Matoušek, J. (1997). On discrepancy bounds via dual shatter function, *Mathematika* 44(1): 42-49.

Matoušek, J., Welzl, E. and Wernisch, L. (1993). Discrepancy and approximation for bounded VC dimension, *Combinatorica* 13: 455-466.

Matula, D. W. (1976). The largest clique size in a random graph, Technical report, Southern Methodist University, Dallas, TX.

Maurey, B. (1979). Construction de suites symétriques, *C. R. Acad. Sci. Paris* 288: 679-681.

Men'shikov, M. V. (1986). Coincidence of critical points in percolation problems, *Dokl. Akad. Nauk SSSR* 288(6): 1308-1311.

Milman, V. D. and Schechtman, G. (1986). *Asymptotic Theory of Finite Dimensional Normed Spaces*, Vol. 1200 of *Lecture Notes in Mathematics*, Springer Verlag, New York.

Molloy, M. and Reed, B. (1999). Graph colouring via the probabilistic method, *Graph Theory and Combinatorial Biology (Balatonlelle, 1996)*, Vol. 7 of *Bolyai Soc. Math. Stud.*, János Bolyai Math. Soc., Budapest, pp. 125-155.

Moon, J. W. (1968). *Topics on Tournaments*, Holt, Reinhart and Winston, New York.

Nakayama, A. and Peroche, B. (1987). Linear arboricity of digraphs, *Networks* 17: 39-53.

Naor, J. and Naor, M. (1990). Small-bias probability spaces: efficient constructions and applications, *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*, ACM Press, New York, pp. 213-223.

Nilli, A. (1991). On the second eigenvalue of a graph, *Discrete Math.* 91: 207-210.

Pach, J. and Agarwal, P. K. (1995). *Combinatorial Geometry*, Wiley, Hoboken, NJ.

Pach, J. and Woeginger, G. (1990). Some new bounds for ϵ -nets, *Proceedings of the 6 th Annual Symposium on Computational Geometry*, ACM Press, New York, pp. 10-15.

Paturi, R. and Simon, J. (1984). Probabilistic communication complexity, *Proceedings of the 25 th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, New York, pp. 118-126.

Paul, W. J. (1977). A $2.5 n$ lower bound on the combinational complexity of boolean functions, *SIAM J. Computing* 6: 427-443.

Pinsker, M. (1973). On the complexity of a concentrator, *Proceedings of the 7 th International Teletraffic Conference*, Stockholm, pp. 318/1-318/4.

Pippenger, N. and Spencer, J. H. (1989). Asymptotic behaviour of the chromatic index for hypergraphs, *J. Combin. Theory, Ser. A* 51: 24-42.

Rabin, M. O. (1980). Probabilistic algorithms for testing primality, *J. Number Theory* 12: 128-138.

Radhakrishnan, J. and Srinivasan, A. (2000). Improved bounds and algorithms for hypergraph two-coloring, *Random Structures Algorithms* 16: 4-32.

Raghavan, P. (1988). Probabilistic construction of deterministic algorithms: approximating packing integer programs, *J. Comput. Sys. Sci.* 37: 130-143.

- Ramsey, F. P. (1929). On a problem of formal logic, *Proc. London Math. Soc.* 30(2): 264-286.
- Razborov, A. A. (1985). Lower bounds on the monotone complexity of some boolean functions, *Dokl. Akad. Nauk SSSR* 281(4): 798-801. (In Russian). English translation in *Sov. Math. Dokl.* 31: 354-357.
- Razborov, A. A. (1987). Lower bounds on the size of bounded depth networks over a complete basis with logical addition, *Mat. Zametki* 41(4): 598-607. (In Russian). English translation in *Math. Notes Acad. Sci. USSR* 41(4):333-338.
- Rödl, V. (1985). On a packing and covering problem, *Eur. J. Combinatorics* 6: 69-78.
- Rödl, V. and Duke, R. A. (1985). On graphs with small subgraphs of large chromatic number, *Graphs Combin.* 1(1): 91-96.
- Ron, D. (2001). Property testing, in: *Handbook of Randomized Computing*, Vol. I, II, Vol. 9 of *Comb. Optim.*, Kluwer Academic, Dordrecht, pp. 597-649.
- Rubinfeld, R. and Sudan, M. (1996). Robust characterizations of polynomials with applications to program testing, *SIAM J. Comput.* 25(2): 252-271.
- Ruzsa, I. Z. and Szemerédi, E. (1978). Triple systems with no six points carrying three triangles, *Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976)*, Vol. II, Vol. 18 of *Colloq. Math. Soc. János Bolyai*, North-Holland, Amsterdam, pp. 939-945.
- Sauer, N. (1972). On the density of families of sets, *J. Combin. Theory, Ser. A* 13: 145-147.
- Schrijver, A. (1978). A short proof of Minc's conjecture, *J. Combin. Theory, Ser. A* 25: 80-83.
- Shamir, E. and Spencer, J. H. (1987). Sharp concentration of the chromatic number in random graphs $G_{n,p}$, *Combinatorica* 7(1): 121-129.
- Shearer, J. B. (1983). A note on the independence number of triangle-free graphs, *Discrete Math.* 46: 83-87.
- Shearer, J. B. (1985). On a problem of Spencer, *Combinatorica* 5: 241-245.
- Shearer, J. B. (1995). On the independence number of sparse graphs, *Random Structures Algorithms* 7: 269-271.
- Shelah, S. and Spencer, J. H. (1988). Zero-one laws for sparse random graphs, *J. Am. Math. Soc.* 1: 97-115.
- Shepp, L. A. (1982). The XYZ-conjecture and the *FKG*-inequality, *Ann. Probab.* 10: 824-827.
- Slade, G. (2006). *The Lace Expansion and Its Applications*, Vol. 1879 of *Lecture Notes in Mathematics*, Springer-Verlag, Berlin. Lectures from the 34th Summer School on Probability Theory held in Saint-Flour, July 6-24, 2004, edited and with a foreword by Jean Picard.
- Smolensky, R. (1987). Algebraic methods in the theory of lower bounds for boolean circuit complexity, *Proceedings of the 19th Annual ACM Symposium on Theory of Computing (STOC)*, ACM Press, New York, pp. 77-82.
- Spencer, J. H. (1977). Asymptotic lower bounds for Ramsey functions, *Discrete Math.* 20: 69-76.
- Spencer, J. H. (1985a). Probabilistic methods, *Graphs Combin.* 1: 357-382.
- Spencer, J. H. (1985b). Six standard deviations suffice, *Trans. Am. Math. Soc.* 289: 679-706.
- Spencer, J. H. (1987). *Ten Lectures on the Probabilistic Method*, SIAM, Philadelphia.
- Spencer, J. H. (1990a). Counting extensions, *J. Combin. Theory, Ser. A* 55: 247-255.
- Spencer, J. H. (1990b). Threshold functions for extension statements, *J. Combin. Theory, Ser. A* 53: 286-305.
- Spencer, J. H. (1995). Asymptotic packing via a branching process, *Random Structures Algorithms* 7: 167-172.
- Subbotovskaya, B. A. (1961). Realizations of linear functions by formulas using $+$, $\cdot$, $-$, *Dokl. Akad. Nauk SSSR* 136(3): 553-555. (In Russian). English translation in *Sov. Math.*

Dokl., 2: 110-112.

Suen, W. C. (1990). A correlation inequality and a Poisson limit theorem for nonoverlapping balanced subgraphs of a random graph, *Random Structures Algorithms* 1: 231-242.

Székely, L. (1997). Crossing numbers and hard Erdős problems in discrete geometry, *Combin. Probab. Comput.* 6: 353-358.

Szele, T. (1943). Kombinatorikai vizsgálatok az irányított teljes gráffal, *Kapcsolatban, Mt. Fiz. Lapok* 50: 223-256. For a German translation see: T. Szele, *Publ. Math. Debrecen* 13: 145-168, 1966.

Szemerédi, E. (1978). Regular partitions of graphs, *Problèmes combinatoires et théorie des graphes* (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976), Vol. 260 of *Colloq. Internat. CNRS*, CNRS, Paris, pp. 399-401.

Talagrand, M. (1996). Concentration of measures and isoperimetric inequalities in product spaces, *Publications Mathématiques de l'I.H.E.S.* 81: 73-205.

Tanner, R. M. (1984). Explicit construction of concentrators from generalized N gons, *SIAM J. Alg. Discrete Meth.* 5: 287-293.

Tarjan, R. E. (1983). *Data Structures and Network Algorithms*, SIAM, Philadelphia.

Thomason, A. (1987). Pseudo-random graphs, *Ann. Discrete Math.* 33: 307-331.

Turán, P. (1934). On a theorem of Hardy and Ramanujan, *J. London Math Soc.* 9: 274-276.

Turán, P. (1941). On an extremal problem in graph theory, *Mat. Fiz. Lapok* 48: 436-452.

Valtr, P. (1995). On the minimum number of empty polygons in planar point sets, *Stud. Sci. Math. Hung.* 30: 155-163.

van der Hofstad, R. and Spencer, J. H. (2006). Counting connected graphs asymptotically, *Eur. J. Combin.* 27(8): 1294-1320.

Vapnik, V. N. and Chervonenkis, A. Y. (1971). On the uniform convergence of relative frequencies of events to their probabilities, *Theory Probab. Appl.* 16: 264-280.

Wegener, I. (1987). *The Complexity of Boolean Functions*, Wiley-Teubner, Hoboken, NJ.

Weil, A. (1948). Sur les courbes algébriques et les variétés qui s'en déduisent, *Actualités Sci. Ind.*, no. 1041. iv+85pp.

Wendel, J. G. (1962). A problem in geometric probability, *Math. Scand.* 11: 109-111.

Wright, E. M. (1977). The number of connected sparsely edged graphs, *J. Graph Theory* 1: 317-330.

Yao, A. C. (1985). Separating the polynomial-time hierarchy by oracles, *Proceedings of the 26th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, New York, pp. 1-10.

Part VI

Appendix

A Introduction and Overview

Приведем соображения о подходе к теории вероятностей.

A.1 Extra Motivation

(?? write later)

Востребованность в аналитических работах (?)

(про то, что на собеседованиях нужно, а также аналитики её по идее используют)

Множество приложений

(? пока хз, каких, потом пропишу)

Множество приложений в программировании

(? пока хз, каких, потом пропишу)

Удивительные факты

(позже пропишу)

Результаты теории вероятности

Обсудим, какие способности и знания дает нам теория вероятности.
(позже пропишу)

Мотивация к разделам теории вероятностей

(потом пропишу, пока не так вижу их, да и тут разделов не столько много, сколько в других записях)

A.2 The Mindset of Professionals

A.2.1 Attitude to Probability Theory

об использовании теории вероятностей в жизни

(?? пока не знаю, насколько на самом деле теория вероятностей используется в жизненных задачах, ибо её знаю пока плохо)

О важности алгебры и теории меры

(??? я хз, насколько она нужна, пока изучаю, потом пойму и пропишу. мб нужна, мб и нет.)

Способы догадаться до главных идей теории вероятностей

незаменимая часть нормального понимания теории вероятностей.
хз

Наиболее прикладные темы теории вероятностей (?)

обсудим, какие темы используются более всего.

О взгляде, что теория вероятностей большая и глубокая

(действительно, много книг про неё написано, так что я придерживаюсь такого взгляда. но потом обсудим.)

О взгляде, что теория вероятностей проста и поверхностна

(так некоторые думают, обсудим этот подход. пока хз, мне и в ней хватает работы)

A.2.2 Features of Effective Learning

Необходимые темы для изучения теории вероятностей

Необходимо хорошее знакомство с математическим анализом

(потому что очень много и теорем из него и в целом предмет берется ровно так же, как и матан. потом распишу подробнее, пока еще сам не освоил кучу лем и кучу переходов.)

Необходимо хорошее знакомство с теорией меры

Тоже без минимальных заготовок на неё отлетишь. Однако если нужно будет, её по ходу дела доделывать правильнее, чем готовить изначально.

способы изучения теории вероятностей

важно собрать основы теории, потому что они постоянно нужны

и это не так много, мб даже страниц 50 только - самые основы.
но они не простые.
хотя и не сложные.
сесть просто нужно, и сделать их.
а дальше всё и заживет.

как можно скорее на задачи выходим

в основах можно долго тормозить, поэтому важно скорее выйти на задачи.
а с задачами - есть устойчивость и база, там уже дело техники - выйти на приложения.

A.3 Acknowledgements

Currently, no one except me has worked on the sections of this note (with the exception of sections taken from books).

A.4 Literature

Литературы аж навалом, так что важно понимать основную и не урываться в неё вечно.

A.4.1 Main Literature

Basic Educational Literature

Самарова С.С. Учебные пособие для студентов МФТИ по теории вероятностей.
Прекрасные пособия, для начала даже не думаю, что еще что-то нужно. Гуглятся.

В. А. Колемаев, В. Н. Калинина и др. Теория Вероятностей В Примерах И Задачах
Крутая книга, где компактно всё написано,

[1] Зубков А. М. Конспект лекций по теории вероятностей

Небольшой сборник лекций, где достаточно теории, чтобы с неё начать. Однако правильнее вообще забыть на это, потому что куда полезнее начать задачи решать.

[4] Горяйнов

добавка в помощь для Зубкова

С.Е. Демин, Е.Л. Демина Теория вероятностей в трех частях. Часть 1. Случайные события

Хорошая книга с многими хорошими задачами разобрана, посмотрю потом.

М.Е. Жуковский Теория вероятностей. Конспект лекций (записал Александр Марков, 30.01.2017)

60 страниц компактного конспекта, но похоже, что там больше бесполезная теория, чем полезные методы. Гуглится, чуть смотрел.

Books With Many Solved Problems

Чистяков Сборник задач

Типичный задачник, можно месяц о нем думать.

Кондратенко А. Е. задачи с зачетов по теории вероятностей

13 страниц задач-решений, полезный файл.

Математическая помощь

[5] Гусев мера

справочник по теории меры. не уверен, насколько она нужна.

A.4.2 Extra Literature

Другая обучающая

<https://www.maths.ed.ac.uk/~jmf/Teaching/mi4a.html> José Figuerola-O’Farrill 21
Lectures on Probability and Random Processes

Nice great lectures, nice summary of typical topics.

Особо профессиональная и подробная литература

[3] Гнеденко Б. В. курс теории вероятностей
(!!!!) вроде хороший вмеру подробный курс. Уже завершаю по Зубкову и Горяйнову, так что его скоро и начну и он очень сильно повлияет на мою структуру!!! Пока еще пара шагов до него остается.

Феллер Введение в теорию вероятностей

Амосов сказал, что хорошая большая книга.

Статьи при интересные модели теории вероятностей

[6] Семенов А. Г. О случайном блуждании «пьяной компании»
помню, что это интересная статья о континуальном интеграле и случайных блужданиях, соберу эти теории - прочитаю её, ибо интересно очень.

По отдельным конструкциям

(пока их не смотрел, но помню, что супер интересная, чуть что посмотрю)

[2] Алон Н. Спенсер Дж. Вероятностный метод
топ книга, но сперва азы нормально прописать нужно!! профессиональные тут всякие теоремы и подходы. Мне не до этого пока.

Про квантовую теорию вероятностей

(просто знаю, что такое есть, курс Амосова ни о чем был, так что забить можно на копание в этом. чуть что - здесь укажу эти книги.)

По основам случайных процессов

Жуковский МФТИ лекции по случайным процессам
помню, что такая есть, потом добавлю тоже мб ей, пока что не актуально.
Лекции по случайным процессам Под редакцией А. В. Гасникова
нашел в интернете, хотелось бы тоже взять их, но уже после многих других книг, что выше я писал, начну их.

По основам математической статистики

Боровков математическая статистика
помню такую книгу, отдельная большая история. мб когда-то добавлю какие-то основы из неё в раздел про основы мат статистики, но пока не до неё.

A.4.3 Literature About Applications

Про жизненные применения теории вероятностей

Существует множество популярных или упрощенно прикладных книг по теорверу, многие из которых обсудим ниже, они же вписаны в многие части (будут).

(Каждое приложение - отдельная запись, здесь основную литературу по ним приведем.)

Тарасов Закономерности окружающего мира. Случайность, необходимость, вероятность

наверное, такая о жизни книга. как основы сделаю - дополню ей много баланса. Очень круто она в помощь для прикладного понимания теорвера.

Мандельброт Фракталы случай и финансы

очень наверное крутые связки там раскрыты. прямо очень интересно, и как раз куча приложений будет начато рассматриваться. так что очень крутая книга, дойти бы!

По приложениям в программировании и анализе данных

тоже много книг в таком направлении, так что обсудим и их тут.

A.4.4 History in a Nutshell

A.5 Puzzles for Different Situations

Обсудим в порядке интересности задачки и вопросы

A.5.1 Puzzles for Interesting Discussion

A.5.2 Puzzles About Everyday Life

(???? what fun questions can arise about this subject in everyday life?)

A.5.3 Selected Fun Typical Problems

(all this should also be in the problem-solution section!)

B Mathematics for Probability Theory

Обсудим все математические конструкции, применяющиеся в теории вероятностей.
(Многое потом добавлю, пока не успел. по актуальности соберу)

B.1 Another Mathematics for Probability Theory

(странно, что до сих пор ничего про нее не написано)

B.1.1 Constructions of the Theory of Measure

(по потребности буду их добавлять)

B.1.2 Convergence Theorems

При доказательстве различных теорем теории вероятностей часто используются два утверждения, фактически относящиеся к математическому анализу: теоремы о монотонной и мажорируемой сходимости. (в отдельный раздел потом перенесу)

Теорема о монотонной сходимости

Если последовательность $\{\xi_n\}$ случайных величин такова, что $0 \leq \xi_n \uparrow \xi$ при $n \uparrow \infty$, то

$$\lim_{n \rightarrow \infty} M\xi_n = M\xi$$

Эта теорема отличается от аналогичного утверждения при обосновании корректности определения математического ожидания тем, что случайные величины ξ_n не предполагаются простыми.

Доказательство. В силу свойства монотонности математического ожидания из условия $0 \leq \xi_n \leq \xi$ следует, что $0 \leq M\xi_n \leq M\xi$ и

$$\lim_{n \rightarrow \infty} M\xi_n \leq M\xi$$

Введем монотонные последовательности простых случайных величин $\{\xi_{nk}\}$ так, что $0 \leq \xi_{nk} \uparrow \xi_n$ при $k \uparrow \infty$. Случайные величины

$$\eta_k = \max_{1 \leq n \leq k} \xi_{nk} \leq \max_{1 \leq n \leq k} \xi_n = \xi_k$$

тоже являются простыми, причем последовательность η_k не убывает:

$$0 \leq \eta_k = \max \{\xi_{1k}, \xi_{2k}, \dots, \xi_{kk}\} \leq \max \{\xi_{1k+1}, \xi_{2k+1}, \dots, \xi_{k+1k+1}\} = \eta_{k+1}$$

Значит, для всех $\omega \in \Omega$ существует $\lim_{k \rightarrow \infty} \eta_k = \eta$. Так как $\eta_k \leq \xi_k$ при всех k и η_k - простые случайные величины, то

$$\lim_{k \rightarrow \infty} \eta_k = M\eta \leq \lim_{k \rightarrow \infty} M\xi_k$$

С другой стороны, $\xi_{nk} \leq \eta_k \leq \eta$; переходя к пределу по $k \rightarrow \infty$, находим: $\xi_n \leq \eta$ при всех n , откуда $\xi \leq \eta$ и $M\xi \leq M\eta$. Таким образом,

$$M\xi \leq M\eta \leq \lim_{k \rightarrow \infty} M\xi_k \leq M\xi$$

и Теорема доказана.

Без условия монотонного возрастания монотонной сходимости может не быть (?????)

Без условия монотонного возрастания утверждение теоремы может быть неправильным.

пример монотонно убывающей последовательности неотрицательных случайных величин с бесконечными математическими ожиданиями, сходящейся к случайной величине, тождественно равной 0.

(ХЗ, сказали построить, но я тупой)

Следствия теоремы о монотонной сходимости

Если случайные величины ξ_n неотрицательны, то

$$M \sum_{n \geq 1} \xi_n = \sum_{n \geq 1} M\xi_n$$

Для доказательства достаточно рассмотреть монотонно не убывающую последовательность случайных величин $\eta_k = \sum_{n=1}^k \xi_n$ и применить теорему о монотонной сходимости.

Следствие 2. Если M η конечно и последовательность событий $A_n \downarrow \emptyset$ при $n \uparrow \infty$, то

$$\lim_{n \rightarrow \infty} M\eta I \{A_n\} = 0$$

Доказательство. По условию $M|\eta| < \infty$. Представим $|\eta|$ в виде $|\eta| = \eta_n + \eta'_n$, где $\eta_n = |\eta| I \{A_n\}$, $\eta'_n = |\eta| I \{\overline{A_n}\}$. Тогда

$$M|\eta| = M\eta_n + M\eta'_n$$

Так как $\overline{A_n} \uparrow \Omega$, то $0 \leq \eta'_n \uparrow |\eta|$ и по теореме о монотонной сходимости

$$M\eta'_n \uparrow M|\eta|, \text{ значит } M\eta_n \rightarrow 0 \text{ при } n \rightarrow \infty$$

Но тогда $|M\eta I \{A_n\}| \leq M|\eta I \{A_n\}| = M\eta_n \rightarrow 0$.

Теорема Лебега о мажорируемой сходимости

(?? зачем она вообще в мат ожидании??)

Если $\lim_{n \rightarrow \infty} \xi_n(\omega) = \xi(\omega)$ и $|\xi_n(\omega)| \leq \eta$ при всех $\omega \in \Omega$, причем $M\eta < \infty$, то

$$\lim_{n \rightarrow \infty} M\xi_n = M\xi$$

Доказательство.

Выберем произвольно $\varepsilon > 0$ и рассмотрим последовательность событий

$$A_n = \left\{ \omega \in \Omega : \sup_{m > n} |\xi_m(\omega) - \xi(\omega)| < \varepsilon \right\}, \quad n = 1, 2, \dots$$

Из условия теоремы следует, что $\overline{A_n} \downarrow \emptyset$ при $n \uparrow \infty$ для любого $\varepsilon > 0$. Тогда, как в следствии 2,

$$\xi_n = \xi_n I\{A_n\} + \xi_n I\{\overline{A_n}\}$$

И

$$\begin{aligned} \xi I\{A_n\} - \varepsilon &\leq \xi_n I\{A_n\} \leq \xi I\{A_n\} + \varepsilon \\ -\eta I\{\overline{A_n}\} &\leq \xi_n I\{\overline{A_n}\} \leq \eta I\{A_n\} \end{aligned}$$

Складывая почленно эти неравенства с учетом равенства $\xi I\{A_n\} = \xi - \xi I\{\overline{A_n}\}$ получаем

$$\xi - \xi I\{\overline{A_n}\} - \varepsilon - \eta I\{\overline{A_n}\} \leq \xi_n \leq \xi - \xi I\{\overline{A_n}\} + \varepsilon + \eta I\{\overline{A_n}\}$$

далее вычислим математическое ожидание от всех частей неравенства:

$$M\xi - \varepsilon - 2M\eta I\{\overline{A_n}\} \leq \xi_n \leq M\xi + \varepsilon + 2M\eta I\{\overline{A_n}\}$$

Переходя здесь к пределу по $n \rightarrow \infty$ и учитывая следствие 2, получим:

$$M\xi - \varepsilon \leq \liminf_{n \rightarrow \infty} \xi_n \leq \limsup_{n \rightarrow \infty} \xi_n \leq M\xi + \varepsilon$$

Ввиду того, что $\varepsilon > 0$ можно выбирать произвольно малым, теорема доказана.

B.1.3 Calculation of Integrals

(?? с этим могут быть проблемы или особенности)

B.2 Elements of Discrete Mathematics**B.2.1 Math.-phys.**

остановимся на леммах теории множеств, которые часто нужны в задачах про вероятность. Укажем также, где они нужны.

B.2.2 Combinatorics

запишем подробно разные методы комбинаторики, которые нам нужны для задач на теорию вероятностей.

размещения это...

сочетания это..

типа там для каждого набора из k элементов с учетом порядка можно получить $k!$ разных наборов, меняя элементы местами. !! кстати, тут уже интересный момент, сколько перестановок? их $A_n^n = n!$. тут-то и пригождаются наши рассуждения.

круто, теперь с возвращением с учетом порядка: n^k . тоже

(место где рассуждаем про возвращение и без учета порядка) честно - пока пофиг, потом добыю, пока просто понимаю, что это есть.

B.3 Programming for Probability Theory (!!???)

(тоже часто нужно какие-то посчитать вещи, напишу тут эти методы потом.)

B.4 Math

B.4.1 A. 1 Sets

A set is a Many that allows itself to be thought of as a One.

- Georg Cantor

A set is a collection of objects. The objects can be anything: numbers, people, cats, courses, even other sets! The language of sets allows us to talk precisely about events. If S is a set, then the notation $x \in S$ indicates that x is an element or member of the set S (and $x \notin S$ indicates that x is not in S). We can think of the set as a club, with precisely defined criteria for membership. For example:

1. $\{1, 3, 5, 7, \dots\}$ is the set of all odd numbers;
2. $\mathbb{R}$ is the set of all real numbers;
3. $\{\text{Worf, Jack, Tobey}\}$ is the set of Joe's cats;
4. $[3, 7]$ is the closed interval consisting of all real numbers between 3 and 7;
5. $\{HH, HT, TH, TT\}$ is the set of all possible outcomes if a coin is flipped twice (where, for example, HT means the first flip lands Heads and the second lands Tails).

To describe a set (when it's tedious or impossible to list out its elements), we can give a rule that says whether each possible object is or isn't in the set. For example, $\{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$ is the disk of radius 1 in the plane $\mathbb{R}^2$, centered at the origin. In this expression, the ":" is read as "such that".

A.1.1 The Empty set

The smallest set, which is both subtle and important, is the empty set, which is the set that has no elements whatsoever. It is denoted by $\emptyset$ or by $\{\}$. Make sure not to confuse $\emptyset$ with $\{\emptyset\}$. The former has no elements, while the latter has one element. If we visualize the empty set as an empty paper bag, then we can visualize $\{\emptyset\}$ as a paper bag inside of a paper bag.

A.1.2 Subsets

If A and B are sets, then we say A is a subset of B (and write $A \subseteq B$) if every element of A is also an element of B . For example, the set of all integers is a subset of the set of all real numbers. It is always true that $\emptyset$ and A itself are subsets of A ; these are the extreme cases for subsets. A general strategy for showing that $A \subseteq B$ is to let x be an arbitrary element of A , and then show that x must also be an element of B . A general strategy for showing that $A = B$ for two sets A and B is to show that each is a subset of the other.

A.1.3 Unions, Intersections, and Complements

The union of two sets A and B , written as $A \cup B$, is the set of all objects that are in A or B (or both). The intersection of A and B , written as $A \cap B$, is the set of all objects that are in both A and B . We say that A and B are disjoint if $A \cap B = \emptyset$. For n sets $A_1, \dots, A_n$, the union $A_1 \cup A_2 \cup \dots \cup A_n$ is the set of all objects that are in at least one of the A_j 's, while the intersection $A_1 \cap A_2 \cap \dots \cap A_n$ is the set of all objects that are in all of the A_j 's.

In many applications, all the sets we're working with are subsets of some set S (in probability, this may be the set of all possible outcomes of some experiment). When S is clear from the context, we define the complement of a set A to be the set of all objects in S that are not in A ; this is denoted by A^c .

Unions, intersections, and complements can be visualized easily using Venn diagrams, such as the one shown below. The union is the entire shaded region, while the intersection is the football-shaped region of points that are in both A and B . The complement of A is all points in the rectangle that are outside of A .




FIGURE A. 1

A Venn diagram. Note that the area of the region $A \cup B$ is the area of A plus the area of B , minus the area of $A \cap B$ (this is a basic form of the inclusion-exclusion principle).

De Morgan's laws give an elegant, useful duality between unions and intersections:

$$\begin{aligned}(A_1 \cup A_2 \cup \dots \cup A_n)^c &= A_1^c \cap A_2^c \cap \dots \cap A_n^c \\ (A_1 \cap A_2 \cap \dots \cap A_n)^c &= A_1^c \cup A_2^c \cup \dots \cup A_n^c\end{aligned}$$

It is much more important to understand De Morgan's laws than to memorize them! The first law says that not being in at least one of the A_j is the same thing as not being in A_1 , nor being in A_2 , nor being in A_3 , etc.

For example, let A_j be the set of people who like the j th Star Wars prequel, for $j \in \{1, 2, 3\}$. Then $(A_1 \cup A_2 \cup A_3)^c$ is the set of people for whom it is not the case that they like at least one of the prequels, but that is the same as $A_1^c \cap A_2^c \cap A_3^c$, the set of people who don't like The Phantom Menace, don't like Attack of the Clones, and don't like Revenge of the Sith.

The second law says that not being in all of the A_j is the same thing as being outside at least one of the A_j . For example, let the A_j be defined as in the previous paragraph. If it is not the case that you like all of the Star Wars prequels (making you a member of the set $(A_1 \cap A_2 \cap A_3)^c$), then there must be at least one prequel that you don't like (making you a member of the set $A_1^c \cup A_2^c \cup A_3^c$), and vice versa.

Proving the following facts about sets (not just drawing Venn diagrams, though they are very helpful for building intuition) is good practice:

1. $A \cap B$ and $A \cap B^c$ are disjoint, with $(A \cap B) \cup (A \cap B^c) = A$.
2. $A \cap B = A$ if and only if $A \subseteq B$.

3. $A \subseteq B$ if and only if $B^c \subseteq A^c$.

A.1.4 Partitions

A collection of subsets $A_1, \dots, A_n$ of a set S is a partition of S if $A_1 \cup \dots \cup A_n = S$ and $A_i \cap A_j = \emptyset$ for all $i \neq j$. In words, a partition of a set is a collection of disjoint subsets whose union is the entire set. For example, the set of even numbers $\{0, 2, 4, \dots\}$ and the set of odd numbers $\{1, 3, 5, \dots\}$ form a partition of the set of nonnegative integers.

A.1.5 Cardinality

A set may be finite or infinite. If A is a finite set, we write $|A|$ for the number of elements in A , which is called its size or cardinality. For example, $|\{2, 4, 6, 8, 10\}| = 5$ since there are 5 elements in this set. A very useful fact is that A and B are finite sets, then

$$|A \cup B| = |A| + |B| - |A \cap B|.$$

This is a form of the inclusion-exclusion theorem from Chapter 1. It says that to count how many elements are in the union of A and B , we can add the separate counts for each, and then adjust for the fact that we have double-counted the elements (if any) that are in both A and B .

Two sets A and B are said to have the same size or same cardinality if they can be put into one-to-one correspondence, i.e., if each element of A can be paired up with exactly one element of B , with no unpaired elements in either set. We say that A is smaller than B if there is not a one-to-one correspondence between A and B , but there is a one-to-one correspondence between A and some subset of B .

For example, suppose that we want to count the number of people in a movie theater with 100 seats. Assume that no one in the theater is standing, and no seat has more than one person in it. The obvious thing to do is to go around counting people one by one (though it's surprisingly easy to miss someone or accidentally count someone twice). But if every seat is occupied, then a much easier method is to note that there must be 100 people, since there are 100 seats and there is a one-to-one correspondence between people and seats. If some seats are empty, then there must be fewer than 100 people there.

This idea of looking at one-to-one correspondences makes sense both for finite and for infinite sets. Consider the perfect squares $1^2, 2^2, 3^2, \dots$. Galileo pointed out the paradoxical result that on the one hand it seems like there are fewer perfect squares than positive integers (since every perfect square is a positive integer, but lots of positive integers aren't perfect squares), but on the other hand it seems like these two sets have the same size since they can be put into one-to-one correspondence: pair 1^2 with 1, pair 2^2 with 2, pair 3^2 with 3, etc.

The resolution of Galileo's paradox is to realize that intuitions about finite sets don't necessarily carry over to infinite sets. By definition, the set of all perfect squares and the set of all positive integers do have the same size. Another famous example of this is Hilbert's hotel. For any hotel in the real world the number of rooms is finite. If every room is occupied, there is no way to accommodate more guests, other than by cramming more people into already occupied rooms.

Now consider an imaginary hotel with an infinite sequence of rooms, numbered $1, 2, 3, \dots$. Assume that all the rooms are occupied, and that a weary traveler arrives, looking for a room. Can the hotel give the traveler a room, without leaving any of the current guests without a room? Yes, one way is to have the guest in room n move to room $n + 1$, for all $n = 1, 2, 3, \dots$. This frees up room 1, so the traveler can stay there.

What if infinitely many travelers arrive at the same time, such that their cardinality is the same as that of the positive integers (so we can label the travelers as traveler 1, traveler 2, ...) ?

The hotel could fit them in one by one by repeating the above procedure over and over again, but it would take forever (infinitely many moves) to accommodate everyone, and it would be bad for business to make the current guests keep moving over and over again. Can the room assignments be updated just once

so that everyone has a room? Yes, one way is to have the guest in room n move to room $2n$ for all $n = 1, 2, 3, \dots$, and then have traveler n move into room $2n - 1$. In this way, the current guests occupy all the even-numbered rooms, and the new guests occupy all the odd-numbered rooms.

An infinite set is called countably infinite if it has the same cardinality as the set of all positive integers. A set is called countable if it is finite or countably infinite, and uncountable otherwise. The mathematician Cantor showed that not all infinite sets are the same size. In particular, the set of all real numbers is uncountable, as is any interval in the real line of positive length.

A.1.6 Cartesian Product

The Cartesian product of two sets A and B is the set

$$A \times B = \{(a, b) : a \in A, b \in B\}.$$

For example, $[0, 1] \times [0, 1]$ is the square $\{(x, y) : x, y \in [0, 1]\}$, and $\mathbb{R} \times \mathbb{R} = \mathbb{R}^2$ is two-dimensional Euclidean space.

B.4.2 A. 2 Functions

Let A and B be sets. A function from A to B is a deterministic rule that, given an element of A as input, provides an element of B as an output. That is, a function from A to B is a machine that takes an x in A and "maps" it to some y in B . Different x 's can map to the same y , but each x only maps to one y . Here A is called the domain and B is called the target. The notation $f : A \rightarrow B$ says that f is a function mapping A into B . The range of f is $\{y \in B : f(x) = y \text{ for some } x \in A\}$.

It is important to distinguish between f (the function) and $f(x)$ (the value of the function when evaluated at x). That is, f is a rule, while $f(x)$ is a number for each number x . The function g given by $g(x) = e^{-x^2/2}$ is exactly the same as the function g given by $g(t) = e^{-t^2/2}$; what matters is the rule, not the name we use for the input. For this example, the rule is that x gets mapped to $e^{-x^2/2}$; this rule is also denoted by $x \mapsto e^{-x^2/2}$.

A function f from the real line to the real line is continuous if $f(x) \rightarrow f(a)$ as $x \rightarrow a$, for any value of a . It is called right-continuous if this is true when approaching from the right, i.e., $f(x) \rightarrow f(a)$ as $x \rightarrow a$ while ranging over values with $x > a$.

The set A may not be a set of numbers. In probability, it is extremely common and useful to consider functions whose domains are the set of all possible outcomes of some experiment. It may be difficult to write down a formula for the function, but the function is still valid as long as it's defined unambiguously.

A.2.1 One-to-one Functions

Let f be a function from A to B . Then f is a one-to-one function if $f(x) \neq f(y)$ whenever $x \neq y$. That is, any two distinct inputs in A get mapped to two distinct outputs in B ; for each y in B , there can be at most one x in A that maps to it.

Let f be a one-to-one function from A to B , and let C be the range of f , i.e.,

$$C = \{b \in B : f(a) = b \text{ for some } a \in A\}.$$

Then there is an inverse function $f^{-1} : C \rightarrow A$, defined by letting $f^{-1}(y)$ be the unique element $x \in A$ such that $f(x) = y$. A function that has an inverse is called an invertible function.

For example, let $f(x) = x^2$ for all real x . This is not a one-to-one function since, for example, $f(3) = f(-3)$. But now assume instead that the domain of f is chosen to be $[0, \infty)$, so we are defining f as a function from $[0, \infty)$ to $[0, \infty)$. Then f is one-to-one, and its inverse function is given by $f^{-1}(y) = \sqrt{y}$ for all $y \in [0, \infty)$.

The pigeonhole principle says that if A and B are finite sets with $|A| > |B|$, then there does not exist a one-to-one function from A to B . In other words, if k objects are placed into n boxes, where $k > n$, then there must be at least one box that contains more than one object.

For example, in the birthday problem from Chapter 1 (assuming there are 365 days in a year), with 366 or more people there is guaranteed to be at least one birthday match. There is no one-to-one function assigning a birthday to each of the people, when there are more people than there are days in a year.

A.2.2 Increasing and Decreasing Functions

Let $f : A \rightarrow \mathbb{R}$, where A is a set of real numbers. Then f is an increasing function if $x \leq y$ implies $f(x) \leq f(y)$ (for all $x, y \in A$). Note that this definition allows there to be regions where f is flat, e.g., the constant function that is equal to 42 everywhere is an increasing function, and any CDF (see Section 3.6) is an increasing function. We say that f is strictly increasing if $x < y$ implies $f(x) < f(y)$. For example, $f : \mathbb{R} \rightarrow \mathbb{R}$ with $f(x) = x^3$ is a strictly increasing function.

Similarly, f is a decreasing function if $x \leq y$ implies $f(x) \geq f(y)$, and is a strictly decreasing function if $x < y$ implies $f(x) > f(y)$. For example, $f : (0, \infty) \rightarrow (0, \infty)$ with $f(x) = 1/x$ is a strictly decreasing function.

A monotone function is a function that is either increasing or decreasing. A strictly monotone function is a function that is either strictly increasing or strictly decreasing. Note that any strictly monotone function is one-to-one.

A.2.3 Even and odd Functions

Let f be a function from $\mathbb{R}$ to $\mathbb{R}$. We say f is an even function if $f(x) = f(-x)$ for all x , and we say f is an odd function if $-f(x) = f(-x)$ for all x . If neither of these conditions is satisfied, then f is neither even nor odd. Figure A. 2 shows the graphs of two even functions and two odd functions.




FIGURE A. 2

Even and odd functions. The graphs on the left are even functions: $f(x) = x^2$ on the top and $f(x) = e^{-x^2/2}$ on the bottom. The graphs on the right are odd functions: $f(x) = x^3$ on the top and $f(x) = xe^{-x^2/2}$ on the bottom.

Even and odd functions have nice symmetry properties. The graph of an even function remains the same if you reflect it about the vertical axis, and the graph of an odd function remains the same if you rotate it 180 degrees around the origin.

Even functions have the property that for any a ,

$$\int_{-a}^a f(x)dx = 2 \int_0^a f(x)dx,$$

assuming the integral exists. This is because the area under the function from $-a$ to 0 is equal to the area under the function from 0 to a . Odd functions have the property that for any a ,

$$\int_{-a}^a f(x)dx = 0,$$

again assuming the integral exists. This is because the area under the function from $-a$ to 0 cancels the area under the function from 0 to a .

A.2.4 Convex and Concave Functions

A function g whose domain is an interval I is convex if

$$g(px_1 + (1-p)x_2) \leq pg(x_1) + (1-p)g(x_2)$$

for all $x_1, x_2 \in I$ and $p \in (0, 1)$. Geometrically, this says that if we draw a line segment connecting two points on the graph of g , then the line segment lies above the graph of g . If the derivative g' exists, then an equivalent definition is that every tangent line to the graph of g lies below the graph. If g'' exists, then an equivalent definition is that $g''(x) \geq 0$ for all $x \in I$. An example is shown in Figure A.3. A simple example is $g(x) = x^2$, whose second derivative is

$$g''(x) = 2 > 0.$$




FIGURE A. 3 Graph of a convex function g . We have $g''(x) \geq 0$. Any line segment connecting two points on the curve lies above the curve. Any tangent line lies below the curve.

A function g is concave if $-g$ is convex. If g'' exists, then g is concave if and only if $g''(x) \leq 0$ for all x in the domain. For example, $g(x) = \log(x)$ defines a concave function on $(0, \infty)$ since $g''(x) = -1/x^2 < 0$ for all $x \in (0, \infty)$.

A.2.5 Exponential and Logarithmic Functions

An exponential function is a function of the form $f(x) = a^x$ for some number $a > 0$. If $a > 1$, the function is increasing, and if $0 < a < 1$, the function is decreasing. The most common exponential function we'll work with is $f(x) = e^x$, and a very useful limit result to know is that

$$\left(1 + \frac{x}{n}\right)^n \rightarrow e^x$$

as $n \rightarrow \infty$, for any real number x . This has an interpretation in terms of a bank paying compound interest on a deposit: as compounding occurs more and more times per year, the growth rate approaches exponential growth.

Here are some properties of exponential functions:

1. $a^x a^y = a^{x+y}$.
2. $a^x b^x = (ab)^x$.
3. $(a^x)^y = a^{xy}$.

The inverse of an exponential function is a logarithmic function: for positive y , $\log_a y$ is defined to be the number x such that $a^x = y$. Throughout this book, when we write $\log y$ without explicitly specifying the base, we are referring to the natural logarithm (base e).

Here are some properties of logarithms:

1. $\log_a x + \log_a y = \log_a xy$.
2. $\log_a x^n = n \log_a x$.
3. $\log_a x = \frac{\log x}{\log a}$.

A.2.6 Floor Function and Ceiling Function

The floor function is defined by letting $\lfloor x \rfloor$ be the greatest integer less than or equal to x . That is, it says to round down to an integer. For example, $\lfloor 3.14 \rfloor = 3$, $\lfloor -1.3 \rfloor = -2$, and $\lfloor 5 \rfloor = 5$. (Some books denote the floor function by $[x]$ but this is bad notation since it does not suggest a corresponding notation for the ceiling function, and since square brackets are also used for other purposes.)

The ceiling function is defined by letting $\lceil x \rceil$ be the smallest integer greater than or equal to x . For example, $\lceil 3.14 \rceil = 4$, $\lceil -1.3 \rceil = -1$, and $\lceil 5 \rceil = 5$.

A.2.7 Factorial Function and Gamma Function

The factorial function takes a positive integer n and returns the product of the integers from 1 to n , denoted $n!$ and read " n factorial":

$$n! = n(n-1)(n-2) \cdots 1$$

Also, we define $0! = 1$. This convention makes sense since if we think of $n!$ as the number of ways in which n people can line up, there is 1 way for $n = 0$ (this just means there is no one there, so the line is empty). It is also very helpful since, for example, we can then say that $n!/(n-1)! = n$ for all positive integers n without running into trouble when $n = 1$.

The factorial function grows extremely quickly as n grows. A famous, useful approximation for factorials is Stirling's formula,

$$n! \approx \sqrt{2\pi n} \left(\frac{n}{e}\right)^n$$

The ratio of the two sides converges to 1 as $n \rightarrow \infty$. For example, direct calculation gives $52! \approx 8.066 \times 10^{67}$, while Stirling's formula says $52! \approx 8.053 \times 10^{67}$.

The gamma function Γ generalizes the factorial function to positive real numbers; it is defined by

$$\Gamma(a) = \int_0^\infty x^a e^{-x} \frac{dx}{x}, \quad a > 0$$

and has the property that

$$\Gamma(n) = (n-1)!$$

for all positive integers n . Also, $\Gamma(1/2) = \sqrt{\pi}$.

An important property of the gamma function, which generalizes the fact that $n! = n \cdot (n-1)!$, is that

$$\Gamma(a+1) = a\Gamma(a)$$

for all $a > 0$. See Chapter 8 for more about the gamma function.

B.4.3 A. 3 Matrices

Neo: What is the Matrix?

Trinity: The answer is out there, Neo, and it's looking for you, and it will find you if you want it to.

- The Matrix (film from 1999)

A matrix is a rectangular array of numbers, such as $\begin{pmatrix} 3 & 1/e \\ 2\pi & 1 \end{pmatrix}$ or $\begin{pmatrix} 1 & 1 & 0 \\ 1 & 2 & 3 \end{pmatrix}$. We say that the dimensions of a matrix are m by n (also written $m \times n$) if it has m rows and n columns (so the former example is 2 by 2, while the latter is 2 by 3). The matrix is called square if $m = n$. If $m = 1$, we have a row vector; if $n = 1$, we have a column vector. We write $A = (a_{ij})$ to indicate that the row i , column j entry of A is a_{ij} .

A.3.1 Matrix Addition and Multiplication

To add two matrices A and B with the same dimensions, just add the corresponding entries, e.g.,

$$\begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \end{pmatrix} = \begin{pmatrix} 2 & 1 & 0 \\ 2 & 2 & 1 \end{pmatrix}.$$

When we multiply an $m \times n$ matrix A by an $n \times r$ matrix B , we obtain an $m \times r$ matrix AB . The product AB is undefined if the number of columns of A does not equal the number of rows of B .

The row i , column j entry of AB is $\sum_{k=1}^n a_{ik}b_{kj}$, where a_{ij} and b_{ij} are the row i , column j entries of A and B , respectively. For example, here is how to multiply a 2×3 matrix by a 3×1 vector:

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \begin{pmatrix} 7 \\ 8 \\ 9 \end{pmatrix} = \begin{pmatrix} 1 \cdot 7 + 2 \cdot 8 + 3 \cdot 9 \\ 4 \cdot 7 + 5 \cdot 8 + 6 \cdot 9 \end{pmatrix} = \begin{pmatrix} 50 \\ 122 \end{pmatrix}$$

Note that AB may not equal BA , even if both are defined. To multiply a matrix A by a scalar, just multiply each entry by that scalar.

The transpose of a matrix A is the matrix whose row i , column j entry is the row j , column i entry of A . It is denoted by A' and read as " A transpose". The rows of A are the columns of A' , and the columns of A are the rows of A' . A square matrix A is symmetric if $A' = A$. A useful property of transposes is that if A and B are matrices such that the product AB is defined, then $(AB)' = B'A'$.

The determinant of a 2×2 matrix is defined by

$$\begin{vmatrix} a & b \\ c & d \end{vmatrix} = ad - bc$$

Determinants can also be defined for $n \times n$ matrices, in a recursive manner not reviewed here.

A.3.2 Eigenvalues and Eigenvectors

An eigenvalue of an $n \times n$ matrix A is a number λ such that

$$A\mathbf{v} = \lambda\mathbf{v}$$

for some $n \times 1$ column vector $\mathbf{v}$, where the elements of $\mathbf{v}$ are not all zero. The vector $\mathbf{v}$ is called an eigenvector of A , or a right eigenvector. (Similarly, a left eigenvector of A is a row vector $\mathbf{w}$ satisfying $\mathbf{w}A = \lambda\mathbf{w}$ for some λ .) This definition says that when A and $\mathbf{v}$ are multiplied, $\mathbf{v}$ just gets stretched by the constant λ .

Some matrices have no real eigenvalues, but the Perron-Frobenius theorem tells us that in a special case that is of particular interest to us in Chapter 11, eigenvalues exist and have nice properties. Let A be a square matrix whose entries are nonnegative and whose rows sum to 1

. Further assume that for all i and j , there exists $k \geq 1$ such that the row i , column j entry of A^k is positive. Then the Perron-Frobenius theorem says that 1 is an eigenvalue of A , which is in fact the largest eigenvalue of A , and there is a corresponding eigenvector whose entries are all positive.

B.4.4 A. 4 Difference Equations

A difference equation describes a sequence of numbers recursively in terms of earlier terms in the sequence. For example, $a_0, a_1, \dots$ is a Fibonacci sequence if

$$a_i = a_{i-1} + a_{i-2}$$

for all $i \geq 2$. There are infinitely many Fibonacci sequences, but such a sequence is uniquely determined after a_0 and a_1 are specified (these are called the initial conditions or boundary conditions). For example, $a_0 = 0, a_1 = 1$ yields the Fibonacci sequence 0, 1, 1, 2, 3, 5, 8, 13, ...

Difference equations often arise in probability, especially when applying LOTP. In this section, we will show how to solve difference equations of the form

$$p_i = p \cdot p_{i+1} + q \cdot p_{i-1}$$

where $p \neq 0$ and $q = 1 - p$. (This equation comes up in the gambler's ruin problem.) The first step is to guess a solution of the form $p_i = x^i$. Plugging this into the above, we have

$$x^i = p \cdot x^{i+1} + q \cdot x^{i-1}$$

which reduces to $px^2 - x + q = 0$. This is called the characteristic equation, and the solution to the difference equation depends on whether the characteristic equation has one or two distinct roots. If there are two distinct roots r_1 and r_2 , then the solution is of the form

$$p_i = ar_1^i + br_2^i$$

for some constants a and b . If there is only one distinct root r , then the solution is of the form

$$p_i = ar^i + bir^i$$

In our case, the characteristic equation has roots 1 and q/p , since

$$\frac{1 \pm \sqrt{1 - 4p(1 - p)}}{2p} = \frac{1 \pm \sqrt{(2p - 1)^2}}{2p} = \frac{1 \pm |2p - 1|}{2p}$$

is $(1 + 2p - 1)/(2p) = 1$ or $(2 - 2p)/(2p) = q/p$. The roots are distinct if $p \neq q$ and are both equal to 1 if $p = q$. So we have

$$p_i = \begin{cases} a + b \left(\frac{q}{p}\right)^i, & p \neq q \\ a + bi, & p = q \end{cases}$$

This is called the general solution of the difference equation, since we have not yet specified the constants a and b . To get a specific solution, we need to know two points in the sequence in order to solve for a and b .

B.4.5 A. 5 Differential Equations

Differential equations are the continuous version of difference equations. A differential equation uses derivatives to describe a function or collection of functions. For example, the differential equation

$$\frac{dy}{dx} = 3y$$

describes a collection of functions that have the following property: the instantaneous rate of change of the function at any point (x, y) is equal to $3y$. This is an example of a separable differential equation because we can separate the x 's and y 's, putting them on opposite sides of the equation:

$$\frac{dy}{y} = 3dx$$

Now we can integrate both sides, giving $\log y = 3x + c$, or equivalently,

$$y = Ce^{3x}$$

where C is any constant. This is called the general solution of the differential equation, and it tells us that all functions satisfying the differential equation are of the form $y = Ce^{3x}$ for some C . To get a specific solution, we need to specify one point on the graph, which allows us to solve for C .

B.4.6 A. 6 Partial Derivatives

If you can do ordinary derivatives, you can do partial derivatives: just hold all the other input variables constant except for the one you're differentiating with respect to. For example, let $f(x, y) = y \sin(x^2 + y^3)$. Then the partial derivative with respect to x is

$$\frac{\partial f(x, y)}{\partial x} = 2xy \cos(x^2 + y^3),$$

and the partial derivative with respect to y is

$$\frac{\partial f(x, y)}{\partial y} = \sin(x^2 + y^3) + 3y^3 \cos(x^2 + y^3).$$

These are first order partial derivatives. If we take a partial derivative of a first order partial derivative, we get a second order partial derivative, such as

$$\frac{\partial^2 f(x, y)}{\partial y \partial x} = \frac{\partial}{\partial y} \left(\frac{\partial f(x, y)}{\partial x} \right) = 2x (\cos(x^2 + y^3) - 3y^3 \sin(x^2 + y^3)).$$

We get the same result if we take the partial derivatives in the opposite order (with respect to y and then with respect to x):

$$\frac{\partial^2 f(x, y)}{\partial x \partial y} = \frac{\partial}{\partial x} \left(\frac{\partial f(x, y)}{\partial y} \right) = 2x (\cos(x^2 + y^3) - 3y^3 \sin(x^2 + y^3)).$$

Under mild technical assumptions, it is also true for general f that

$$\frac{\partial^2 f(x, y)}{\partial x \partial y} = \frac{\partial^2 f(x, y)}{\partial y \partial x}.$$

The Jacobian of a transformation which maps $(x_1, \dots, x_n)$ to $(y_1, \dots, y_n)$ is the $n \times n$ matrix of all possible first order partial derivatives, given by

$$\frac{\partial \mathbf{y}}{\partial \mathbf{x}} = \begin{pmatrix} \frac{\partial y_1}{\partial x_1} & \frac{\partial y_1}{\partial x_2} & \cdots & \frac{\partial y_1}{\partial x_n} \\ \vdots & \vdots & & \vdots \\ \frac{\partial y_n}{\partial x_1} & \frac{\partial y_n}{\partial x_2} & \cdots & \frac{\partial y_n}{\partial x_n} \end{pmatrix}$$

B.4.7 A. 7 Multiple Integrals

If you can do single integrals, you can do multiple integrals: just do more than one integral, holding variables other than the current variable of integration constant.

For example,

$$\begin{aligned} \int_0^1 \int_0^y (x-y)^2 dx dy &= \int_0^1 \int_0^y (x^2 - 2xy + y^2) dx dy \\ &= \int_0^1 \left((x^3/3 - x^2y + xy^2) \Big|_0^y \right) dy \\ &= \int_0^1 (y^3/3 - y^3 + y^3) dy \\ &= \frac{1}{12} \end{aligned}$$

A.7.1 Change of Order of Integration

We can also integrate in the other order, $dydx$ rather than $dx dy$, as long as we are careful about the limits of integration. Since we're integrating over all (x, y) with x and y between 0 and 1 such that $x \leq y$, to integrate the other way we write

$$\begin{aligned} \int_0^1 \int_x^1 (x-y)^2 dy dx &= \int_0^1 \int_x^1 (x^2 - 2xy + y^2) dy dx \\ &= \int_0^1 \left((x^2y - xy^2 + y^3/3) \Big|_x^1 \right) dx \\ &= \int_0^1 (x^2 - x + 1/3 - x^3 + x^3 - x^3/3) dx \\ &= \frac{1}{12} \end{aligned}$$

A.7.2 Change of Variables

In making a change of variables with multiple integrals, a Jacobian is needed. To simplify notation, we will discuss the two-dimensional version; the n -dimensional version is analogous. Suppose we make an invertible transformation from (x, y) to (u, v) , where the inverse transformation is given by

$$x = g(u, v), y = h(u, v)$$

Let A be a region in the (x, y) -plane and B be the corresponding region in the (u, v) -plane. Then the change of variables formulas says that

$$\iint_A f(x, y) dx dy = \iint_B f(g(u, v), h(u, v)) \cdot \left\| \frac{\partial(x, y)}{\partial(u, v)} \right\| du dv$$

where $\left| \left| \frac{\partial(x,y)}{\partial(u,v)} \right| \right|$ is the absolute value of the determinant of the Jacobian matrix. A few technical assumptions are needed for this result, such as that the partial derivatives that appear in the Jacobian matrix exist and are continuous, and that the determinant of the Jacobian matrix is never 0 .

For example, let's find the area of a circle of radius 1 . To find the area of a region, we just need to integrate 1 over that region (so the part that may be tricky is the limits of integration; the integrand is just 1). So the area is

$$\iint_{x^2+y^2 \leq 1} 1 dx dy = \int_{-1}^1 \int_{-\sqrt{1-y^2}}^{\sqrt{1-y^2}} 1 dx dy = 2 \int_{-1}^1 \sqrt{1-y^2} dy$$

Note that the limits for the inner variable (x) of the double integral can depend on the outer variable (y), while the outer limits are constants. The last integral can be done with a trigonometric substitution, but instead let's simplify the problem by transforming to polar coordinates: let

$$x = r \cos \theta, y = r \sin \theta,$$

where r is the distance from (x, y) to the origin and $\theta \in [0, 2\pi)$ is the angle. The Jacobian matrix of this transformation is

$$\frac{\partial(x,y)}{\partial(r,\theta)} = \begin{pmatrix} \cos \theta & -r \sin \theta \\ \sin \theta & r \cos \theta \end{pmatrix}$$

so the absolute value of the determinant of the Jacobian is $r (\cos^2 \theta + \sin^2 \theta) = r$. That is, $dx dy$ becomes $r dr d\theta$. So the area of the circle is

$$\int_0^{2\pi} \int_0^1 r dr d\theta = \int_0^{2\pi} \frac{1}{2} d\theta = \pi$$

For a circle of radius r , it follows immediately that the area is πr^2 since we can convert our units of measurement to the unit for which the radius is 1 .

This may seem like a lot of work just to get such a familiar result, but it served as illustration and with similar methods we can get the volume of a ball in any number of dimensions! It turns out that the volume of a ball of radius 1 in $\mathbb{R}^n$ is

$$V_n = \frac{\pi^{n/2}}{\Gamma(n/2 + 1)},$$

where Γ is the gamma function (see Section A.2.7).

B.4.8 A. 8 Sums

There are several kinds of sums that come up frequently in probability.

A.8.1 Binomial Theorem

The binomial theorem states that

$$(x + y)^n = \sum_{k=0}^n \binom{n}{k} x^k y^{n-k}$$

for any nonnegative integer n . Here $\binom{n}{k}$ is a binomial coefficient, defined as the number of ways to choose k objects out of n , with order not mattering. An explicit formula for $\binom{n}{k}$ in terms of factorials for $0 \leq k \leq n$ is

$$\binom{n}{k} = \frac{n!}{(n-k)!k!}$$

A proof of the binomial theorem is given in Example 1.4.19.

Sometimes 0^0 is said to be undefined, but there are several strong arguments for defining it to be 1, and we will take $0^0 = 1$ in this book. For one simple reason for doing so, consider the case $x = 0, y = 1$ in the binomial theorem. Then the left-hand side is $1^n = 1$, and we need to have $0^0 = 1$ in order to make the right-hand side also equal to 1.

A.8.2 Geometric Series

A series of the form $\sum_{n=0}^{\infty} x^n$ is called a geometric series. For $|x| < 1$, the series converges and we have

$$\sum_{n=0}^{\infty} x^n = \frac{1}{1-x}$$

The series diverges if $|x| \geq 1$. A series of the same form except with finitely many terms is called a finite geometric series. For $x \neq 1$, the sum is

$$\sum_{k=0}^n x^k = \frac{1-x^{n+1}}{1-x}$$

Note that the $n+1$ in the exponent is the number of terms in the sum. If the sum starts at $k = m$ rather than $k = 0$ (where m is an integer with $0 \leq m \leq n$), then we can easily reduce to the case where the sum starts at 0. Making the substitution $j = k - m$, we have

$$\sum_{k=m}^n x^k = \sum_{j=0}^{n-m} x^{j+m} = x^m \sum_{j=0}^{n-m} x^j = x^m \cdot \frac{1-x^{n-m+1}}{1-x}.$$

A.8.3 Taylor Series for e^x

The Taylor series for e^x is

$$\sum_{n=0}^{\infty} \frac{x^n}{n!} = e^x, \text{ for all } x$$

A.8.4 Harmonic Series and Other Sums with a Fixed Exponent

It is also useful to know that $\sum_{n=1}^{\infty} 1/n^c$ converges for $c > 1$ and diverges for $c \leq 1$. For $c = 1$, this is called the harmonic series. The sum of the first n terms of the harmonic series can be approximated using

$$\sum_{k=1}^n \frac{1}{k} \approx \log(n) + \gamma$$

for n large, where $\gamma \approx 0.577$ is the Euler-Mascheroni constant.

The sum of the first n positive integers is

$$\sum_{k=1}^n k = n(n+1)/2 = \binom{n+1}{2}.$$

It follows from this identity, for example, that the sum of the first n odd positive integers is

$$\sum_{k=1}^n (2k-1) = 2 \sum_{k=1}^n k - \sum_{k=1}^n 1 = n(n+1) - n = n^2.$$

For squares of integers, we have

$$\sum_{k=1}^n k^2 = n(n+1)(2n+1)/6.$$

For cubes of integers, amazingly, the sum is the square of the sum of the first n positive integers! That is,

$$\sum_{k=1}^n k^3 = (n(n+1)/2)^2$$

Exercise 22 from Chapter 1 asks for a story proof of this result.

B.4.9 A. 9 Pattern Recognition

Much of math and statistics is really about pattern recognition: seeing the essential structure of a problem, recognizing when one problem is essentially the same as another problem (just in a different guise), noticing symmetry, and so on. We will see many examples of this kind of thinking in this book. For example, suppose we have the series $\sum_{k=0}^{\infty} e^{tk} e^{-\lambda} \lambda^k / k!$, with λ a positive constant. The $e^{-\lambda}$ can be taken out from the sum, and then the structure of the series exactly matches up with the structure of the Taylor series for e^x . Therefore

$$\sum_{k=0}^{\infty} \frac{e^{tk} e^{-\lambda} \lambda^k}{k!} = e^{-\lambda} \sum_{k=0}^{\infty} \frac{(\lambda e^t)^k}{k!} = e^{-\lambda} e^{\lambda e^t} = e^{\lambda(e^t-1)},$$

valid for all real t .

Similarly, suppose we want the Taylor series for $1/(1-x^3)$ about $x=0$. It would be tedious to start taking derivatives of this function. Instead, note that this function is reminiscent of the result of summing a geometric series. Therefore

$$\frac{1}{1-x^3} = \sum_{n=0}^{\infty} x^{3n}$$

for $|x^3| < 1$ (which is equivalent to $|x| < 1$). What matters is the structure, not what names we use for variables!

B.4.10 A. 10 Common Sense and Checking Answers

It is very easy to make mistakes in probability, so checking answers is especially important. Some useful strategies for checking answers are:

- seeing whether the answer makes sense intuitively (though as we have often seen in this book, probability has many results that seem counterintuitive at first);
- making sure your answer isn't a category error;
- making sure you're avoiding biohazards;
- trying out simple cases;

- trying out extreme cases;
- looking for alternative methods to solve the problem (including methods that may only give a bound or approximation, such as applying one of the inequalities from Chapter 10 or using R to run a simulation).

B.5 R

B.5.1 B. 1 Vectors

Command	What it does
<code>c(1,1,0,2.7,3.1)</code>	creates the vector (1, 1, 0, 2.7, 3.1)
<code>1:100</code>	creates the vector (1, 2, ..., 100)
<code>(1 : 100)^3</code>	creates the vector (1 ³ , 2 ³ , ..., 100 ³)
<code>rep(0,50)</code>	creates the vector (0, 0, ..., 0) of length 50
<code>seq (0, 99, 3)</code>	creates the vector (0, 3, 6, 9, ..., 99)
<code>v [5]</code>	5 th entry of vector v (index starts at 1)
<code>v [-5]</code>	all but the 5th entry of v
<code>v[c(3,1,4)]</code>	3rd, 1st, 4th entries of vector v
<code>v[v > 2]</code>	entries of v that exceed 2
<code>which(v>2)</code>	indices of v such that entry exceeds 2
<code>which(v == 7)</code>	indices of v such that entry equals 7
<code>min (v)</code>	minimum of v
<code>max (v)</code>	maximum of v
<code>which.max(v)</code>	indices where $\max(v)$ is achieved
<code>sum (v)</code>	sum of the entries in v
<code>cumsum(v)</code>	cumulative sums of the entries in v
<code>prod(v)</code>	product of the entries in v
<code>rank(v)</code>	ranks of the entries in v
<code>length(v)</code>	length of vector v
<code>sort(v)</code>	sorts vector v (in increasing order)
<code>unique(v)</code>	lists each element of v once, without duplicates
<code>tabulate(v)</code>	tallies how many times each element of v occurs
<code>table(v)</code>	same as <code>tabulate(v)</code> , except in table format
<code>c(v,w)</code>	concatenates vectors v and w
<code>union (v , w)</code>	union of v and w as sets
<code>intersect(v,w)</code>	intersection of v and w as sets
<code>v+w</code>	adds v and w entrywise (recycling if needed)
<code>v*w</code>	multiplies v and w entrywise (recycling if needed)

B.5.2 B. 2 Matrices

Command	What it does
<code>matrix (c(1, 3, 5, 7), nrow = 2, ncol = 2)</code>	binds vectors u, v, w into a matrix, as rows binds vectors u row 2, column 3 entry of matrix A row 2 of matrix A (as a submatrix of A , keeping rows 1,3 and columns 2,4 row sum eigenvalues and eigenvectors of matrix A A^{-1} solves $A\mathbf{x} = \mathbf{b}$ for $\mathbf{x}$ (where $\mathbf{b}$ is a column vector) matrix m

B.5.3 B. 3 Math

Command	What it does
abs(x)	$ x $
exp(x)	e^x
log (x)	$\log(n)$
log(x, b)	$\log_b(n)$
sqrt(x)	$\sqrt{x}$
floor(x)	$\lfloor x \rfloor$
ceiling(x)	$\lceil x \rceil$
factorial(n)	$n!$
lfactorial(n)	$\log(n!)$ (helps prevent overflow)
gamma(a)	$\Gamma(a)$
lgamma(a)	$\log(\Gamma(a))$ (helps prevent overflow)
choose(n,k)	binomial coefficient $\binom{n}{k}$
pbirthday(k)	solves birthday problem for k people
if (x > 0)x^2 else x^3	x^2 if $x > 0$, x^3 otherwise (piecewise)
f <- function(x) exp(-x)	defines the function f by $f(x) = e^{-x}$
integrate(f, lower=0, upper=Inf)	finds $\int_0^\infty f(x)dx$ numerically
optimize(f , lower = 0, upper = 5, maximum=TRUE)	maximizes f numerically on $[0, 5]$
uniroot(f, lower = 0, upper = 5)	searches numerically for a zero of f in $[0, 5]$

B.5.4 B. 4 Sampling and Simulation

Command	What it does
sample(7)	random permutation of 1, 2, ..., 7
sample (52, 5)	picks 5 times from 1, 2, ..., 52 (don't replace)
sample(letters,5)	picks 5 random letters of the alphabet (don't replace)
sample(3,5, replace = TRUE , prob=p)	picks 5 times from 1, 2, 3 with probabilities p (replace)
replicate(10^4, experiment)	simulates 10^4 runs of experiment

B.5.5 B. 5 Plotting

Command	What it does
curve (f, from=a, to=b)	graphs the function f from a to b
plot(x,y)	creates scatter plot of the points (x_i, y_i)
plot(x,y,type="l")	creates line plot of the points (x_i, y_i)
points(x, y)	adds the points (x_i, y_i) to the plot
abline(a,b)	adds line segments through the (x_i, y_i) to the plot
hist(x, breaks=b, col="blue")	blue histogram of the values in x , with b bins suggested
par(new=TRUE)	tells R not to clear the palette when we make our next plot
par(mfrow=c(1,2))	tells R we want 2 side-by-side plots (a 1 by 2 array of plots)

B.5.6 B. 6 Programming

Command	What it does
<code>x <- pi</code>	sets x equal to π
<code>x > 3 && x < 5</code>	Is $x > 3$ and $x < 5$? (TRUE/FALSE)
<code>x > 3 x < 5</code>	Is $x > 3$ or $x < 5$? (TRUE/FALSE)
<code>if (n > 3)x <- -x + 1</code>	adds 1 to x if $n > 3$
<code>if (n == 0) x <- -x + 1 else x <- -x + 2</code>	adds 1 to x if $n = 0$, else adds 2
<code>v <- rep(0, 50); for (k in 1:50) v[k]<-rbinom(1, 50, 0.5)</code>	solves birthday problem up to 50 people

B.5.7 B. 7 Summary Statistics

Command	What it does
<code>mean (v)</code>	sample mean of vector v
<code>var(v)</code>	sample variance of vector v
<code>sd(v)</code>	sample standard deviation of vector v
<code>median(v)</code>	sample median of vector v
<code>summary(v)</code>	min, 1st quartile, median, mean, 3rd quartile, max of v
<code>quantile(v,p)</code>	p th sample quantile of vector v
<code>cov (v,w)</code>	sample covariance of vectors v and w
<code>cor(v, w)</code>	sample correlation of vectors v and w

B.5.8 B. 8 Distributions

Command	What it does
<code>help(distributions)</code>	shows documentation on distributions
<code>dbinom(k,n,p)</code>	PMF $P(X = k)$ for $X \sim \text{Bin}(n, p)$
<code>pbinom(x,n,p)</code>	CDF $P(X \leq x)$ for $X \sim \text{Bin}(n, p)$
<code>qbinom(a,n,p)</code>	quantile $\min\{x : P(X \leq x) \geq a\}$ for $X \sim \text{Bin}(n, p)$
<code>rbinom(r,n,p)</code>	vector of r i.i.d. $\text{Bin}(n, p)$ r.v.s
<code>dgeom(k,p)</code>	PMF $P(X = k)$ for $X \sim \text{Geom}(p)$
<code>dhyper(k,w,b,n)</code>	PMF $P(X = k)$ for $X \sim \text{HGeom}(w, b, n)$
<code>dnbinom(k,r,p)</code>	PMF $P(X = k)$ for $X \sim \text{NBin}(r, p)$
<code>dpois(k,r)</code>	PMF $P(X = k)$ for $X \sim \text{Pois}(r)$
<code>dbeta(x,a,b)</code>	PDF $f(x)$ for $X \sim \text{Beta}(a, b)$
<code>dcauchy(x)</code>	PDF $f(x)$ for $X \sim \text{Cauchy}$
<code>dchisq(x,n)</code>	PDF $f(x)$ for $X \sim \chi_n^2$
<code>dexp(x,r)</code>	PDF $f(x)$ for $X \sim \text{Expo}(r)$
<code>dgamma(x,a,r)</code>	PDF $f(x)$ for $X \sim \text{Gamma}(a, r)$
<code>dlnorm(x,m,s)</code>	PDF $f(x)$ for $X \sim \mathcal{LN}(m, s^2)$
<code>dnorm(x,m,s)</code>	PDF $f(x)$ for $X \sim \mathcal{N}(m, s^2)$
<code>dt(x, n)</code>	PDF $f(x)$ for $X \sim t_n$
<code>dunif(x,a,b)</code>	PDF $f(x)$ for $X \sim \text{Unif}(a, b)$
<code>dweibull(x,a,b)</code>	PDF $f(x)$ for $X \sim \text{Weibull}(b^{-a}, a)$

There are commands that are completely analogous to `pbinom`, `qbinom`, and `rbinom` for the other distributions appearing in the above table. For example, `pnorm`, `qnorm`, and `rnorm` can be used to get the CDF, the quantiles, and random generation for the Normal (note that the mean and standard deviation need to be provided as the parameters, rather than the mean and variance).

For the Multinomial, `dmultinom` can be used for calculating the joint PMF and `rmultinom` can be used for generating random vectors. For the Multivariate Normal, after installing and

loading the mvtnorm package, dmvnorm can be used for calculating the joint PDF and rmvnorm can be used for random generation.

B.6 Table of Distributions

Name	Param.	PMF or PDF	Mean	Variance
Bernoulli	p	$P(X = 1) = p, P(X = 0) = q$	p	pq
Binomial	n, p	$\binom{n}{k} p^k q^{n-k}, k \in \{0, 1, \dots, n\}$	np	npq
FS	p	$pq^{k-1}, k \in \{1, 2, \dots\}$	$1/p$	q/p^2
Geom	p	$pq^k, k \in \{0, 1, 2, \dots\}$	q/p	q/p^2
NBin	r, p	$\binom{r+k-1}{r-1} p^r q^k, k \in \{0, 1, 2, \dots\}$	rq/p	rq/p^2
HGeom	w, b, n	$\frac{\binom{w}{k} \binom{b}{n-k}}{\binom{w+b}{n}}, k \in \{0, 1, \dots, n\}$	$\mu = \frac{nw}{w+b}$	$\left(\frac{w+b-n}{w+b-1}\right) \mu \left(1 - \frac{\mu}{n}\right)$
NHGeom	w, b, r	$\frac{\binom{r+k-1}{r-1} \binom{w+b-r-k}{w-r}}{\binom{w+b}{w}}, k \in \{0, 1, \dots, b\}$	$\frac{rb}{w+1}$	$\frac{rb(w+b+1)(w-r+1)}{(w+1)^2(w+2)}$
Poisson	λ	$\frac{e^{-\lambda} \lambda^k}{k!}, k \in \{0, 1, 2, \dots\}$	λ	λ
Uniform	$a < b$	$\frac{1}{b-a}, x \in (a, b)$	$\frac{a+b}{2}$	$\frac{(b-a)^2}{12}$
Normal	μ, σ^2	$\frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/(2\sigma^2)}$	μ	σ^2
Log-Normal	μ, σ^2	$\frac{1}{x\sigma\sqrt{2\pi}} e^{-(\log x - \mu)^2/(2\sigma^2)}, x > 0$	$\theta = e^{\mu + \sigma^2/2}$	$\theta^2 (e^{\sigma^2} - 1)$
Expo	λ	$\lambda e^{-\lambda x}, x > 0$	$1/\lambda$	$1/\lambda^2$
Weibull	λ, γ	$\gamma \lambda e^{-\lambda x^\gamma} x^{\gamma-1}, x > 0$	$\mu = \frac{\Gamma(1+1/\gamma)}{\lambda^{1/\gamma}}$	$\frac{\Gamma(1+2/\gamma)}{\lambda^{2/\gamma}} - \mu^2$
Gamma	a, λ	$\Gamma(a)^{-1} (\lambda x)^a e^{-\lambda x} x^{-1}, x > 0$	a/λ	a/λ^2
Beta	a, b	$\frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} x^{a-1} (1-x)^{b-1}, 0 < x < 1$	$\mu = \frac{a}{a+b}$	$\frac{\mu(1-\mu)}{a+b+1}$
Chi-Square	n	$\frac{1}{2^{n/2} \Gamma(n/2)} x^{n/2-1} e^{-x/2}, x > 0$	n	$2n$
Student- t	n	$\frac{\Gamma((n+1)/2)}{\sqrt{n\pi} \Gamma(n/2)} (1 + x^2/n)^{-(n+1)/2}$	0 if $n > 1$	$\frac{n}{n-2}$ if $n > 2$

B.7 References

- [1] Donald J. Albers and Gerald L. Alexanderson. More Mathematical People: Contemporary Conversations. Academic Press, 1990.
- [2] Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. Handbook of Markov Chain Monte Carlo. CRC Press, 2011.
- [3] Jian Chen and Jeffrey S. Rosenthal. Decrypting classical cipher text using Markov chain Monte Carlo. Statistics and Computing, 22(2):397-413, 2012.
- [4] William G. Cochran. The effectiveness of adjustment by subclassification in removing bias in observational studies. Biometrics, 1968.
- [5] Persi Diaconis. Statistical problems in ESP research. Science, 201(4351):131136, 1978.
- [6] Persi Diaconis. The Markov chain Monte Carlo revolution. Bulletin of the American Mathematical Society, 46(2):179-205, 2009.
- [7] Persi Diaconis, Susan Holmes, and Richard Montgomery. Dynamical bias in the coin toss. SIAM Review, 49(2):211-235, 2007.
- [8] Bradley Efron and Ronald Thisted. Estimating the number of unseen species: How many words did Shakespeare know? Biometrika, 63(3):435, 1976.
- [9] Bradley Efron and Ronald Thisted. Did Shakespeare write a newly-discovered poem? Biometrika, 74:445-455, 1987.
- [10] Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. Bayesian Data Analysis. CRC Press, 2013.
- [11] Andrew Gelman and Deborah Nolan. You can load a die, but you can't bias a coin. The

American Statistician, 56(4):308-311, 2002.

[12] Andrew Gelman, Boris Shor, Joseph Bafumi, and David K. Park. Red State, Blue State, Rich State, Poor State: Why Americans Vote the Way They Do (Expanded Edition). Princeton University Press, 2009.

[13] Gerd Gigerenzer and Ulrich Hoffrage. How to improve Bayesian reasoning without instruction: Frequency formats. Psychological Review, 102(4):684, 1995.

[14] Prakash Gorroochurn. Classic Problems of Probability. John Wiley & Sons, 2012.

[15] Richard Hamming. You and your research. IEEE Potentials, pages 37-40, October 1993.

[16] David P. Harrington. The randomized clinical trial. Journal of the American Statistical Association, 95(449):312-315, 2000.

[17] David J.C. MacKay. Information Theory, Inference, and Learning Algorithms. Cambridge University Press, 2003.

[18] Richard McElreath. Statistical Rethinking: A Bayesian Course with Examples in *R* and Stan. Chapman & Hall/CRC Press, 2015.

[19] Pierre Rémond de Montmort. Essay d'Analyse sur les Jeux de Hazard. Quilau, Paris, 1708.

[20] John Allen Paulos. Innumeracy: Mathematical Illiteracy and Its Consequences. Macmillan, 1988.

[21] Horst Rinne. The Weibull Distribution: A Handbook. CRC Press, 2008.

[22] James G. Sanderson. Testing ecological patterns. American Scientist, 88:332-339, 2000.

[23] Tom W. Smith, Peter Marsden, Michael Hout, and Jibum Kim. General social surveys, 1972-2012. Sponsored by National Science Foundation. NORC, Chicago: National Opinion Research Center, 2013.

[24] Stephen M. Stigler. Isaac Newton as a probabilist. Statistical Science, 21(3):400-403, 2006.

[25] Tom Stoppard. Rosencrantz & Guildenstern Are Dead. Samuel French, Inc., 1967.

[26] R.J. Stroeker. On the sum of consecutive cubes being a perfect square. Compositio Mathematica, 97:295-307, 1995.

[27] Amos Tversky and Daniel Kahneman. Causal schemas in judgments under uncertainty. In Daniel Kahneman, Paul Slovic, and Amos Tversky, editors, Judgment under Uncertainty: Heuristics and Biases. Cambridge University Press, 1982.

[28] Herbert S. Wilf. generatingfunctionology. A K Peters/CRC Press, 3rd edition, 2005.

[29] Elizabeth Wrigley-Field. Length-biased sampling. Online article at <https://www.edge.org/response-detail/27022>, 2017.

[30] Lisa Zyga. Why too much evidence can be a bad thing. Online article at <https://phys.org/news/2016-01-evidence-bad.html>, 2016.

C Bibliography

References

- [1] А.М.Зубков. *Конспект лекций по теории вероятностей Механико-математический факультет МГУ, 4-й семестр.* 2014.
- [2] Спенсер Дж. Алон Н. *Вероятностный метод.* Бином, 2007.
- [3] Б. В. Гнеденко. *Курс теории вероятностей.* URSS, 2010.
- [4] В. В. Горяйнов. *Лекции по теории вероятностей МФТИ.* 2019.
- [5] Н.А. Гусев. *Заметки к курсу Мера и интеграл Лебега.* 2020.
- [6] А.Г. Семенов. *Уравнение Ланжевена и функциональный интеграл.* заметки к лекциям, 2010.

References From “2.8 Probability Theory - a Comprehensive Course” by Klenke

1. Adler RJ (1990) An introduction to continuity, extrema, and related topics for general Gaussian processes. Institute of Mathematical Statistics lecture notes-monograph series, vol 12. Institute of Mathematical Statistics, Hayward
2. Aizenman M, Kesten H, Newman CM (1987) Uniqueness of the infinite cluster and continuity of connectivity functions for short and long range percolation. Commun Math Phys 111(4):505-531
3. Aizenman M, Kesten H, Newman CM (1987) Uniqueness of the infinite cluster and related results in percolation. In: Percolation theory and ergodic theory of infinite particle systems, Minneapolis, Minn, 1984-1985. IMA volumes in mathematics and its applications, vol 8. Springer, New York, pp 13-20
4. Aldous DJ (1985) Exchangeability and related topics. In: École d'été de probabilités de SaintFlour, XIII-1983. Lecture notes in mathematics, vol 1117. Springer, Berlin, pp 1-198
5. Athreya KB, Ney PE (1972) Branching processes. Springer, Berlin
6. Azéma J, Yor M (1979) Le problème de Skorokhod: compléments à "Une solution simple au problème de Skorokhod". In: Séminaire de probabilités, XIII, Univ Strasbourg, Strasbourg, 1977/78. Lecture notes in mathematics, vol 721. Springer, Berlin, pp 625-633
7. Azéma J, Yor M (1979) Une solution simple au problème de Skorokhod. In: Séminaire de probabilités, XIII, Univ. Strasbourg, Strasbourg, 1977/78. Lecture notes in mathematics, vol 721. Springer, Berlin, pp 90-115
8. Baum LE, Katz M (1965) Convergence rates in the law of large numbers. Trans Am Math Soc 120:108-123
9. Baxter M, Rennie R (1997) Financial calculus. Cambridge University Press, Cambridge
10. Berry AC (1941) The accuracy of the Gaussian approximation to the sum of independent variates. Trans Am Math Soc 49:122-136

11. Beutelspacher A (2009) Kryptologie, 9th edn. Vieweg+Teubner, Wiesbaden
12. Billingsley P (1968) Convergence of probability measures. Wiley, New York
13. Billingsley P (1971) Weak convergence of measures: applications in probability. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1971. Conference board of the mathematical sciences regional conference series in applied mathematics, No 5
14. Billingsley P (1999) Convergence of probability measures, 2nd edn. Wiley series in probability and statistics: probability and statistics. Wiley, New York. A Wiley-Interscience publication
15. Binder K, Heermann DW (1997) Monte Carlo simulation in statistical physics: an introduction, 3rd edn. Springer series in solid-state sciences, vol 80. Springer, Berlin
16. Birkhoff GD (1931) Proof of the ergodic theorem. Proc Natl Acad Sci 17:656-660
17. Blackwell D, Kendall D (1964) The martin boundary of Pólya's urn scheme, and an application to stochastic population growth. J Appl Probab 1:284-296
18. Blumenthal RM (1957) An extended Markov property. Trans Am Math Soc 85:52-72
19. Bochner S (1932) Vorlesungen über Fouriersche Integrale. Chelsea, New York. Reprinted 1948
20. Breiman L (1968) Probability. Addison-Wesley, Reading
21. Brémaud P (1999) Markov chains. Texts in applied mathematics, vol 31. Springer, New York. Gibbs fields, Monte Carlo simulation, and queues
22. Brüggemann D (2002) Starke Gesetze der großen Zahlen bei blockweisen Unabhängigkeitsbedingungen. PhD thesis, Universität zu Köln
23. Burton RM, Keane M (1989) Density and uniqueness in percolation. Commun Math Phys 121(3):501-505
24. Choquet G, Deny J (1960) Sur l'équation de convolution $\mu = \mu * \sigma$. C R Acad Sci Paris 250:799-801
25. Chow YS, Teicher H (1997) Probability theory: independence, interchangeability, martingales, 3rd edn. Springer texts in statistics. Springer, New York
26. Chung KL, Fuchs WHJ (1951) On the distribution of values of sums of random variables. Memoirs of the American Mathematical Society, vol 6
27. Chung KL (1960) Markov chains with stationary transition probabilities. Die Grundlehren der mathematischen Wissenschaften, vol 104. Springer, Berlin
28. Clifford P, Sudbury A (1973) A model for spatial conflict. Biometrika 60:581-588
29. Cramér H (1938) Sur un nouveau théorème-limite de la théorie des probabilités. Actual Sci Ind 763:5-23. Colloque consacré à la théorie des probabilités
30. Delbaen F, Schachermayer W (1994) A general version of the fundamental theorem of asset pricing. Math Ann 300(3):463-520

31. Dembo A, Zeitouni O (1998) Large deviations techniques and applications, 2nd edn. Applications of mathematics, vol 38. Springer, New York
32. Deuschel J-D, Stroock DW (1989) Large deviations. Pure and applied mathematics, vol 137. Academic Press, Boston
33. Diaconis P, Freedman D (1980) Finite exchangeable sequences. *Ann Probab* 8(4):745-764
34. Dieudonné J (1960) Foundations of modern analysis. Pure and applied mathematics, vol X. Academic Press, New York
35. Donsker MD (1951) An invariance principle for certain probability limit theorems. *Memoirs of the American Mathematical Society*, vol 6
36. Doyle PG, Snell JL (1984) Random walks and electric networks. Carus mathematical monographs, vol 22. Mathematical Association of America, Washington
37. Dudley RM (2002) Real analysis and probability. Cambridge studies in advanced mathematics, vol 74. Cambridge University Press, Cambridge. Revised reprint of the 1989 original
38. Dunford N, Schwartz JT (1958) Linear operators. I. General theory. With the assistance of WG Bade and RG Bartle. Pure and applied mathematics, vol 7. Interscience Publishers, New York
39. Durrett R (2010) Probability: theory and examples, 4th edn. Cambridge series in statistical and probabilistic mathematics. Cambridge University Press, Cambridge
40. Dvoretzky A, Erdős P, Kakutani S (1961) Nonincrease everywhere of the Brownian motion process. In: *Proceedings of the 4th Berkeley symposium on mathematics, statistics and probability*, vol II. University of California Press, Berkeley, pp 103-116
41. Egoroff D (1911) Sur les suites des fonctions mesurables. *C R Acad Sci Paris* 152:135-157
42. Elliott RJ, Kopp PE (1999) Mathematics of financial markets. Springer finance. Springer, New York
43. Ellis RS (1985) Entropy, large deviations, and statistical mechanics. *Grundlehren der Mathematischen Wissenschaften*, vol 271. Springer, New York
44. Elstrodt J (2011) Maß- und Integrationstheorie, 7th edn. Springer, New York
45. Erdős P, Graham RL (1972) On a linear diophantine problem of Frobenius. *Acta Arith* 21:399-408
46. Esseen C-G (1942) On the Liapounoff limit of error in the theory of probability. *Ark Mat Astron Fys A* 28(9):1-19
47. Etemadi N (1981) An elementary proof of the strong law of large numbers. *Z Wahrscheinlichkeitstheor Verw Geb* 55(1):119-122
48. Etheridge A (2002) A course in financial calculus. Cambridge University Press, Cambridge
49. Ethier SN, Kurtz TG (1986) Markov processes: characterization and convergence. Wiley series in probability and mathematical statistics: probability and mathematical statistics. Wiley, New York

- 50. Evans SN, Zhou X (1999) Identifiability of exchangeable sequences with identically distributed partial sums. *Electron Commun Probab* 4:9-13 (electronic)
- 51. Feller W (1935) Über den zentralen Grenzwertsatz der Wahrscheinlichkeitstheorie I. *Math Z* 40:521-559
- 52. Feller W (1937) Über den zentralen Grenzwertsatz der Wahrscheinlichkeitstheorie II. *Math Z* 42:301-312
- 53. Feller W (1968) An introduction to probability theory and its applications, vol I, 3rd edn. Wiley, New York
- 54. Feller W (1971) An introduction to probability theory and its applications, vol II, 2nd edn. Wiley, New York
- 55. Fill JA (1998) An interruptible algorithm for perfect sampling via Markov chains. *Ann Appl Probab* 8(1):131-162
- 56. Fill JA, Machida M, Murdoch DJ, Rosenthal JS (2000) Extension of Fill's perfect rejection sampling algorithm to general chains. *Random Struct Algorithms* 17(3-4):290-316. Proceedings of the ninth international conference "Random structures and algorithms" (Poznan, 1999)
- 57. Föllmer H, Schied A (2004) Stochastic finance, 2nd edn. de Gruyter studies in mathematics, vol 27. Walter de Gruyter, Berlin
- 58. Freedman DA (1965) Bernard Friedman's urn. *Ann Math Stat* 36:956-970
- 59. Georgii H-O (2012) Stochastics: introduction to probability theory and statistics. de Gruyter Lehrbuch, 2nd edn. Walter de Gruyter, Berlin
- 60. Gibbs AL, Su FE (2002) On choosing and bounding probability metrics. *Int Stat Rev* 70(3):419-435
- 61. Glasser ML, Zucker IJ (1977) Extended Watson integrals for the cubic lattices. *Proc Natl Acad Sci USA* 74(5):1800-1801
- 62. Gnedenko BV, Kolmogorov AN (1968) Limit distributions for sums of independent random variables. Addison-Wesley, Reading
- 63. Grimmett G (1999) Percolation, 2nd edn. Grundlehren der Mathematischen Wissenschaften, vol 321. Springer, Berlin
- 64. Grimmett GR, Stirzaker DR (2001) Probability and random processes, 3rd edn. Oxford University Press, New York
- 65. Grosswald E (1976) The Student t -distribution of any degree of freedom is infinitely divisible. *Z Wahrscheinlichkeitstheor Verw Geb* 36(2):103-109
- 66. Häggström O (2002) Finite Markov chains and algorithmic applications. London Mathematical Society student texts, vol 52. Cambridge University Press, Cambridge
- 67. Hara T, Slade G (1990) Mean-field critical behaviour for percolation in high dimensions. *Commun Math Phys* 128(2):333-391
- 68. Harrison JM, Pliska SR (1981) Martingales and stochastic integrals in the theory of continuous trading. *Stoch Process Appl* 11(3):215-260

- 69. Hartman P, Wintner A (1941) On the law of the iterated logarithm. *Am J Math* 63:169-176
- 70. Hastings WK (1970) Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* 57:97-109
- 71. Hewitt E, Ross KA (1970) Abstract harmonic analysis, vol II: Structure and analysis for compact groups. Analysis on locally compact Abelian groups. *Die Grundlehren der mathematischen Wissenschaften*, vol 152. Springer, New York
- 72. Hewitt E, Savage LJ (1955) Symmetric measures on Cartesian products. *Trans Am Math Soc* 80:470-501
- 73. Heyde CC (1963) On a property of the lognormal distribution. *J R Stat Soc B* 29:392-393
- 74. den Hollander F (2000) Large deviations. *Fields Institute monographs*, vol 14. American Mathematical Society, Providence
- 75. Holley RA, Liggett TM (1975) Ergodic theorems for weakly interacting infinite systems and the voter model. *Ann Probab* 3(4):643-663
- 76. Hughes BD (1995) Random walks and random environments, vol 1. Oxford science publications. The Clarendon Press/Oxford University Press, New York. Random walks
- 77. Hughes BD (1996) Random walks and random environments, vol 2. Oxford science publications. The Clarendon Press/Oxford University Press, New York. Random environments
- 78. Ikeda N, Watanabe S (1989) Stochastic differential equations and diffusion processes, 2nd edn. *North-Holland mathematical library*, vol 24. North-Holland, Amsterdam
- 79. Jost J (2013) Partial differential equations, 3rd edn. *Graduate texts in mathematics*, vol 214. Springer, New York
- 80. Joyce GS (2003) Singular behaviour of the lattice Green function for the d -dimensional hypercubic lattice. *J Phys A* 36(4):911-921
- 81. Kakutani S (1973) Examples of ergodic measure preserving transformations which are weakly mixing but not strongly mixing. In: *Recent advances in topological dynamics (Proceedings of the conference at Yale University in honor of Gustav Arnold Hedlund)*, New Haven, CT, 1972. *Lecture notes in mathematics*, vol 318. Springer, Berlin, pp 143-149
- 82. Kallenberg O (1986) Random measures, 4th edn. Akademie, Berlin
- 83. Kallenberg O (2002) Foundations of modern probability, 2nd edn. Probability and its applications. Springer, New York
- 84. Kantorovič LV, Rubinšteĭn GŠ (1958) On a space of completely additive functions. *Vestn Leningr Univ* 13(7):52-59
- 85. Karatzas I, Shreve SE (1991) Brownian motion and stochastic calculus, 2nd edn. *Graduate texts in mathematics*, vol 113. Springer, New York
- 86. Karatzas I, Shreve SE (1998) Methods of mathematical finance. Applications of mathematics, vol 39. Springer, New York

- 87. Kato T (1976) Perturbation theory for linear operators, 2nd edn. Grundlehren der mathematischen Wissenschaften, vol 132. Springer, Berlin
- 88. Keller G (1998) Equilibrium states in ergodic theory. London Mathematical Society student texts, vol 42. Cambridge University Press, Cambridge
- 89. Keller G (2003) Wahrscheinlichkeitstheorie. Lecture notes (German). Universität Erlangen
- 90. Kelley JL (1975) General topology. Graduate texts in mathematics, vol 27. Springer, New York. Reprint of the 1955 edn. Van Nostrand, Toronto
- 91. Kemeny JG, Snell JL (1976) Finite Markov chains. Undergraduate texts in mathematics. Springer, New York. Reprinting of the 1960 original
- 92. Kenyon RW, Propp JG, Wilson DB (2000) Trees and matchings. Electron J Combin 7: Research Paper 25, 34 pp (electronic)
- 93. Kesten H, Kozlov MV, Spitzer F (1975) A limit law for random walk in a random environment. Compos Math 30:145-168
- 94. Kesten H (1975) Sums of stationary sequences cannot grow slower than linearly. Proc Am Math Soc 49:205-211
- 95. Kesten H (1980) The critical probability of bond percolation on the square lattice equals $\frac{1}{2}$. Commun Math Phys 74(1):41-59
- 96. Kesten H, Stigum BP (1966) A limit theorem for multidimensional Galton-Watson processes. Ann Math Stat 37:1211-1223
- 97. Khintchine A (1923) Über dyadische Brüche. Math Z 18:109-116
- 98. Kingman JFC (1978) Uses of exchangeability. Ann Probab 6(2):183-197
- 99. Kingman JFC (1993) Poisson processes. Oxford studies in probability, vol 3. The Clarendon Press/Oxford University Press, New York. Oxford Science publications
- 100. Klenke A, Mattner L (2010) Stochastic ordering of classical discrete distributions. Adv Appl Probab 42(2):392-410
- 101. Kolmogorov AN (1933) Sulla determinazione empirica di una legge di distribuzione. G Ist Ital Attuari 4:83-91
- 102. Korn R, Korn E (2001) Option pricing and portfolio optimization. Graduate studies in mathematics, vol 31. American Mathematical Society, Providence. Modern methods of financial mathematics, translated from the 1999 German original by the authors
- 103. Krengel U (1985) Ergodic theorems. de Gruyter studies in mathematics, vol 6. Walter de Gruyter, Berlin
- 104. Kullback S, Leibler RA (1951) On information and sufficiency. Ann Math Stat 22:79-86
- 105. Lauritzen SL (1988) Extremal families and systems of sufficient statistics. Lecture notes in statistics, vol 49. Springer, New York
- 106. Lévy P (1937) Théorie de l'Addition des Variables Aléatoires. Gauthier-Villars, Paris

- 107. Lévy *P* (1948) Processus stochastiques et mouvement Brownien. Suivi d'une note de M. Loève. Gauthier-Villars, Paris
- 108. Lindeberg JW (1922) Eine neue Herleitung des Exponentialgesetzes in der Wahrscheinlichkeitsrechnung. *Math Z* 15:211-225
- 109. Lindvall T (1972) Convergence of critical Galton-Watson branching processes. *J Appl Probab* 9:445-450
- 110. Lyons R, Pemantle R, Peres Y (1995) Conceptual proofs of $L \log L$ criteria for mean behavior of branching processes. *Ann Probab* 23(3):1125-1138
- 111. Lyons R, Peres Y (2012) Probability on trees. Preliminary version on the Internet at <http://mypage.iu.edu/~rdlyons/>
- 112. Madras N (2002) Lectures on Monte Carlo methods. Fields Institute monographs, vol 16. American Mathematical Society, Providence
- 113. Menchoff DE (1923) Sur les séries des fonctions orthogonales (première partie). *Fundam Math* 4:92-105
- 114. Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E (1953) Equation of state calculations by fast computing machines. *J Chem Phys* 21:1087-1092
- 115. Meyer P-A (1966) Probability and potentials. Blaisdell Ginn, Waltham
- 116. Meyn SP, Tweedie RL (1993) Markov chains and stochastic stability. Communications and control engineering series. Springer, London
- 117. Móricz F, Tandori K (1996) An improved Menshov-Rademacher theorem. *Proc Am Math Soc* 124(3):877-885
- 118. Mörters P, Peres Y (2010) Brownian motion. Cambridge series in statistical and probabilistic mathematics. Cambridge University Press, Cambridge. With an appendix by Oded Schramm and Wendelin Werner
- 119. Motwani R, Raghavan P (1995) Randomized algorithms. Cambridge University Press, Cambridge
- 120. Müller A, Stoyan D (2002) Comparison methods for stochastic models and risks. Wiley series in probability and statistics. Wiley, Chichester
- 121. Musiela M, Rutkowski M (2005) Martingale methods in financial modelling, 2nd edn. Stochastic modelling and applied probability, vol 36. Springer, Berlin
- 122. von Neumann J (1932) Proof of the quasi-ergodic hypothesis. *Proc Natl Acad Sci* 18:70-82
- 123. Norris JR (1998) Markov chains. Cambridge series in statistical and probabilistic mathematics. Cambridge University Press, Cambridge. Reprint of the 1997 edition
- 124. Nummelin E (1984) General irreducible Markov chains and nonnegative operators. Cambridge tracts in mathematics, vol 83. Cambridge University Press, Cambridge
- 125. Paley REAC, Wiener N (1987) Fourier transforms in the complex domain. American Mathematical Society colloquium publications, vol 19. American Mathematical Society, Providence. Reprint of the 1934 original

- 126. Paley REAC, Wiener N, Zygmund A (1933) Note on random functions. *Math Z* 37:647-668
- 127. Peierls RF (1936) On Ising's model of ferromagnetism. *Proc Camb Philos Soc* 32:477-481
- 128. Petrov VV (1975) Sums of independent random variables. *Ergebnisse der Mathematik und ihrer Grenzgebiete*, vol 82. Springer, New York
- 129. Pitman J (1995) Exchangeable and partially exchangeable random partitions. *Probab Theory Relat Fields* 102(2):145-158
- 130. Pitman J (2006) Combinatorial stochastic processes. *Lecture notes in mathematics*, vol 1875. Springer, Berlin, pp 7-24. Lectures from the 32nd summer school on probability theory held in Saint-Flour, July 7-24, 2002, with a foreword by Jean Picard
- 131. Pitman J, Yor M (1981) Bessel processes and infinitely divisible laws. In: *Stochastic integrals (Proceedings of the symposium at the University of Durham, Durham, 1980)*. *Lecture notes in mathematics*, vol 851. Springer, Berlin, pp 285-370
- 132. Pitman J, Yor M (1997) The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator. *Ann Probab* 25(2):855-900
- 133. Pitman J, Yor M (2001) On the distribution of ranked heights of excursions of a Brownian bridge. *Ann Probab* 29(1):361-384
- 134. Pólya G (1921) Über eine Aufgabe der Wahrscheinlichkeitsrechnung betreffend die Irrfahrt im Straßennetz. *Math Ann* 84:149-160
- 135. Pólya G (1931) Sur quelques points de la théorie de probabilités. *Ann Inst Henri Poincaré* 1:117-161
- 136. Prohorov YuV (1956) Convergence of random processes and limit theorems in probability theory. *Teor Veroâtn Ee Primen* 1:177-238. Russian with English summary
- 137. Propp J, Wilson D (1998) Coupling from the past: a user's guide. In: *Microsurveys in discrete probability*, Princeton, NJ, 1997. DIMACS series in discrete mathematics and theoretical computer science, vol 41. American Mathematical Society, Providence, pp 181-192
- 138. Propp JG, Wilson DB (1996) Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Struct Algorithms* 9(1-2):223-252
- 139. Propp JG, Wilson DB (1998) How to get a perfectly random sample from a generic Markov chain and generate a random spanning tree of a directed graph. *J Algorithms* 27(2):170-217. 7th Annual ACM-SIAM Symposium on Discrete Algorithms (Atlanta, GA, 1996)
- 140. Protter PE (2004) *Stochastic integration and differential equations*, 2nd edn. Applications of mathematics (New York), vol 21. Springer, Berlin. Stochastic modelling and applied probability
- 141. Rademacher H (1922) Einige Sätze über Reihen von allgemeinen Orthogonalfunktionen. *Math Ann* 87:112-138
- 142. Révész P (2005) *Random walk in random and non-random environments*, 2nd edn. World Scientific, Hackensack

- 143. Revuz D (1984) Markov chains, 2nd edn. North-Holland mathematical library, vol 11. NorthHolland, Amsterdam
- 144. Revuz D, Yor M (1999) Continuous martingales and Brownian motion, 3rd edn. Grundlehren der Mathematischen Wissenschaften, vol 293. Springer, Berlin
- 145. Rockafellar RT (1970) Convex analysis. Princeton mathematical series, vol 28. Princeton University Press, Princeton
- 146. Rogers LCG, Williams D (2000) Diffusions, Markov processes, and martingales, vol 1: Foundations. Cambridge mathematical library. Cambridge University Press, Cambridge. Reprint of the 2nd edition from 1994
- 147. Rogers LCG, Williams D (2000) Diffusions, Markov processes, and martingales, vol 2: Itô calculus. Cambridge mathematical library. Cambridge University Press, Cambridge. Reprint of the 2nd edition from 1994
- 148. Rudin W (1976) Principles of mathematical analysis, 3rd edn. International series in pure and applied mathematics. McGraw-Hill, New York
- 149. Sanov IN (1957) On the probability of large deviations of random magnitudes. Mat Sb (NS) 42(84):11-44 (in Russian)
- 150. Sanov IN (1961) On the probability of large deviations of random variables. Sel Transl Math Stat Probab 1:213-244
- 151. Schilling RL, Partzsch L (2012) Brownian motion. De Gruyter, Berlin. An introduction to stochastic processes, with a chapter on simulation by Björn Böttcher
- 152. Seneta E (2006) Non-negative matrices and Markov chains. Springer series in statistics. Springer, New York. Revised reprint of the second (1981) edition
- 153. Shiga T, Shimizu A (1980) Infinite-dimensional stochastic differential equations and their applications. J Math Kyoto Univ 20 (3):395-416
- 154. Shiryaev AN (1996) Probability, 2nd edn. Graduate texts in mathematics, vol 95. Springer, New York. Translation of the Russian edition from 1980
- 155. Sinai Ja (1959) On the concept of entropy for a dynamic system. Dokl Akad Nauk SSSR 124:768-771
- 156. Smirnov NV (1939) Sur les écarts de la Courbe de distribution empirique. Mat Sb2 : 3–16. Russian with French summary
- 157. Solomon F(1975) Random walks in a random environment. Ann Probab 3:1-31
- 158. Spitzer F (1976) Principles of random walks, 2nd edn. Graduate texts in mathematics, vol 34. Springer, New York
- 159. Steele JM (2001) Stochastic calculus and financial applications. Applications of mathematics (New York), vol 45. Springer, New York
- 160. Strassen V (1965) The existence of probability measures with given marginals. Ann Math Stat 36:423-439
- 161. Stroock DW, Varadhan SRS (1979) Multidimensional diffusion processes. Grundlehren der Mathematischen Wissenschaften, vol 233. Springer, Berlin

- 162. Sylvester JJ (1884) Mathematical questions with their solutions. *Educ Times* 41:171-178
- 163. Tandori K (1957) Über die orthogonalen Funktionen, I. *Acta Sci Math*, Szeged 18:57-130
- 164. Tandori K (1961) Über die Divergenz der Orthogonalreihen. *Publ Math (Debr)* 8:291-307
- 165. Tandori K (1986) Bemerkung über die paarweise unabhängigen zufälligen Größen. *Acta Math Hung* 48(3-4):357-359
- 166. Varadhan SRS (1966) Asymptotic probabilities and differential equations. *Commun Pure Appl Math* 19:261-286
- 167. Walters P (1982) An introduction to ergodic theory. Graduate texts in mathematics, vol 79. Springer, New York
- 168. Watson GN (1939) Three triple integrals. *Q J Math, Oxf Ser* 10:266-276
- 169. Williams D (1991) Probability with martingales. Cambridge mathematical textbooks. Cambridge University Press, Cambridge
- 170. Wilson DB, Propp JG (1996) How to get an exact sample from a generic Markov chain and sample a random spanning tree from a directed graph, both within the cover time. In: Proceedings of the seventh annual ACM-SIAM symposium on discrete algorithms, Atlanta, GA, 1996. ACM, New York, pp 448-457
- 171. Wright S (1931) Evolution in Mendelian populations. *Genetics* 16:97-159
- 172. Yamada T, Watanabe S (1971) On the uniqueness of solutions of stochastic differential equations. *J Math Kyoto Univ* 11:155-167
- 173. Yosida K (1995) Functional analysis. Classics in mathematics. Springer, Berlin. Reprint of the sixth edition from 1980
- 174. Zeitouni O (2004) Random walks in random environment. In: Lectures on probability theory and statistics. Lecture notes in mathematics, vol 1837. Springer, Berlin, pp 189-312