CONTENTS

PREFACE

> The last thing one discovers in composing a work is what to put first.
>
> —Blaise Pascal

The intended audience for this book are students who like probability. With that prerequisite, I am confident that you will love stochastic processes.

Stochastic, or random, processes, is the dynamic side of probability. What differential equations is to calculus, stochastic processes is to probability. The material appeals to those who like applications and to those who like theory. It is both excellent preparation for future study, as well as a *terminal* course, in the sense that we don't have to tell students to wait until the next class or the next year before seeing the good stuff. This *is* the good stuff! Stochastic processes, as a branch of probability, speaks in the language of rolling dice, flipping coins and gambling games, but in the service of applications as varied as the spread of infectious diseases, the evolution of genetic sequences, models for climate change, and the growth of the worldwide web.

The book assumes the reader has taken a calculus-based probability course and is familiar with matrix algebra. Conditional probability and conditional expectation, which are essential tools, are offered in the introductory chapter, but may be skimmed over depending upon students' background. Some topics assume a greater knowledge of linear algebra than basic matrices (such as eigenvalues and eigenvectors) but these are optional, and relevant sections are starred. The book does not assume background in combinatorics, differential equations, or real analysis. Necessary mathematics are introduced as needed.

A focus of this book is the use of simulation. I have chosen to use the popular statistical freeware `R`, which is an accessible interactive computing environment. The use of simulation, important in its own right for applied work and mathematical research, is a powerful pedagogical tool for making theoretical concepts come alive with practical, hands-on demonstrations. It is not necessary to use `R` in order to use this book; code and script files are supplemental. However, the software is easy — and fun — to learn, and there is a tutorial and exercises in an appendix for bringing students up to speed.

The book contains more than enough material for a standard one-semester course. Several topics may lend themselves to individual or group projects, such as card shuffling, perfect sampling (coupling from the past), queueing theory, stochastic calculus, martingales, and stochastic differential equations. Such specialized material is contained in starred sections.

An undergraduate textbook poses many challenges. I have struggled with trying to find the right balance between theory and application, between conceptual understanding and formal proof. There are, of course, some things that cannot be said. Continuous-time processes, in particular, require advanced mathematics based on measure theory to be made precise. Where these subjects are presented I have emphasized intuition over rigor.

Following is a synopsis of the book's nine chapters.

Chapter 1 introduces stochastic and deterministic models, the generic features of stochastic processes, and simulation. This is essential material. The second part of the chapter treats conditional probability and conditional expectation, which can be reviewed at a fast pace.

The main features of discrete-time Markov chains are covered in Chapters 2 and 3. Many examples of Markov chains are introduced, some of which are referenced throughout the book. Numerical and simulation-based methods motivate the discussion of limiting behavior. In addition to basic computations, topics include: stationary distributions, strong Markov property, ergodic and absorbing chains, and time reversibility. Several important limit theorems are discussed in detail, with proofs given at the end of the chapter. Instructors may choose to limit how much time is spent on proofs.

Branching processes are the topic of Chapter 4. Although branching processes are Markov chains, the methods of analysis are different enough to warrant a separate chapter. Probability generating functions are introduced, and do not assume prior exposure.

The focus of Chapter 5 is Markov chain Monte Carlo, a relatively new topic but one with exponentially growing application. Instructors will find many subjects to pick and choose. Several case studies make for excellent classroom material, in particular (i) a randomized method for decoding text, from Diaconis (2009), and (ii) an application which combines ecology and counting matrices with fixed row and column totals, based on Cobb and Chen (2003). Other topics include coupling from the past, card shuffling, and rates of convergence of Markov chains.

Chapter 6 is devoted to the Poisson process. The approach emphasizes three alternate definitions and characterizations, based on the (i) counting process, (ii) arrival process, and (iii) infinitesimal description. Additional topics are: spatial processes, nonhomogeneous Poisson processes, embedding, and arrival time paradoxes.

Continuous-time Markov chains are discussed in Chapter 7. For continuous-time stochastic processes, here and in Chapter 8, there is an emphasis on intuition, examples, and applications. In addition to basic material, there are sections on queueing theory (with Little's formula), absorbing processes, and Poisson subordination.

Brownian motion is the topic of Chapter 8. The material is more challenging. Topics include the invariance principle, transformations, Gaussian processes, martingales, and the optional stopping theorem. Examples include scoring in basketball and animal tracking. The Black-Scholes options pricing formula is derived.

Chapter 9 is a gentle introduction to stochastic calculus. *Gentle* means no measure theory, sigma fields, or filtrations, but an emphasis on intuition, examples and applications. I decided to include this material because of its growing popularity and application. Stochastic differential equations are introduced. Simulation and numerical methods help make the topic accessible.

Book appendices include: (i) Getting started with `R`, with exercises, (ii) Probability review, with short sections on the main discrete and continuous probability distributions, (iii) Summary table of common probability distributions, and (iv) Matrix algebra review. Resources for students include a suite of `R` functions and script files for generating many of the processes from the book.

The book contains over 200 examples, and some 500 end-of-chapter exercises. A web site `www.people.carleton.edu/rdobrow/stochbook` is established. It contains errata and relevant files. All the `R` code and script files used in the book are available at this site. A solutions manual with detailed solutions to all exercises is available for instructors.

Much of this book is a reflection of my experience teaching the course over the past ten years. Here is a suggested one-semester syllabus, which I have used.

1. Introduction and review — 1.1, 1.2, 1.3 (quickly skim 1.4 and 1.5)
2. One-day introduction to `R` — Appendix A
3. Markov chains — All of chapters 2 and 3
4. Branching processes — Chapter 4
5. MCMC — 5.1, 5.2
6. Poisson process — 6.1, 6.2, 6.4, 6.5, 6.8
7. Continuous-time Markov chains — 7.1, 7.2, 7.3, 7.4
8. Brownian motion — 8.1, 8.2, 8.3, 8.4, 8.5, 8.7

If instructors have questions on syllabus, homework assignments, exams, or projects, I am happy to share resources and experiences teaching this most rewarding course.

Stochastic Processes is a great mathematical adventure. Bon voyage!