

Stochastic Processes

Theory for Applications

This definitive textbook provides a solid introduction to discrete and continuous stochastic processes, tackling a complex field in a way that instills a deep understanding of the relevant mathematical principles, and develops an intuitive grasp of the way these principles can be applied to modeling real-world systems.

- Basic underlying principles and axioms are made clear from the start, and new topics are developed as needed, encouraging and enabling students to develop an instinctive grasp of the fundamentals.
- Mathematical proofs are made easy for students to understand and remember, helping them quickly learn how to choose and apply the best possible models to real-world situations.
- Includes a review of elementary probability theory; detailed coverage of Poisson, Gaussian and Markov processes; the basic elements of queueing theory; and theory and applications of inference, hypothesis testing, detection and estimation, in addition to more advanced topics.

Written by one of the world's leading information theorists, based on his 20 years' experience of teaching stochastic processes to graduate students, this is an exceptional resource for anyone looking to develop their understanding of stochastic processes.

Robert G. Gallager is a Professor Emeritus at MIT, and one of the world's leading information theorists. He is a Member of the US National Academy of Engineering, and the US National Academy of Sciences, and his numerous awards and honors include the IEEE Medal of Honour (1990) and the Marconi Prize (2003). He was awarded the MIT Graduate Student Teaching Award in 1993, and this book is based on his 20 years' experience of teaching this subject to students.

“The book is a wonderful exposition of the key ideas, models, and results in stochastic processes most useful for diverse applications in communications, signal processing, analysis of computer and information systems, and beyond. The text provides excellent intuition, with numerous beautifully crafted examples, and exercises. Foundations are included in a natural way that enhances clarity and the reader’s ability to apply the results.”

Bruce Hajek, University of Illinois, Urbana-Champaign

“This book provides a beautiful treatment of the fundamentals of stochastic process. Gallager’s clear exposition conveys a deep and intuitive understanding of this important subject. Graduate students and researchers alike will benefit from this text, which will soon be a classic.”

Randall Berry, Northwestern University

“In *Stochastic Processes: Theory for Applications*, Robert Gallager has produced another in his series of outstanding texts. Using a style that is very intuitive and approachable, but without sacrificing the underlying rigor of the subject matter, he has focused his treatment exactly at the level that engineers and applied scientists need to understand in order to have a working knowledge of this field. The breadth and sequencing of the coverage are also excellent. This book will be a useful resource both for students entering the field and for practitioners seeking to deepen their understanding of stochastic methods.”

H. Vincent Poor, Princeton University

“Professor Gallager’s book is the first of a plethora of textbooks on stochastic processes for engineers that strike the perfect balance between broad coverage, rigor, and motivation for applications. With a wealth of illustrative examples and challenging exercises, this book is the ideal text for graduate students in any field that applies stochastic processes.”

Abbas El Gamal, Stanford University

Cambridge University Press
978-1-107-03975-9 - Stochastic Processes: Theory for Applications
Robert G. Gallager
Frontmatter
[More information](#)

Stochastic Processes

Theory for Applications

Robert G. Gallager
MIT



Cambridge University Press
978-1-107-03975-9 - Stochastic Processes: Theory for Applications
Robert G. Gallager
Frontmatter
[More information](#)

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning, and research at the highest international levels of excellence.

www.cambridge.org

Information on this title: www.cambridge.org/9781107039759

© Cambridge University Press 2013

This publication is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published 2013

Reprinted 2015

Printing in the United Kingdom by TJ International Ltd, Padstow, Cornwall

A catalogue record for this publication is available from the British Library

Library of Congress Cataloguing in Publication data

Gallager, Robert G.

Stochastic processes: theory for applications / Robert G. Gallager, MIT.

pages cm

ISBN 978-1-107-03975-9 (hardback)

I. Stochastic processes – Textbooks. I. Title.

QA274.G344 2013

519.2'3–dc23

2013005146

ISBN 978-1-107-03975-9 Hardback

Additional resources for this publication at www.cambridge.org/stochasticprocesses

Cambridge University Press has no responsibility for the persistence or accuracy of URLs for external or third-party internet websites referred to in this publication, and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

Cambridge University Press
978-1-107-03975-9 - Stochastic Processes: Theory for Applications
Robert G. Gallager
Frontmatter
[More information](#)

**To Marie, with thanks for her love and
encouragement while I finished this book**

Cambridge University Press
978-1-107-03975-9 - Stochastic Processes: Theory for Applications
Robert G. Gallager
Frontmatter
[More information](#)

Contents

	<i>Preface</i>	page xv
	<i>Suggestions for instructors and self study</i>	xix
	<i>Acknowledgements</i>	xxi
1	Introduction and review of probability	1
	1.1 Probability models	1
	1.1.1 The sample space of a probability model	3
	1.1.2 Assigning probabilities for finite sample spaces	4
	1.2 The axioms of probability theory	5
	1.2.1 Axioms for events	7
	1.2.2 Axioms of probability	8
	1.3 Probability review	9
	1.3.1 Conditional probabilities and statistical independence	9
	1.3.2 Repeated idealized experiments	11
	1.3.3 Random variables	12
	1.3.4 Multiple random variables and conditional probabilities	14
	1.4 Stochastic processes	16
	1.4.1 The Bernoulli process	17
	1.5 Expectations and more probability review	19
	1.5.1 Random variables as functions of other random variables	23
	1.5.2 Conditional expectations	25
	1.5.3 Typical values of random variables; mean and median	28
	1.5.4 Indicator random variables	29
	1.5.5 Moment generating functions and other transforms	29
	1.6 Basic inequalities	31
	1.6.1 The Markov inequality	32
	1.6.2 The Chebyshev inequality	32
	1.6.3 Chernoff bounds	33
	1.7 The laws of large numbers	36
	1.7.1 Weak law of large numbers with a finite variance	36
	1.7.2 Relative frequency	39
	1.7.3 The central limit theorem (CLT)	39
	1.7.4 Weak law with an infinite variance	44
	1.7.5 Convergence of random variables	45
	1.7.6 Convergence with probability 1	48

1.8	Relation of probability models to the real world	51
1.8.1	Relative frequencies in a probability model	52
1.8.2	Relative frequencies in the real world	52
1.8.3	Statistical independence of real-world experiments	55
1.8.4	Limitations of relative frequencies	56
1.8.5	Subjective probability	57
1.9	Summary	57
1.10	Exercises	58
2	Poisson processes	72
2.1	Introduction	72
2.1.1	Arrival processes	72
2.2	Definition and properties of a Poisson process	74
2.2.1	Memoryless property	75
2.2.2	Probability density of S_n and joint density of S_1, \dots, S_n	78
2.2.3	The probability mass function (PMF) for $N(t)$	79
2.2.4	Alternative definitions of Poisson processes	80
2.2.5	The Poisson process as a limit of shrinking Bernoulli processes	82
2.3	Combining and splitting Poisson processes	84
2.3.1	Subdividing a Poisson process	86
2.3.2	Examples using independent Poisson processes	87
2.4	Non-homogeneous Poisson processes	89
2.5	Conditional arrival densities and order statistics	92
2.6	Summary	96
2.7	Exercises	97
3	Gaussian random vectors and processes	105
3.1	Introduction	105
3.2	Gaussian random variables	105
3.3	Gaussian random vectors	107
3.3.1	Generating functions of Gaussian random vectors	108
3.3.2	IID normalized Gaussian random vectors	108
3.3.3	Jointly-Gaussian random vectors	109
3.3.4	Joint probability density for Gaussian n -rv s (special case)	112
3.4	Properties of covariance matrices	114
3.4.1	Symmetric matrices	114
3.4.2	Positive definite matrices and covariance matrices	115
3.4.3	Joint probability density for Gaussian n -rv s (general case)	117
3.4.4	Geometry and principal axes for Gaussian densities	118
3.5	Conditional PDFs for Gaussian random vectors	120
3.6	Gaussian processes	124
3.6.1	Stationarity and related concepts	126
3.6.2	Orthonormal expansions	128
3.6.3	Continuous-time Gaussian processes	130
3.6.4	Gaussian sinc processes	132

5	Renewal processes	214
5.1	Introduction	214
5.2	The strong law of large numbers and convergence with probability 1	217
5.2.1	Convergence with probability 1 (WP1)	217
5.2.2	Strong law of large numbers	219
5.3	Strong law for renewal processes	221
5.4	Renewal–reward processes; time averages	226
5.4.1	General renewal–reward processes	230
5.5	Stopping times for repeated experiments	233
5.5.1	Wald’s equality	236
5.5.2	Applying Wald’s equality to $E[N(t)]$	238
5.5.3	Generalized stopping trials, embedded renewals, and $G/G/1$ queues	239
5.5.4	Little’s theorem	242
5.5.5	$M/G/1$ queues	245
5.6	Expected number of renewals; ensemble averages	249
5.6.1	Laplace transform approach	250
5.6.2	The elementary renewal theorem	251
5.7	Renewal–reward processes; ensemble averages	254
5.7.1	Age and duration for arithmetic processes	255
5.7.2	Joint age and duration: non-arithmetic case	258
5.7.3	Age $Z(t)$ for finite t : non-arithmetic case	259
5.7.4	Age $Z(t)$ as $t \rightarrow \infty$: non-arithmetic case	262
5.7.5	Arbitrary renewal–reward functions: non-arithmetic case	264
5.8	Delayed renewal processes	266
5.8.1	Delayed renewal–reward processes	268
5.8.2	Transient behavior of delayed renewal processes	269
5.8.3	The equilibrium process	270
5.9	Summary	270
5.10	Exercises	271
6	Countable-state Markov chains	287
6.1	Introductory examples	287
6.2	First-passage times and recurrent states	289
6.3	Renewal theory applied to Markov chains	294
6.3.1	Renewal theory and positive recurrence	294
6.3.2	Steady state	296
6.3.3	Blackwell’s theorem applied to Markov chains	299
6.3.4	Age of an arithmetic renewal process	300
6.4	Birth–death Markov chains	302
6.5	Reversible Markov chains	303
6.6	The $M/M/1$ sampled-time Markov chain	307
6.7	Branching processes	309
6.8	Round-robin service and processor sharing	312

		317
	6.9 Summary	317
	6.10 Exercises	319
7	Markov processes with countable-state spaces	324
	7.1 Introduction	324
	7.1.1 The sampled-time approximation to a Markov process	328
	7.2 Steady-state behavior of irreducible Markov processes	329
	7.2.1 Renewals on successive entries to a given state	330
	7.2.2 The limiting fraction of time in each state	331
	7.2.3 Finding $\{p_j(i); j \geq 0\}$ in terms of $\{\pi_j; j \geq 0\}$	332
	7.2.4 Solving for the steady-state process probabilities directly	335
	7.2.5 The sampled-time approximation again	336
	7.2.6 Pathological cases	336
	7.3 The Kolmogorov differential equations	337
	7.4 Uniformization	341
	7.5 Birth–death processes	342
	7.5.1 The M/M/1 queue again	342
	7.5.2 Other birth–death systems	343
	7.6 Reversibility for Markov processes	344
	7.7 Jackson networks	350
	7.7.1 Closed Jackson networks	355
	7.8 Semi-Markov processes	357
	7.8.1 Example – the M/G/1 queue	360
	7.9 Summary	361
	7.10 Exercises	363
8	Detection, decisions, and hypothesis testing	375
	8.1 Decision criteria and the maximum a posteriori probability (MAP) criterion	376
	8.2 Binary MAP detection	379
	8.2.1 Sufficient statistics I	381
	8.2.2 Binary detection with a one-dimensional observation	381
	8.2.3 Binary MAP detection with vector observations	385
	8.2.4 Sufficient statistics II	391
	8.3 Binary detection with a minimum-cost criterion	395
	8.4 The error curve and the Neyman–Pearson rule	396
	8.4.1 The Neyman–Pearson detection rule	402
	8.4.2 The min–max detection rule	403
	8.5 Finitely many hypotheses	403
	8.5.1 Sufficient statistics with $M \geq 2$ hypotheses	407
	8.5.2 More general minimum-cost tests	409
	8.6 Summary	409
	8.7 Exercises	410

9	Random walks, large deviations, and martingales	417
9.1	Introduction	417
9.1.1	Simple random walks	418
9.1.2	Integer-valued random walks	419
9.1.3	Renewal processes as special cases of random walks	419
9.2	The queueing delay in a $G/G/1$ queue	420
9.3	Threshold crossing probabilities in random walks	423
9.3.1	The Chernoff bound	423
9.3.2	Tilted probabilities	425
9.3.3	Large deviations and compositions	428
9.3.4	Back to threshold crossings	431
9.4	Thresholds, stopping rules, and Wald's identity	433
9.4.1	Wald's identity for two thresholds	434
9.4.2	The relationship of Wald's identity to Wald's equality	435
9.4.3	Zero-mean random walks	435
9.4.4	Exponential bounds on the probability of threshold crossing	436
9.5	Binary hypotheses with IID observations	438
9.5.1	Binary hypotheses with a fixed number of observations	438
9.5.2	Sequential decisions for binary hypotheses	442
9.6	Martingales	444
9.6.1	Simple examples of martingales	444
9.6.2	Scaled branching processes	446
9.6.3	Partial isolation of past and future in martingales	446
9.7	Submartingales and supermartingales	447
9.8	Stopped processes and stopping trials	450
9.8.1	The Wald identity	452
9.9	The Kolmogorov inequalities	453
9.9.1	The SLLN	455
9.9.2	The martingale convergence theorem	456
9.10	A simple model for investments	458
9.10.1	Portfolios with constant fractional allocations	461
9.10.2	Portfolios with time-varying allocations	465
9.11	Markov modulated random walks	468
9.11.1	Generating functions for Markov random walks	469
9.11.2	Stopping trials for martingales relative to a process	471
9.11.3	Markov modulated random walks with thresholds	471
9.12	Summary	472
9.13	Exercises	473
10	Estimation	488
10.1	Introduction	488
10.1.1	The squared-cost function	489
10.1.2	Other cost functions	490
10.2	MMSE estimation for Gaussian random vectors	491

	10.2.1 Scalar iterative estimation	494
	10.2.2 Scalar Kalman filter	496
10.3	LLSE estimation	498
10.4	Filtered vector signal plus noise	500
	10.4.1 Estimate of a single random variable in IID vector noise	502
	10.4.2 Estimate of a single random variable in arbitrary vector noise	502
	10.4.3 Vector iterative estimation	503
	10.4.4 Vector Kalman filter	504
10.5	Estimation for circularly-symmetric Gaussian rv s	505
10.6	The vector space of random variables; orthogonality	507
10.7	MAP estimation and sufficient statistics	512
10.8	Parameter estimation	513
	10.8.1 Fisher information and the Cramer–Rao bound	516
	10.8.2 Vector observations	518
	10.8.3 Information	519
10.9	Summary	521
10.10	Exercises	523
	<i>References</i>	528
	<i>Index</i>	530

Cambridge University Press
978-1-107-03975-9 - Stochastic Processes: Theory for Applications
Robert G. Gallager
Frontmatter
[More information](#)

Preface

This text has evolved over some 20 years, starting as lecture notes for two first-year graduate subjects at MIT, namely, *Discrete Stochastic Processes* (6.262) and *Random Processes, Detection, and Estimation* (6.432). The two sets of notes are closely related and have been integrated into one text. Instructors and students can pick and choose the topics that meet their needs, and suggestions for doing this follow this preface.

These subjects originally had an application emphasis, the first on queueing and congestion in data networks and the second on modulation and detection of signals in the presence of noise. As the notes have evolved, it has become increasingly clear that the mathematical development (with minor enhancements) is applicable to a much broader set of applications in engineering, operations research, physics, biology, economics, finance, statistics, etc.

The field of stochastic processes is essentially a branch of probability theory, treating probabilistic models that evolve in time. It is best viewed as a branch of mathematics, starting with the axioms of probability and containing a rich and fascinating set of results following from those axioms. Although the results are applicable to many areas, they are best understood initially in terms of their mathematical structure and interrelationships.

Applying axiomatic probability results to a real-world area requires creating a probability model for the given area. Mathematically precise results can then be derived within the model and translated back to the real world. If the model fits the area sufficiently well, real problems can be solved by analysis within the model. However, since models are almost always simplified approximations of reality, precise results within the model become approximations in the real world.

Choosing an appropriate probability model is an essential part of this process. Sometimes an application area will have customary choices of models, or at least structured ways of selecting them. For example, there is a well-developed taxonomy of queueing models. A sound knowledge of the application area, combined with a sound knowledge of the behavior of these queueing models, often lets one choose a suitable model for a given issue within the application area. In other cases, one can start with a particularly simple model and use the behavior of that model to gain insight about the application, and use this to iteratively guide the selection of more general models.

An important aspect of choosing a probability model for a real-world area is that a prospective choice depends heavily on prior understanding, at both an intuitive and mathematical level, of results from the range of mathematical models that might be involved. This partly explains the title of the text – *Theory for Applications*. The aim is

to guide the reader in both the mathematical and intuitive understanding necessary in developing and using stochastic process models in studying application areas.

Application-oriented students often ask why it is important to understand axioms, theorems, and proofs in mathematical models when the precise results in the model become approximations in the real-world system being modeled. One answer is that a deeper understanding of the mathematics leads to the required intuition for understanding the differences between model and reality. Another answer is that theorems are transferable between applications, and thus enable insights from one application area to be transferred to another.

Given the need for precision in the theory, however, why is an axiomatic approach needed? Engineering and science students learn to use calculus, linear algebra, and undergraduate probability effectively without axioms or rigor. Why does this not work for more advanced probability and stochastic processes?

Probability theory has more than its share of apparent paradoxes, and these show up in very elementary arguments. Undergraduates are content with this, since they can postpone these questions to later study. For the more complex issues in graduate work, however, reasoning without a foundation becomes increasingly frustrating, and the axioms provide the foundation needed for sound reasoning without paradoxes.

I have tried to avoid the concise and formal proofs of pure mathematics, and instead use explanations that are longer but more intuitive while still being precise. This is partly to help students with limited exposure to pure mathematics, and partly because intuition is vital when going back and forth between a mathematical model and a real-world problem. In doing research, we grope toward results, and successful groping requires both a strong intuition and precise reasoning.

The text neither uses nor develops measure theory. Measure theory is undoubtedly important in understanding probability at a deep level, but most of the topics useful in many applications can be understood without measure theory. I believe that the level of precision here provides a good background for a later study of measure theory.

The text does require some background in probability at an undergraduate level. Chapter 1 presents this background material as a review, but it is too concentrated and deep for most students without prior background. Some exposure to linear algebra and analysis (especially concrete topics like vectors, matrices, and limits) is helpful, but the text develops the necessary results. The most important prerequisite is the mathematical maturity and patience to couple precise reasoning with intuition.

The organization of the text, after the review in Chapter 1 is as follows: Chapters 2, 3, and 4 treat three of the simplest and most important classes of stochastic processes, first Poisson processes, next Gaussian processes, and finally finite-state Markov chains. These are beautiful processes where almost everything is known, and they contribute insights, examples, and initial approaches for almost all other processes. Chapter 5 then treats renewal processes, which generalize Poisson processes and provide the foundation for the rest of the text.

Chapters 6 and 7 use renewal theory to generalize Markov chains to countable state spaces and continuous time. Chapters 8 and 10 then study decision making and estimation, which in a sense gets us out of the world of theory and back to using the theory.

Chapter 9 treats random walks, large deviations, and martingales and illustrates many of their applications.

Most results here are quite old and well established, so I have not made any effort to attribute results to investigators. My treatment of the material is indebted to the texts by Bertsekas and Tsitsiklis [2], Sheldon Ross [22] and William Feller [8] and [9].

Cambridge University Press
978-1-107-03975-9 - Stochastic Processes: Theory for Applications
Robert G. Gallager
Frontmatter
[More information](#)

Suggestions for instructors and self study

The subject of stochastic processes contains many beautiful and surprising results at a relatively simple level. These results should be savored and contemplated rather than rushed. The urge to go too quickly, to sacrifice understanding for shallow bottom lines, and to cover all the most important topics should be resisted. This text covers all the material in two full term graduate subjects at MIT, plus many other topics added for enrichment, so it cannot be ‘covered’ in one term.

My conviction is that if a student acquires a deep understanding of any, say, 20% of the material, then that student will be able to read and understand the rest with relative ease at a later time. Better still, a full appreciation of that 20% will make most students eager to learn more. In other words, instructors have a good deal of freedom, subject to a prerequisite structure, to choose topics of interest to them and their students to cover in a one term course.

One of the two MIT courses leading to this text covers Chapters 1, 2, 4, 5, 6, 7, and 9, skipping many of the more detailed parts of the latter five chapters. The other course covers Chapters 1, 3, 8, and 10, again omitting many topics. The first course is largely discrete and the second largely continuous, and a different mix is probably more appropriate for a student taking only one subject.

The topics in Chapter 1 are largely covered in good elementary probability subjects, but students are usually better at doing plug and chug exercises on these topics than having the depth of understanding required by the subsequent topics. Thus instructors should spend some time reviewing these topics.

It is difficult to be precise about the extent to which one topic is a prerequisites of another. The table below lists the prerequisites of each section. Most sections have only one prerequisite, but that recursively includes the prerequisites of the prerequisite. Instructors and students are encouraged to use their own judgement here.

Suggestions for instructors and self study

Sect.	Prereq.	Sect.	Prereq.	Sect.	Prereq.	Sect.	Prereq.	Sect.	Prereq.
1.2	1.1	3.4	3.3	5.6	5.5	7.5	7.2	9.7	9.6
1.3	1.2	3.5	3.4	5.7	5.6	7.6	7.5	9.8	9.7
1.4	1.3	3.6	3.5	5.8	5.7	7.7	7.6	9.9	9.8
1.5	1.4	3.7	3.5	6.1	5.5, 4.3	7.8	7.2	9.10	9.9
1.6	1.5	4.1	1.5	6.2	6.1	8.1	1.8	9.11	9.9
1.7	1.6	4.2	4.1	6.3	6.2	8.2	8.1, 3.3	10.1	1.8, 3.5
1.8	1.7	4.3	4.2	6.4	6.3	8.3	8.2	10.2	10.1
2.1	1.5	4.4	4.3	6.5	6.4	8.4	8.3	10.3	10.2
2.2	2.1	4.5	4.4	6.6	6.5	8.5	8.4	10.4	10.3
2.3	2.2	4.6	4.5	6.7	6.3	9.1	5.5	10.5	10.4, 3.7
2.4	2.3	5.1	1.7, 2.2	6.8	6.3	9.2	9.1	10.6	10.3
2.5	1.2	5.2	5.1	7.1	6.3	9.3	9.2	10.7	10.3
3.1	1.7	5.3	5.2	7.2	7.1	9.4	9.3	10.8	10.3
3.2	3.1	5.4	5.3	7.3	7.2	9.5	9.4, 8.2		
3.3	3.2	5.5	5.4	7.4	7.2	9.6	9.5		

Acknowledgements

This book has its roots in a book called *Discrete Stochastic Processes* that I wrote back in 1996, some lecture notes on continuous random processes from about the same time, and lecture notes that I have been writing at MIT from 2007 to 2012 for a subject also entitled *Discrete Stochastic Processes*.

I am deeply grateful to Professors John Wyatt, John Tsitsiklis, and Lizhong Zheng who have used these notes in teaching the *Discrete Stochastic Process* course in recent years. Their many general observations about the value and teachability of various topics and the suggestions of alternative approaches have been invaluable. Their ability (particularly in John Tsitsiklis' case) to catch minor flaws in proofs and suggest cleaner approaches has saved me from many errors. They are also responsible for a number of excellent new exercises.

Natasha Blitvic, Mina Karzand, and Fabian Kozynski who were teaching assistants in the course for the last five years, were also very helpful both in creating and improving the wording in a number of exercises, but also in explaining why students were having difficulty and how to improve the presentation.

I am also indebted to a number of people in the MIT community who have been helpful in the evolution of this book. Professors Dimitri Berksekas and Yuri Polyansky have helped in discussing topics for the book and in reading various sections. Shan-Yuan Ho has been enormously helpful in reading the entire manuscript and catching many things, ranging from typos to poorly presented concepts. She also wrote the solution manual for the old book and has been helpful through the whole evolution of this project.

A number of friends, students, and even people who found the text on the web have also been helpful in catching errors and inconsistencies and suggesting better approaches to various topics. Murat Azizoglu and Baris Nakiboglu have been particularly helpful in this regard. Others are Ivan Bersenco, Dimitris Bisias, Nathan Jones, Tarek Lahlou, Vahid Montazerhodju, Emre Teletar, Roy Yates, and Andrew Young. Finally, I am grateful to the many students who have taken the course in the last five years and who look puzzled (or sleepy) when something is explained badly, and who ask questions when it is almost clear.