

Editors

J.E. Marsden
L. Sirovich
M. Golubitsky
W. Jäger

Advisors

G. Iooss
P. Holmes
D. Barkley
M. Dellnitz
P. Newton

Springer

New York
Berlin
Heidelberg
Hong Kong
London
Milan
Paris
Tokyo

Texts in Applied Mathematics

1. *Sirovich*: Introduction to Applied Mathematics.
2. *Wiggins*: Introduction to Applied Nonlinear Dynamical Systems and Chaos, 2nd ed.
3. *Hale/Koçak*: Dynamics and Bifurcations.
4. *Chorin/Marsden*: A Mathematical Introduction to Fluid Mechanics, 3rd ed.
5. *Hubbard/West*: Differential Equations: A Dynamical Systems Approach: Ordinary Differential Equations.
6. *Sontag*: Mathematical Control Theory: Deterministic Finite Dimensional Systems, 2nd ed.
7. *Perko*: Differential Equations and Dynamical Systems, 3rd ed.
8. *Seaborn*: Hypergeometric Functions and Their Applications.
9. *Pipkin*: A Course on Integral Equations.
10. *Hoppensteadt/Peskin*: Modeling and Simulation in Medicine and the Life Sciences, 2nd ed.
11. *Braun*: Differential Equations and Their Applications, 4th ed.
12. *Stoer/Bulirsch*: Introduction to Numerical Analysis, 3rd ed.
13. *Renardy/Rogers*: An Introduction to Partial Differential Equations.
14. *Banks*: Growth and Diffusion Phenomena: Mathematical Frameworks and Applications.
15. *Brenner/Scott*: The Mathematical Theory of Finite Element Methods, 2nd ed.
16. *Van de Velde*: Concurrent Scientific Computing.
17. *Marsden/Ratiu*: Introduction to Mechanics and Symmetry, 2nd ed.
18. *Hubbard/West*: Differential Equations: A Dynamical Systems Approach: Higher-Dimensional Systems.
19. *Kaplan/Glass*: Understanding Nonlinear Dynamics.
20. *Holmes*: Introduction to Perturbation Methods.
21. *Curtain/Zwart*: An Introduction to Infinite-Dimensional Linear Systems Theory.
22. *Thomas*: Numerical Partial Differential Equations: Finite Difference Methods.
23. *Taylor*: Partial Differential Equations: Basic Theory.
24. *Merkin*: Introduction to the Theory of Stability of Motion.
25. *Naber*: Topology, Geometry, and Gauge Fields: Foundations.
26. *Polderman/Willems*: Introduction to Mathematical Systems Theory: A Behavioral Approach.
27. *Reddy*: Introductory Functional Analysis with Applications to Boundary-Value Problems and Finite Elements.
28. *Gustafson/Wilcox*: Analytical and Computational Methods of Advanced Engineering Mathematics.
29. *Tveito/Winther*: Introduction to Partial Differential Equations: A Computational Approach.
30. *Gasquet/Witomski*: Fourier Analysis and Applications: Filtering, Numerical Computation, Wavelets.

(continued after index)

Pierre Brémaud

Markov Chains

Gibbs Fields, Monte Carlo Simulation,
and Queues

With 64 Illustrations



Springer

Pierre Brémaud
CNRS-ESE
Laboratoire des Signaux et Systèmes
Plateau de Moulon
91192 Gif-sur-Yvette
France
bremaud@lss.supelec.fr

Series Editors

J.E. Marsden
Control and Dynamical Systems, 107-81
California Institute of Technology
Pasadena, CA 91125
USA

L. Sirovich
Division of Applied Mathematics
Brown University
Providence, RI 02912
USA

M. Golubitsky
Department of Mathematics
University of Houston
Houston, TX 77204-3476
USA

W. Jäger
Department of Applied Mathematics
Universität Heidelberg
Im Neuenheimer Feld 294
69120 Heidelberg
Germany

Mathematics Subject Classification (1991): 60-01, 93Exx

Library of Congress Cataloging-in-Publication Data

Brémaud, Pierre.

Markov chains : Gibbs fields, Monte Carlo simulation, and queues /

Pierre Brémaud.

p. cm. — (Texts in applied mathematics : 31)

Includes bibliographical references (p. -) and index.

ISBN 0-387-98509-3 (acid-free paper)

I. Markov processes. I. Title. II. Series.

QA274.8.B74 1998

519.233—dc21

98-17539

Printed on acid-free paper.

© 1999 Springer-Verlag New York, Inc.

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer-Verlag New York, Inc., 175 Fifth Avenue, New York, NY 10010, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use of general descriptive names, trade names, trademarks, etc., in this publication, even if the former are not especially identified, is not to be taken as a sign that such names, as understood by the Trade Marks and Merchandise Marks Act, may accordingly be used freely by anyone.

Production managed by A. Orrantia; manufacturing supervised by Nancy Wu.

Typeset by The Bartlett Press, Inc., Marietta, GA, from the author's TeX files.

Printed and bound by Hamilton Printing, Co., Rensselaer, NY.

Printed in the United States of America.

9 8 7 6 5 4 3

ISBN 0-387-98509-3

SPIN 10946541

Springer-Verlag New York Berlin Heidelberg

A member of BertelsmannSpringer Science+Business Media GmbH

To Marion

Series Preface

Mathematics is playing an ever more important role in the physical and biological sciences, provoking a blurring of boundaries between scientific disciplines and a resurgence of interest in the modern as well as the classical techniques of applied mathematics. This renewal of interest, both in research and teaching, has led to the establishment of the series: *Texts in Applied Mathematics (TAM)*.

The development of new courses is a natural consequence of high level of excitement on the research frontier as newer techniques, such as numerical and symbolic computer systems, dynamical systems, and chaos, mix with and reinforce the traditional methods of applied mathematics. Thus, purpose of this textbook series is to meet the current and future needs of these advances and encourage the teaching of new courses.

TAM will publish textbooks suitable for use in advanced undergraduate and beginning graduate courses, and will complement the *Applied Mathematical Sciences (AMS)* series, which will focus on advanced textbooks and research level monographs.

Preface

From Pushkin to Monte Carlo

When Markov introduced his famous model in 1906, he was not preoccupied with applications. He just wanted to show that independence is not necessary for the law of large numbers (the weak law of large numbers at that time, since Borel proved the strong law for heads and tails only in 1909). An example that he considered was the alternation of consonants and vowels in Pushkin's *Eugene Onegin*, which he described as a two-state chain. (This, however, does not say much about the plot!) Almost at the same time, and faithful to the French tradition of gambling probabilists, Poincaré studied Markov chains on finite groups, with applications to card shuffling. The Austrian physicists Paul and Tatiana Ehrenfest proposed in 1907 a Markov chain model that very much helped to clarify the controversial issue of thermodynamic irreversibility. Sir Francis Galton, a cousin of Darwin, who was interested in the probability of the survival of the English peerage, was the inventor of the branching process, another famous Markov model with many applications besides the original one. He posed the problem in the *Educational Times* in 1873, and in the same year and same journal, Reverend Watson proposed the method of solution that became a textbook classic.

The dates mentioned above show that Markov models were already around even before Markov started the systematic study of this class of random sequences. However, the work of Markov challenged the best probabilists, such as Kolmogorov, Doeblin, and Fréchet, just to mention the leading pioneering figures. The outcome was a clean and sound theory ready for applications, and today, Markov chains are omnipresent in the applied sciences. For instance, *biology* is an important consumer of Markov models, many of them concerning genetics and population theory. In the *social sciences*, social mobility can be described in Markovian terms. Quantitative *psychology* uses Markovian models of learning. *Physics* is a major patron of Markov chain theory, and Markov models (for instance, the Ehrenfest diffusion model, the annealing model, and the Ising–Peierls model of phase transition) have been very useful in understanding qualitatively complex phenomena. Markov chains have found many applications in *electrical engineering*, for instance in the performance analysis of multiple access communications protocols and of communications networks, in coded modulation, and in image processing. Recently, Markov chain theory has received an additional impetus from the advent of Monte Carlo Markov chain simulation. Markov chains have found a privileged domain of application in *operations research*, in reliability theory and queuing theory for instance.

A Useful, Simple, and Beautiful Theory

The list of applications of Markov chains is virtually infinite, and one is entitled to say that it is the single most successful class of stochastic processes, its success being due to the relative simplicity of its theory and to the fact that simple Markov models can exhibit extremely varied and complex behavior. The modeling power of Markov chains may well be compared to that of ordinary differential equations.

The theory of Markov chains with a countable state space is an ideal introduction to stochastic processes; it is not protected by a wall of technicalities, and therefore the student has quick access to the main results. Indeed, the mathematical equipment required for a rewarding study of this class of models consists only of the notion of conditional independence and the strong law of large numbers.

Another pleasant feature of Markov chain theory is that this classic topic can be presented in terms of the elegant concepts of the modern theory of stochastic processes, such as reversibility, martingales, and coupling.

Basic Theory and Advanced Topics

This book is devoted to the study of homogeneous Markov chains (HMCs) with a countable state space, in discrete time and in continuous time. The table of contents reflects the recent advances of the theory and responds to a growing need for a unified treatment of related topics such as finite Gibbs fields, nonhomogeneous Markov chains, discrete-time regenerative processes, Monte Carlo simulation, simulated annealing, and queuing theory.

About half of the book is devoted to the basic theory (Chapters 2, 3, 4, and 8). This part of the book introduces *discrete-time Markov models* (Chapter 2) and gives an account of their *recurrence and ergodicity* (Chapter 3) properties as well as of their *long-run behavior* (Chapter 4). It also treats the *continuous-time Markov models* (Chapter 8). Continuous-time HMCs are essentially discrete-time HMCs with a random time scale. The time separating two successive transitions is not one unit, but it is an exponential time depending on the current state. (This quick description, of course, jumps over the fine technical details.) Another informal statement is that continuous-time HMCs are discrete-time HMCs with a little dose of Poisson processes. We shall say no more at this point except that the traditional semigroup approach of Chapter 8 is completed in Chapter 9, on *Poisson calculus and queues*, by a probabilistic approach, where continuous-time HMCs are described by a family of independent Poisson processes associated with a deterministic rule selecting state transitions. This approach has many advantages, but it is not yet standard in textbooks, and besides, it requires some mathematical maturity. It is one of the advanced topics of this book (Chapters 5, 6, 7, and 9).

At least one of three requirements must be satisfied for a topic to be called advanced in this book. It has to be mathematically subtle, or technical, or simply not a textbook standard.

The chapter on *eigenvalues and nonhomogeneous Markov chains* (Chapter 6) is, for instance, a little technical, but it is important and concerns an area of intense research activity. It deals with the topic of effective computation of convergence rates, both for convergence to steady state of an ergodic chain and for convergence of ergodic estimates, in view of applications to simulation, for instance. It also includes the basic results of nonhomogeneous Markov chain theory, which will be applied to the analysis of the simulated annealing algorithm. The chapter on *Lyapunov functions and martingales* (Chapter 5) is more mathematical, and contains a brief introduction to potential theory. It gives the powerful theorem of Foster, a sufficient condition of positive recurrence, and places it in its natural environment among martingales and potentials. It is a first contact with martingale theory whose efficiency is demonstrated by a few examples concerning the absorption problem. *Gibbs fields and Monte Carlo simulation* (Chapter 7) are very important topics, of interest in physics, image processing, and optimization. There are quite a few good reasons to include Gibbs fields in a book devoted to Markov chains, besides the observation that they generalize in a natural way the Markov definition on the real line to arbitrary discrete index sets. One of them is that they are a privileged domain of application of the *Monte Carlo* Markov chain simulation algorithms. The latter is the last topic of Chapter 7, which also includes an introduction to simulated annealing.

Chapter 1 is a *probability review*, and it has a structure slightly different from the others in that it contains worked out exercises that the reader is expected to try a few minutes before looking at the solution. An *appendix* provides the results in number theory, analysis, and linear algebra directly useful to the theory of Markov chains. Each chapter is closed by a problems section.

Acknowledgments

My colleagues David Aldous, of the University of California at Berkeley, Jim Fill, of the Johns Hopkins University, Michael Frey, of the University of Ulm, Takis Konstantopoulos, of the University of Texas at Austin, Torgny Lindvall, of the Chalmers University of Technology and the Göteborg University, and Laurent Massoulié, of the Centre National d'Etudes en Télécommunications, helped me with their corrections and suggestions. My friends in the Laboratory of Signals and Systems of the Centre National de la Recherche Scientifique Hervé Carfantan, Philippe Ciuciu, Jean-François Giovannelli, and Jérôme Idier have supported me in many ways (L^AT_EX and Xfig, corrections, coffee break cookies).

To all of them I wish to express my sincere gratitude, as well as to Pam Williams, who diligently typed a first draft of the manuscript, and to my students at the Ecole Polytechnique, at the Swiss Federal Institute of Technology of Lausanne, and at the Ecole Nationale Supérieure de Techniques Avancées, who have contributed to many improvements of the manuscript.

Gif-sur-Yvette
November, 1998

Pierre Brémaud

Contents

Series Preface	vii
Preface	ix
1 Probability Review	1
1 Basic Concepts	1
1.1 Events	1
1.2 Random Variables	3
1.3 Probability	4
2 Independence and Conditional Probability	7
2.1 Independence of Events and of Random Variables	7
2.2 Bayes's Rules	9
2.3 Markov Property	11
3 Expectation	14
3.1 Cumulative Distribution Function	14
3.2 Expectation, Mean, and Variance	15
3.3 Famous Random Variables	20
4 Random Vectors	25
4.1 Absolutely Continuous Random Vectors	25
4.2 Discrete Random Vectors	28
4.3 Product Formula for Expectation	29
5 Transforms of Probability Distributions	29
5.1 Generating Functions	29
5.2 Characteristic Functions	34
6 Transformations of Random Vectors	36
6.1 Smooth Change of Variables	36
6.2 Order Statistics	38
7 Conditional Expectation of Discrete Variables	39
7.1 Definition and Basic Properties	39
7.2 Successive Conditioning	41
8 The Strong Law of Large Numbers	42
8.1 Borel–Cantelli Lemma	42
8.2 Almost-Sure Convergence	43

8.3	Markov's Inequality	45
8.4	Proof of Kolmogorov's SLLN	47
2	Discrete-Time Markov Models	53
1	The Transition Matrix	53
1.1	Markov Property	53
1.2	Distribution of an HMC	56
2	Markov Recurrences	58
2.1	A Canonical Representation	58
2.2	A Few Famous Examples	59
3	First-Step Analysis	65
3.1	Absorption Probability	65
3.2	Mean Time to Absorption	68
4	Topology of the Transition Matrix	71
4.1	Communication	71
4.2	Period	72
5	Steady State	75
5.1	Stationarity	75
5.2	Examples	76
6	Time Reversal	80
6.1	Reversed Chain	80
6.2	Time Reversibility	81
7	Regeneration	83
7.1	Strong Markov Property	83
7.2	Regenerative Cycles	86
3	Recurrence and Ergodicity	95
1	Potential Matrix Criterion	95
1.1	Recurrent and Transient States	95
1.2	Potential Matrix	97
1.3	Structure of the Transition Matrix	100
2	Recurrence and Invariant Measures	100
3	Positive Recurrence	104
3.1	Stationary Distribution Criterion	104
3.2	Examples	105
4	Empirical Averages	110
4.1	Ergodic Theorem	110
4.2	Examples	113
4.3	Renewal Reward Theorem	117
4	Long Run Behavior	125
1	Coupling	125
1.1	Convergence in Variation	125
1.2	The Coupling Method	128

2	Convergence to Steady State	130
2.1	Positive Recurrent Case	130
2.2	Null Recurrent Case	131
2.3	Thermodynamic Irreversibility	133
2.4	Convergence Rates via Coupling	136
3	Discrete-Time Renewal Theory	137
3.1	Renewal Equation	137
3.2	Renewal Theorem	140
3.3	Defective Renewal Sequences	142
4	Regenerative Processes	145
4.1	Renewal Equation of a Regenerative Process	145
4.2	Regenerative Theorem	146
5	Life Before Absorption	149
5.1	Infinite Sojourns	149
5.2	Time to Absorption	153
6	Absorption	154
6.1	Fundamental Matrix	154
6.2	Absorption Matrix	156
5	Lyapunov Functions and Martingales	167
1	Lyapunov Functions	167
1.1	Foster's Theorem	167
1.2	Queuing Applications	173
2	Martingales and Potentials	178
2.1	Harmonic Functions and Martingales	178
2.2	The Maximum Principle	180
3	Applications of Martingales to HMCs	185
3.1	The Two Pillars of Martingale Theory	185
3.2	Transience and Recurrence via Martingales	186
3.3	Absorption via Martingales	189
6	Eigenvalues and Nonhomogeneous Markov Chains	195
1	Finite Transition Matrices	195
1.1	Perron–Frobenius Theorem	195
1.2	Quasi-stationary Distributions	199
2	Reversible Transition Matrices	201
2.1	Eigenstructure and Diagonalization	201
2.2	Spectral Theorem	204
3	Convergence Bounds Without Eigenvectors	207
3.1	Basic Bounds, Reversible Case	207
3.2	Nonreversible Case	211
4	Geometric Bounds	212
4.1	Weighted Paths	212
4.2	Conductance	215

5	Probabilistic Bounds	219
5.1	Separation and Strong Stationary Times	219
5.2	Convergence Rates via Strong Stationary Times	223
6	Fundamental Matrix of Recurrent Chains	226
6.1	Definition of the Fundamental Matrix	226
6.2	Mutual Time–Distance Matrix	230
6.3	Variance of Ergodic Estimates	232
7	The Ergodic Coefficient	235
7.1	Dobrushin’s Inequality	235
7.2	Interaction Coefficients and Coincidence	238
8	Nonhomogeneous Markov Chains	239
8.1	Ergodicity of Nonhomogeneous Markov Chains	239
8.2	Block Criterion of Weak Ergodicity	241
8.3	Sufficient Condition of Strong Ergodicity	242
7	Gibbs Fields and Monte Carlo Simulation	253
1	Markov Random Fields	253
1.1	Neighborhoods and Local Specification	253
1.2	Cliques, Potential, and Gibbs Distributions	256
2	Gibbs–Markov Equivalence	260
2.1	From the Potential to the Local Specification	260
2.2	From the Local Specification to the Potential	261
3	Image Models	268
3.1	Textures	268
3.2	Lines and Points	270
4	Bayesian Restoration of Images	275
4.1	MAP Likelihood Estimation	275
4.2	Penalty Methods	279
5	Phase Transitions	280
5.1	Spontaneous Magnetization	280
5.2	Peierls’s Argument	281
6	Gibbs Sampler	285
6.1	Simulation of Random Fields	285
6.2	Convergence Rate of the Gibbs Sampler	288
7	Monte Carlo Markov Chain Simulation	290
7.1	General Principle	290
7.2	Convergence Rates in MCMC	295
7.3	Variance of Monte Carlo Estimators	299
8	Simulated Annealing	305
8.1	Stochastic Descent and Cooling	305
8.2	Convergence of Simulated Annealing	311

8	Continuous-Time Markov Models	323
1	Poisson Processes	323
1.1	Point Processes	323
1.2	Counting Process of an HPP	324
1.3	Competing Poisson Processes	327
2	Distribution of a Continuous-Time HMC	329
2.1	Transition Semigroup	329
2.2	Infinitesimal Generator	333
3	Kolmogorov's Differential Systems	338
3.1	Finite State Space	338
3.2	General Case	340
3.3	Regular Jumps	344
4	The Regenerative Structure	345
4.1	Strong Markov Property	345
4.2	Embedded Chain and Transition Times	348
4.3	Explosions	350
5	Recurrence	357
5.1	Stationary Distribution Criterion of Ergodicity	357
5.2	Time Reversal	361
6	Long-Run Behavior	363
6.1	Ergodic Chains	363
6.2	Absorbing Chains	364
9	Poisson Calculus and Queues	369
1	Continuous-Time Markov Chains as Poisson Systems	369
1.1	Strong Markov Property of HPPs	369
1.2	From Generator to Markov Chain	372
2	Stochastic Calculus of Poisson Processes	375
2.1	Counting Integrals and the Smoothing Formula	375
2.2	Kolmogorov's Forward System via Poisson Calculus	378
2.3	Watanabe's Characterization of Poisson Processes	380
3	Poisson Systems	383
3.1	The Purely Poissonian Description	383
3.2	The GSMP construction	385
3.3	Markovian Queues as Poisson Systems	388
4	Markovian Queuing Theory	394
4.1	Isolated Markovian Queues	394
4.2	The $M/G/1/\infty$ /FIFO Queue	398
4.3	The $GI/M/1/\infty$ /FIFO Queue	402
4.4	Markovian Queuing Networks	407

Appendix	417
1 Number Theory and Calculus	417
1.1 Greatest Common Divisor	417
1.2 Abel's Theorem	418
1.3 Lebesgue's Theorems for Series	420
1.4 Infinite Products	422
1.5 Tychonov's Theorem	423
1.6 Subadditive Functions	423
2 Linear Algebra	424
2.1 Eigenvalues and Eigenvectors	424
2.2 Exponential of a Matrix	426
2.3 Gershgorin's Bound	427
3 Probability	428
3.1 Expectation Revisited	428
3.2 Lebesgue's Theorems for Expectation	430
Bibliography	433
Author Index	439
Subject Index	441

Probability Review

1 Basic Concepts

The present chapter reviews the notions and results of probability theory directly useful in the sequel. This book is self-contained and the review of probability that follows is sufficient for the purpose of studying Markov chains.

1.1 Events

Probability theory provides a mathematical framework for the study of random phenomena. It requires a precise description of the *outcome* of an observation when such a phenomenon is observed. The collection of all possible outcomes ω is called the *sample space* Ω .

Example 1.1. *A Die*


The experiment consists in tossing a die once. The possible outcomes are $\omega = 1, 2, \dots, 6$ and the sample space is the set $\Omega = \{1, 2, 3, 4, 5, 6\}$ (see Fig. 1.1.1). \diamond

Example 1.2. *Darts*


The experiment consists in throwing a dart at a very large wall. The sample space can be chosen to be the plane \mathbb{R}^2 , an idealization of the wall, an outcome being the position $\omega = (x, y)$ hit by the dart (see Fig. 1.1.1). \diamond

Example 1.3. *Coin Tosses*

The experiment is an infinite sequence of coin tosses. One can take for the sample space Ω the collection of all sequences $\omega = \{x_n\}_{n \geq 1}$, where $x_n = 1$ or 0 , depending on whether the n th toss results in heads or tails (see Fig. 1.1.1). \diamond



The samples in Example 1.1



A sample in Example 1.2

HTTTHHTTTHTHTHT...

A sample in Example 1.3

Figure 1.1.1. Samples

The Logics of Events

Probabilists have their own interpretation of sets. It is reflected in their terminology, which we shall now review.

Any subset A of the sample space Ω can be regarded as a representation of some *event*. In Example 1.1, $A = \{1, 3, 5\}$ is the event “result is odd.” In Example 1.2, subset $A = \{(x, y); x^2 + y^2 \geq 1\}$ could be the event “you missed the dartboard” if the dartboard is centered at 0 and of radius 1. In Example 1.3, subset $A = \{\omega; x_k = 1 \text{ for } k = 1 \text{ to } 1,000\}$ is a very lucky event if you bet on heads.

One says that outcome ω *realizes* event A if $\omega \in A$. Obviously, if ω does not realize A , it realizes \bar{A} , the complement of A in Ω .

Event $A \cap B$ is realized by outcome ω if and only if ω realizes both A and B . Similarly, $A \cup B$ is realized by ω if and only if at least one event among A and B is realized.

Two events A and B are called *incompatible* when $A \cap B = \emptyset$. In other words, event $A \cap B$ is impossible, and so no outcome ω can realize both A and B . For this reason one refers to the empty set ϕ as the *impossible event*. Naturally, Ω is called the *certain event*.




Figure 1.1.2. The logic of events

The notation $A + B$ (the *sum* of A and B) implies by convention that A and B are disjoint, and represents the union $A \cup B$. Similarly, the notation $\sum_{k=1}^{\infty} A_k$ is used for $\cup_{k=1}^{\infty} A_k$ only when the A_k 's are pairwise disjoint. The equality

$$\sum_{k=0}^{\infty} A_k = \Omega$$

means that the events A_1, A_2, \dots are *mutually incompatible and exhaustive*. They are exhaustive in the sense that any outcome ω realizes at least one among them. They are mutually exclusive in the sense that any two events among them are incompatible. Therefore, any ω realizes one and only one of the events A_1, \dots, A_n .

If $B \subset A$, event B is said to *imply* event A , because A is realized by ω whenever B is realized by ω . The notation $A - B$ is used only if $B \subset A$, and it stands for $A \cap \bar{B}$. In particular, if $B \subset A$, then $A = B + (A - B)$ (see Fig. 1.1.2).

Probability theory assigns to an event a number, the *probability* of the said event. For technical reasons, the collection \mathcal{F} of events that are assigned a probability is not always identical to the collection of all subsets of Ω . The requirements on \mathcal{F} are the following:

1. The impossible event \emptyset and the certain event Ω are in \mathcal{F}
2. If A is in \mathcal{F} , then so is \bar{A}
3. If A_1, A_2, \dots are in \mathcal{F} , then so is $\cup_{k=1}^{\infty} A_k$

One calls the collection of subsets \mathcal{F} a *sigma field* on Ω , here the sigma-field of *events*.

If the sample space Ω is finite, one usually considers any subset of Ω to be an event. The same is generally true for a countable sample space.

1.2 Random Variables

Definition 1.1. *Random Variables*

A *random variable* is a function $X : \Omega \rightarrow \bar{\mathbb{R}}$ such that for all $a \in \mathbb{R}$, the event $\{X \leq a\} = \{\omega; X(\omega) \leq a\}$ can be assigned a probability, that is,

$$\{X \leq a\} \in \mathcal{F}.$$

A function $X : \Omega \rightarrow E$ where E is a denumerable set is called a *discrete random variable* if for all $x \in E$

$$\{X = x\} \in \mathcal{F}.$$

Sometimes a random variable is called a *random number*. This is an innocuous habit as long as one is aware that it is not the function X that is random, but the outcome ω , which in turn makes the number $X(\omega)$ random. If $X(\omega)$ is real for all ω , then X is called a *real random variable*.

Since a denumerable set E can be identified with \mathbb{N} , we shall consider a discrete random variable to be a special case of a real random variable.

Example 1.4. *A Die*

This is a continuation of Example 1.1. Take for X the identity $X(\omega) = \omega$. In that sense X is a random number obtained by tossing a die. \diamond

Example 1.5. *Darts*

This is a continuation of Example 1.2. Here $\omega = (x, y)$, where $x, y \in \mathbb{R}$. Define the *coordinate* random variables of Ω , X , and Y by

$$X(\omega) = x, Y(\omega) = y. \quad \diamond$$

Example 1.6. *Coin Tosses*

This is a continuation of Example 1.3. Here $\omega = \{x_n\}_{n \geq 1}$. Define X_n to be the random number obtained at the n th toss:

$$X_n(\omega) = x_n. \quad \diamond$$

1.3 Probability

The *probability* $P(A)$ of an event $A \in \mathcal{F}$ measures the likeliness of its occurrence. As a function defined on \mathcal{F} , the probability P is required to satisfy a few properties, the *axioms of probability*.

Definition 1.2. *Axioms of Probability*

A probability (measure) on (Ω, \mathcal{F}) is a mapping $P : \mathcal{F} \rightarrow \mathbb{R}$ such that

1. $0 \leq P(A) \leq 1$
2. $P(\Omega) = 1$
3. $P(\sum_{k=1}^{\infty} A_k) = \sum_{k=1}^{\infty} P(A_k)$

Property 3 is called *sigma additivity*. The triple (Ω, \mathcal{F}, P) is called a *probability space*, or *probability model*.

The axioms of probability are motivated by the following heuristic interpretation of $P(A)$ as the empirical frequency of occurrence of event A . If n “independent” experiments are performed, among which n_A result in the realization of A , then the empirical frequency

$$F(A) = \frac{n_A}{n}$$

should be close to $P(A)$ if n is “sufficiently large.” Clearly, the function F satisfies the axioms.

The axiomatic presentation of probability theory is nevertheless logically independent of the frequency interpretation. As a matter of fact, its success is due to its apparent ignorance of the frequency interpretation, which blurs the picture because the empirical frequency F depends on too many things: on the number of experiments and on the experiments themselves. The axiomatic theory of probability connects to the frequency interpretation *a posteriori*: the latter appears as a *theorem*, the famous *strong law of large numbers* (SLLN) given in Section 8. To obtain it, all that is needed besides the axioms of probability and clever computations is a good definition of what is meant by independent experiments. This definition will be given in the next section.

Example 1.7. A Die

This is a continuation of Example 1.4. For $A \subset \Omega = \{1, 2, 3, 4, 5, 6\}$, the formula

$$P(A) = \frac{|A|}{6},$$

where $|A|$ is the *cardinal* of A , that is, the number of elements in A , defines a probability P . \diamond

Example 1.8. Darts

This is a continuation of Example 1.5. Take, for instance,

$$P(A) = \frac{1}{2\pi\sigma^2} \int \int_A e^{-\frac{(x^2+y^2)}{2\sigma^2}} dx dy.$$

It can be checked that $P(\Omega) = 1$. Indeed,

$$\begin{aligned} \frac{1}{2\pi\sigma^2} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} e^{-\frac{x^2+y^2}{2\sigma^2}} dx dy &= \frac{1}{\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} e^{-(x^2+y^2)} dx dy \\ &= \frac{1}{2\pi} \int_0^{2\pi} \int_0^{+\infty} e^{-\rho^2} 2\rho d\rho d\theta \\ &= \frac{1}{2\pi} \int_0^{2\pi} \int_0^{+\infty} e^{-u} du d\theta = 1. \end{aligned} \quad \diamond$$

Example 1.9. *Coin Tosses*

This is a continuation of Example 1.6. Choose probability P such that for any event of the form $A = \{x_1 = a_1, \dots, x_k = a_k\}$, where a_1, \dots, a_k are arbitrary in $\{0, 1\}$,

$$P(A) = \frac{1}{2^k}. \quad \diamond$$

The choices for probability P in these examples is arbitrary, and many other choices are possible. That of Example 1.7 suggests an unbiased die, where each outcome 1, 2, 3, 4, 5, or 6 has the same probability. Probability P of Example 1.9 implies an unbiased coin and independent tosses (see Example 2.2).

From the axioms of probability, the following properties are easy to check:

$$P(\bar{A}) = 1 - P(A), \quad (1.1)$$

$$P(\emptyset) = 0, \quad (1.2)$$

$$A \subset B \Rightarrow P(A) \leq P(B), \quad (1.3)$$

$$P(\cup_{k=1}^{\infty} A_k) \leq \sum_{k=1}^{\infty} P(A_k). \quad (1.4)$$

Proof. For (1.1), use additivity: $1 = P(\Omega) = P(A + \bar{A}) = P(A) + P(\bar{A})$. For (1.2), use (1.1) with $A = \Omega$: $P(\emptyset) = 1 - P(\Omega) = 1 - 1 = 0$. For (1.3), write $P(B) = P(A) + P(B - A) \geq P(A)$. Finally, for (1.4), observe that

$$\cup_{k=1}^{\infty} A_k = \sum_{k=1}^{\infty} A'_k,$$

where

$$A'_k = A_k \cap \left[\overline{\cup_{i=1}^{k-1} A_i} \right].$$

Therefore,

$$P(\cup_{k=1}^{\infty} A_k) = P\left(\sum_{k=1}^{\infty} A'_k\right) = \sum_{k=1}^{\infty} P(A'_k).$$

But $A'_k \subset A_k$, and therefore $P(A'_k) \leq P(A_k)$. □

Property (1.3) is the *monotonicity* property, and (1.4) is the *sub- σ -additivity* property.

Theorem 1.1. *Monotone Sequential Continuity*

Let $\{A_n\}_{n \geq 1}$ be a nondecreasing sequence of events, i.e., for all $n \geq 1$,

$$A_{n+1} \supset A_n.$$

Then

$$P(\cup_{n=1}^{\infty} A_n) = \lim_{n \uparrow \infty} \uparrow P(A_n) \quad (1.5)$$

(where $\lim \uparrow a_n$ is a notation emphasizing that the sequence $\{a_n\}$ is nondecreasing). Similarly, if $\{B_n\}_{n \geq 1}$ is a nonincreasing sequence of events, i.e., for all $n \geq 1$,

$$B_{n+1} \subset B_n,$$

then

$$P(\bigcap_{n=1}^{\infty} B_n) = \lim_{n \uparrow \infty} \downarrow P(B_n). \quad (1.6)$$

Proof. For (1.5), write

$$A_n = A_1 + (A_2 - A_1) + \cdots + (A_n - A_{n-1})$$

and

$$\bigcup_{k=1}^{\infty} A_k = A_1 + (A_2 - A_1) + (A_3 - A_2) + \cdots.$$

Therefore,

$$\begin{aligned} P(\bigcup_{k=1}^{\infty} A_k) &= P(A_1) + \sum_{j=1}^{\infty} P(A_j - A_{j-1}) \\ &= \lim_{n \uparrow \infty} \left\{ P(A_1) + \sum_{j=1}^n P(A_j - A_{j-1}) \right\} = \lim_{n \uparrow \infty} P(A_n). \end{aligned}$$

For (1.6), write

$$P(\bigcap_{n=1}^{\infty} B_n) = 1 - P(\overline{\bigcap_{n=1}^{\infty} B_n}) = 1 - P(\bigcup_{n=1}^{\infty} \bar{B}_n)$$

and apply (1.5) with $A_n = \bar{B}_n$. □

2 Independence and Conditional Probability

2.1 Independence of Events and of Random Variables

Definition 2.1. *Independence*

Two events A and B are called *independent* if

$$P(A \cap B) = P(A)P(B). \quad (2.1)$$

Two random variables X and Y are called independent if for all $a, b \in \mathbb{R}$,

$$P(X \leq a, Y \leq b) = P(X \leq a)P(Y \leq b). \quad (2.2)$$

In the above definition, we used the following notation: $P(X \leq a) = P(\{X \leq a\})$ where $\{X \leq a\} = \{\omega ; X(\omega) \leq a\}$. Also, the left-hand side of (2.2) is $P(\{X \leq a\} \cap \{Y \leq b\})$. This is a general feature of the notational system: commas often replace intersection signs. For instance, $P(A, B)$ is the probability that both events A and B occur.

Example 2.1. Darts

This is a continuation of Example 1.8. The coordinate random variables X and Y are independent. Indeed, $\{(x, y) \in \mathbb{R}^2; x \leq a, y \leq b\} = \{X \leq a\} \cap \{Y \leq b\}$, and therefore

$$\begin{aligned} P(\{X \leq a\} \cap \{Y \leq b\}) &= \frac{1}{2\pi\sigma^2} \int_{-\infty}^a \int_{-\infty}^b e^{-\frac{x^2+y^2}{2\sigma^2}} dx dy \\ &= \left(\frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^a e^{-\frac{x^2}{2\sigma^2}} dx \right) \left(\frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^b e^{-\frac{y^2}{2\sigma^2}} dy \right) \\ &= P(X \leq a)P(Y \leq b), \end{aligned}$$

where the last equality comes from

$$P(X \leq a) = \lim_{n \uparrow \infty} P(X \leq a, Y \leq n)$$

(sequential continuity) and the identity

$$\frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{+\infty} e^{-\frac{y^2}{2\sigma^2}} dy = 1$$

(see Example 1.8). ◇

Exercise 2.1. Incompatibility Is Not Independence

Show that two incompatible events A and B are independent if and only if at least one of them has zero probability.

Solution. $P(A \cap B) = P(\emptyset) = 0$. Therefore (2.1) holds if and only if $P(A)P(B) = 0$.

The notion of independence carries over to families of events.

Definition 2.2. Independence of Families

A family $\{A_i\}_{i \in I}$ of events indexed by an arbitrary set I is called independent if for any *finite* collection of *distinct* events A_{i_1}, \dots, A_{i_k} from this family,

$$P\left(\bigcap_{j=1}^k A_{i_j}\right) = \prod_{j=1}^k P(A_{i_j}). \quad (2.3)$$

A family $\{X_i\}_{i \in I}$ of real random variables indexed by an arbitrary set I is called independent if for any *finite* collection of *distinct* random variables X_{i_1}, \dots, X_{i_k} from this family,

$$P\left(\bigcap_{j=1}^k \{X_{i_j} \leq a_j\}\right) = \prod_{j=1}^k P(X_{i_j} \leq a_j) \quad (2.4)$$

for all real numbers a_1, \dots, a_k . The family of real random variables $\{X_i\}_{i \in I}$ is said to be independent from the family of real random variables $\{Y_j\}_{j \in J}$ if

$$P\left(\left(\bigcap_{\ell=1}^r \{X_{i_\ell} \leq a_\ell\}\right) \cap \left(\bigcap_{m=1}^s \{Y_{j_m} \leq b_m\}\right)\right) = P\left(\bigcap_{\ell=1}^r \{X_{i_\ell} \leq a_\ell\}\right) P\left(\bigcap_{m=1}^s \{Y_{j_m} \leq b_m\}\right) \quad (2.5)$$

for all indices $i_1, \dots, i_r \in I$ and $j_1, \dots, j_s \in J$, and all real numbers $a_1, \dots, a_r, b_1, \dots, b_s$. Here I and J are arbitrary index sets.

Example 2.2. Coin Tosses

This is a continuation of Example 1.9. Event $\{X_k = a_k\}$ is the direct sum of events $\{X_1 = a_1, \dots, X_{k-1} = a_{k-1}, X_k = a_k\}$ for all possible values of (a_1, \dots, a_{k-1}) . Since there are 2^{k-1} such values and each one has probability 2^{-k} , we have $P(X_k = a_k) = 2^{k-1} 2^{-k}$, that is,

$$P(X_k = 1) = P(X_k = 0) = \frac{1}{2}.$$

Therefore,

$$P(X_1 = a_1, \dots, X_k = a_k) = P(X_1 = a_1)P(X_k = a_k)$$

for all $a_1, \dots, a_k \in \{0, 1\}$, from which it follows by definition that X_1, \dots, X_k are independent random variables, and more generally that $\{X_n\}_{n \geq 1}$ is a family of independent random variables. Therefore, we have a model for *independent* tosses of an *unbiased* coin. \diamond

2.2 Bayes's Rules

In the frequency interpretation, the definition of independence (2.1) reads, in rough and imprecise terms and using the notation introduced in Section 1, $n_{A \cap B}/n \approx (n_A/n) \cdot (n_B/n)$, or

$$\frac{n_{A \cap B}}{n_B} \approx \frac{n_A}{n}$$

(here \approx is a “fuzzy” version of the equality sign). Therefore, statistics relative to A do not vary when performed on a neutral sample of population or on a selected sample of population characterized by the property B . For example, the proportion of people with a family name beginning with H is the same among a large population with the usual mix of men and women as it would be among a large all-male population. This is very much the intuitive notion of independence.

Dependence between A and B occurs when $P(A \cap B) \neq P(A)P(B)$. In this case the relative frequency $n_{A \cap B}/n_B \approx P(A \cap B)/P(B)$ is different from the frequency n_A/n .

Definition 2.3. Conditional Probability

The *conditional probability* of A given B is the number

$$P(A | B) \stackrel{\text{def}}{=} \frac{P(A \cap B)}{P(B)}, \quad (2.6)$$

defined when $P(B) > 0$.

The quantity $P(A | B)$ represents our expectation of A being realized when the only available information is that B is realized. Indeed, this expectation would then be based upon the relative frequency $n_{A \cap B}/n_B$ alone.

A symmetric form of (2.6) is

$$P(A \cap B) = P(A | B)P(B) = P(B | A)P(A). \quad (2.7)$$

Of course, if A and B are independent, then

$$P(A | B) = P(A), P(B | A) = P(B). \quad (2.8)$$

Probability theory is primarily concerned with the computation of probabilities of complex events. The following formulas are useful additions to the tool kit already containing rules (1.1) to (1.6).

Theorem 2.1. Bayes's Rules

Bayes's rule of retrodiction. With $P(A) > 0$, we have

$$P(B | A) = \frac{P(A | B)P(B)}{P(A)}, \quad (2.9)$$

a rephrasing of (2.7).

Bayes's rule of exclusive and exhaustive causes. For B_1, B_2, \dots such that

$$\sum_{i=1}^{\infty} B_i = \Omega$$

and for all A ,

$$P(A) = \sum_{i=1}^{\infty} P(A | B_i)P(B_i). \quad (2.10)$$

Bayes's sequential formula. For any sequence of events A_1, \dots, A_n ,

$$P(\cap_{i=1}^k A_i) = P(A_1)P(A_2 | A_1)P(A_3 | A_1 \cap A_2) \cdots P(A_k | \cap_{i=1}^{k-1} A_i). \quad (2.11)$$

Proof. For (2.10), just decompose A :

$$\begin{aligned} P(A) &= P\left(A \cap \left(\sum_{i=1}^{\infty} B_i\right)\right) = P\left(\sum_{i=1}^{\infty} (A \cap B_i)\right) \\ &= \sum_{i=1}^{\infty} P(A \cap B_i) = \sum_{i=1}^{\infty} P(A | B_i)P(B_i). \end{aligned}$$

For (2.11), proceed by induction. Suppose that (2.11) is true for k . Write

$$P(\cap_{i=1}^{k+1} A_i) = P((\cap_{i=1}^k A_i) \cap A_{k+1}) = P(A_{k+1} | \cap_{i=1}^k A_i) P(\cap_{i=1}^k A_i),$$

and replace $P(\cap_{i=1}^k A_i)$ by the assumed equality (2.11). \square

Exercise 2.2. Medical Tests

Doctors apply a test that if the patient is affected by the disease being tracked gives a positive result in 99% of the cases. However, it happens in 2% of the cases that a healthy patient has a positive test. Statistical data show that one individual out of 1000 has the disease. What is the probability for a patient with a positive test to be affected by the disease?

Solution. Let M be the event “patient is ill,” and $+$ and $-$ the events “test is positive” and “test is negative,” respectively. We have the data

$$P(M) = 0.001, \quad P(+ | M) = 0.99, \quad P(+ | \bar{M}) = 0.02,$$

and we must compute $P(M | +)$. By the retrodiction formula,

$$P(M | +) = \frac{P(+ | M)P(M)}{P(+)}.$$

By the formula of exclusive and exhaustive causes,

$$P(+) = P(+ | M)P(M) + P(+ | \bar{M})P(\bar{M}).$$

Therefore,

$$P(M | +) = \frac{(0.99)(0.001)}{(0.99)(0.001) + (0.02)(0.999)} \approx \frac{1}{20}.$$

This is a low probability indeed. However, the situation is not as bad as it seems, because in view of the ambiguous results, all positive patients will be tested once again, this time with a better test (but more expensive, one expects). The gain will be that only a small portion of the population will be given the expensive second test. Indeed, $P(+) = (0.99)(0.001) + (0.02)(0.999) \approx 0.003$. One possible cause of a positive test for a healthy patient is the practice of group testing, when the biological samples of several patients are mixed.

We now proceed to introduce the central concept of Markov chain theory.

2.3 Markov Property

Definition 2.4. *Conditional Independence*

One says that A and B are *conditionally independent given C* if

$$P(A \cap B | C) = P(A | C)P(B | C). \quad (2.12)$$

Let X, Y, Z be random variables taking their values in the denumerable sets E, F, G , respectively. One says that X and Y are *conditionally independent given Z* if for all x, y, z in E, F, G , respectively, events $\{X = x\}$ and $\{Y = y\}$ are conditionally independent given $\{Z = z\}$.

Note that (exercise) for fixed C ,

$$P_C(A) \stackrel{\text{def}}{=} P(A | C) \quad (2.13)$$

defines a probability P_C . Equality (2.12) expresses the independence of A and B with respect to this probability.

Theorem 2.2. *Markov Property for Events*

Let A_1, A_2, A_3 be three events of positive probability. Events A_1 and A_3 are conditionally independent given A_2 if and only if the “Markov property” holds, that is,

$$P(A_3 | A_1 \cap A_2) = P(A_3 | A_2).$$

Proof. Assume conditional independence. Then

$$\begin{aligned} P(A_3 | A_1 \cap A_2) &= P(A_1 \cap A_2 \cap A_3) / P(A_1 \cap A_2) = P(A_1 \cap A_3 | A_2) P(A_2) / P(A_1 \cap A_2) \\ &= P(A_1 | A_2) P(A_3 | A_2) P(A_2) / P(A_1 \cap A_2) \\ &= P(A_1 | A_2) P(A_3 | A_2) / P(A_1 | A_2) \\ &= P(A_3 | A_2). \end{aligned}$$

Similar computations yield the converse implication. □

The next result is a characterization of conditional independence for random variables. It is a simple result that will be frequently used in the sequel. Recall the notation with commas instead of intersection symbols, for instance, $P(A, B | C, D) = P(A \cap B | C \cap D)$.

Theorem 2.3. *Markov Property for Random Variables*

Let $X, Y,$ and Z be three discrete random variables with values in $E, F,$ and $G,$ respectively. If for some function $g : E \times F \rightarrow [0, 1], P(X = x | Y = y, Z = z) = g(x, y)$ for all $x, y, z,$ then $P(X = x | Y = y) = g(x, y)$ for all $x, y,$ and X and Y are conditionally independent given $Z.$

Proof. We have

$$\begin{aligned} P(X = x, Y = y) &= \sum_z P(X = x, Y = y, Z = z) \\ &= \sum_z P(X = x | Y = y, Z = z) P(Y = y, Z = z) \\ &= g(x, y) \sum_z P(Y = y, Z = z) \\ &= g(x, y) P(Y = y). \end{aligned}$$

Therefore,

$$P(X = x | Y = y) = g(x, y) = P(X = x | Y = y, Z = z).$$

The conclusion then follows from Theorem 2.2. □

Exercise 2.3. *Cheap Watches*

Two factories A and B manufacture watches. Factory A produces on the average one defective item out of 100, and B produces one bad watch out of 200. A retailer receives a container of watches from one of the two above factories, but he does not know which. He checks the first watch. It works! What is the probability that the second watch he will check is good?

Solution. Let X_n be the state of the n th watch in the container, with $X_n = 1$ if it works and $X_n = 0$ if it does not. Let Y be the factory of origin. We express our a priori ignorance of where the case comes from by

$$P(Y = A) = P(Y = B) = \frac{1}{2}.$$

Also, we assume that given $Y = A$ (resp., $Y = B$), the states of the successive watches are independent. For instance,

$$P(X_1 = 1, X_2 = 0 | Y = A) = P(X_1 = 1 | Y = A)P(X_2 = 0 | Y = A).$$

We have the data

$$P(X_n = 0 | Y = A) = 0.01, P(X_n = 0 | Y = B) = 0.005.$$

We are required to compute $P(X_2 = 1 | X_1 = 1)$, that is, $P(X_1 = 1, X_2 = 1)/P(X_1 = 1)$. By the formula of exclusive and exhaustive causes, the numerator of the last fraction equals

$$P(X_1 = 1, X_2 = 1 | Y = A)P(Y = A) + P(X_1 = 1, X_2 = 1 | Y = B)P(Y = B),$$

that is, $\frac{1}{2}(99/100)^2 + \frac{1}{2}(199/200)^2$, and the denominator is

$$P(X_1 = 1 | Y = A)P(Y = A) + P(X_1 = 1 | Y = B)P(Y = B),$$

that is, $\frac{1}{2}(99/100) + \frac{1}{2}(199/200)$. Therefore,

$$P(X_2 = 1 | X_1 = 1) = \frac{\left(\frac{99}{100}\right)^2 + \left(\frac{199}{200}\right)^2}{\frac{99}{100} + \frac{199}{200}}.$$

The states of the two watches are not independent. Indeed, if they were, then

$$P(X_2 = 1 | X_1 = 1) = P(X_2 = 1) = \frac{1}{2} \left(\frac{99}{100} \right) + \frac{1}{2} \left(\frac{199}{200} \right),$$

a result different from what we obtained. This shows that even though for some given event C , two events A and B can be conditionally independent given C and conditionally independent given \bar{C} , yet they need *not* be unconditionally independent.

3 Expectation

3.1 Cumulative Distribution Function

From the probabilistic point of view, a random variable X is described by its *cumulative distribution function* (c.d.f.)

$$F(x) = P(X \leq x). \quad (3.1)$$

The cumulative distribution function has the following properties:

- (i) $F : \mathbb{R} \rightarrow [0, 1]$.
- (ii) F is nondecreasing.
- (iii) F is right-continuous.
- (iv) $F(+\infty) \stackrel{\text{def}}{=} \lim_{a \uparrow \infty} F(a) = P(X < \infty)$.
- (v) $F(-\infty) \stackrel{\text{def}}{=} \lim_{a \downarrow -\infty} F(a) = P(X = -\infty)$.

The proofs of the above statements follow:

- (i) Obvious.
- (ii) If $a \leq b$, then $\{X \leq a\} \subset \{X \leq b\}$, and therefore, by (1.3), $P(X \leq a) \leq P(X \leq b)$.
- (iii) Apply (1.6) with $B_n = \{X \leq a + \frac{1}{n}\}$. Since $\bigcap_{n \geq 1} \{X \leq a + \frac{1}{n}\} = \{X \leq a\}$, we have $\lim_{n \uparrow \infty} P(X \leq a + \frac{1}{n}) = P(X \leq a)$.
- (iv) Apply (1.5) with $A_n = \{X \leq n\}$ and observe that $\bigcup_{n=1}^{\infty} \{X \leq n\} = \{X < \infty\}$.
- (v) Apply (1.6) with $B_n = \{X \leq -n\}$ and observe that $\bigcap_{n=1}^{\infty} \{X \leq -n\} = \{X = -\infty\}$. \square

Since F is nondecreasing, the following limit exists for all $x \in \mathbb{R}$:

$$F(x-) \stackrel{\text{def}}{=} \lim_{h \downarrow 0} F(x - h). \quad (3.2)$$

The sequence $B_n = \{a - \frac{1}{n} < X \leq a\}$ is decreasing, and $\bigcap_{n=1}^{\infty} B_n = \{X = a\}$. Therefore, by the monotone sequential continuity property (1.5),

$$P(X = a) = \lim_{n \uparrow \infty} P\left(a - \frac{1}{n} < X \leq a\right) = \lim_{n \uparrow \infty} \left(F(a) - F\left(a - \frac{1}{n}\right)\right),$$

that is to say,

$$P(X = a) = F(a) - F(a-). \quad (3.3)$$

In particular, F is continuous if and only if $P(X = a) = 0$ for every fixed point $a \in \mathbb{R}$. From (3.3) and the definition of F , we have

$$\begin{aligned} P(a < X \leq b) &= F(b) - F(a), \\ P(a \leq X < b) &= F(b-) - F(a-), \\ P(a \leq X \leq b) &= F(b) - F(a-), \\ P(a < X < b) &= F(b-) - F(a). \end{aligned}$$

Also, recall that

$$\begin{aligned} P(X = +\infty) &= 1 - F(+\infty), \\ P(X = -\infty) &= F(-\infty). \end{aligned}$$

A special case of continuous c.d.f. is

$$F(x) = \int_{-\infty}^x f(y)dy \quad (3.4)$$

for some function $f \geq 0$ called the *probability density function* (p.d.f.) of X . Note that a continuous c.d.f. F need not admit a representation such as (3.4). If (3.4) holds, the c.d.f. F and the random variable X are called *absolutely continuous* (a.c.).

3.2 Expectation, Mean, and Variance

The discontinuity points of F are denumerable, as is the case for any bounded nondecreasing function. Calling $\{d_n\}_{n \geq 1}$ the sequence of discontinuity points, define

$$F_d(t) = \sum_{d_n \leq t} \{F(d_n) - F(d_{n-})\} \quad (3.5)$$

and

$$F_c(t) = F(t) - F_d(t), \quad (3.6)$$

the discontinuous and continuous components of F , respectively. We shall use the symbol

$$\int_{-\infty}^{+\infty} g(x)dF(x) \quad (3.7)$$

to represent the Stieltjes–Lebesgue integral with respect to F of the function g (see Appendix). In the special case where the continuous component of the c.d.f. is in fact absolutely continuous, i.e.,

$$F_c(x) = \int_{-\infty}^x f_c(y)dy,$$

it suffices for our purpose to interpret (3.7) by

$$\int_{-\infty}^{+\infty} g(x)dF(x) = \sum_{n=1}^{\infty} g(d_n)(F(d_n) - F(d_{n-})) + \int_{-\infty}^{+\infty} g(x)f_c(x)dx.$$

The most frequent cases arising are the purely discontinuous case where $F(t) = F_d(t)$, for which

$$\int_{-\infty}^{+\infty} g(x)dF(x) = \sum_{n=1}^{\infty} g(d_n)\{F(d_n) - F(d_{n-})\}, \quad (3.8)$$

and the absolutely continuous case, for which

$$\int_{-\infty}^{+\infty} g(x)dF(x) = \int_{-\infty}^{+\infty} g(x)f(x)dx. \quad (3.9)$$

If X takes the value $+\infty$ with positive probability, $+\infty$ should be considered conventionally as a discontinuity point, and a sum such as (3.8) will include the term $g(\infty)(1 - F(\infty))$. Of course, $g(\infty)$ must be well-defined.




Figure 1.3.1. Sum of 2 uniform random variables

Example 3.1. *Sum of Independent Uniforms*

For a random point inside the unit square $[0, 1]^2 = [0, 1] \times [0, 1]$, the following model is chosen: $\Omega = [0, 1]^2$, $P(A) = \text{area of } A$. Let X and Y be the coordinate random variables, and let $Z = X + Y$ (see Fig. 1.3.1).

The c.d.f. of Z is easily computed. For instance, when $a > 1$,

$$F_Z(a) \stackrel{\text{def}}{=} P(Z \leq a) = P(X + Y \leq a) = P(\{\omega; X(\omega) + Y(\omega) \leq a\})$$

is just the area of the domain $[0, 1]^2 \cap \{x + y \leq a\}$ (see Fig. 1.3.1), and

$$F_Z(a) = \begin{cases} \frac{1}{2}a^2 & \text{if } 0 \leq a \leq 1, \\ 1 - \frac{1}{2}(2 - a)^2 & \text{if } 1 \leq a \leq 2. \end{cases}$$

The random variable Z admits a probability density given by

$$f_Z(x) = \frac{dF_Z(x)}{dx}$$

that is

$$f_Z(x) = \begin{cases} x & \text{if } 0 \leq x \leq 1, \\ 2 - x & \text{if } 1 \leq x \leq 2. \end{cases} \quad \diamond$$

When a random variable takes its values in a denumerable set, its c.d.f. reduces to the discontinuous part F_d , and the discontinuity points d_n are the values taken by the random variable. In particular,

$$P(X = d_n) = F(d_n) - F(d_{n-}). \quad (3.10)$$

For such random variables, the *probability distribution* $\{p(d_n)\}_{n \geq 1}$, where

$$p(d_n) = P(X = d_n),$$

is usually preferred to the cumulative distribution function to describe its probabilistic behavior. In this notation,

$$\int_{-\infty}^{+\infty} g(x) dF(x) = \sum_{n=1}^{\infty} g(d_n) p(d_n).$$

Example 3.2. Coin Tosses

This is a continuation of Example 2.2. The number of occurrences of heads in n is $S_n = X_1 + \dots + X_n$ tosses. This random variable takes the integer values from 0 to n . The event $\{S_n = k\}$ is “ k among X_1, \dots, X_n are equal to 1.” There are $\binom{n}{k}$ distinct ways of assigning k values of 1 and $n - k$ values of 0 to X_1, \dots, X_n , and all have the same probability 2^{-n} . Therefore,

$$P(S_n = k) = \binom{n}{k} \frac{1}{2^n}. \quad \diamond$$

Definition 3.1. Expectation

Let X be a random variable with the c.d.f. $F(x) = P(X \leq x)$ and let $g: \bar{\mathbb{R}} \rightarrow \mathbb{R}$ be such that

$$\int_{-\infty}^{+\infty} |g(x)| dF(x) < \infty. \quad (3.11)$$

Then one defines $E[g(X)]$, the expectation of $g(X)$, by the formula

$$E[g(X)] = \int_{-\infty}^{+\infty} g(x) dF(x). \quad (3.12)$$

Example 3.3. Sum of Independent Uniforms

This is a continuation of Example 3.1.

$$E[Z^2] = \int_{-\infty}^{+\infty} x^2 f(x) dx = \int_0^1 x^3 dx + \int_1^2 x^2(2-x) dx.$$

Therefore, $E[Z^2] = 7/6$. \(\diamond\)

Example 3.4. Coin Tosses

This is a continuation of Example 3.2. Consider the sum $S_n = X_1 + \dots + X_n$. Its expectation is

$$\begin{aligned} E[S_n] &= \sum_{k=0}^n k P(S_n = k) \\ &= \frac{1}{2^n} \sum_{k=1}^n k \frac{n!}{k!(n-k)!} \end{aligned}$$

$$\begin{aligned}
&= \frac{n}{2^n} \sum_{k=1}^n \frac{(n-1)!}{(k-1)!((n-1)-(k-1))!} \\
&= \frac{n}{2^n} \sum_{j=0}^{n-1} \frac{(n-1)!}{j!(n-1-j)!} \\
&= \frac{n}{2^n} 2^{n-1}.
\end{aligned}$$

Therefore, $E[S_n] = n/2$. ◇

Expectation inherits the linearity property of the Stieltjes–Lebesgue integral:

$$E[\lambda_1 g_1(X) + \lambda_2 g_2(X)] = \lambda_1 E[g_1(X)] + \lambda_2 E[g_2(X)], \quad (3.13)$$

where $\lambda_1, \lambda_2 \in \mathbb{R}$, and g_1 and g_2 satisfy the integrability condition (3.11). Also, expectation is monotone, in the sense that $g_1(x) \leq g_2(x)$ for all x implies

$$E[g_1(X)] \leq E[g_2(X)]. \quad (3.14)$$

Consider now the *indicator function* of an event A :

$$1_A(\omega) = \begin{cases} 1 & \text{if } \omega \in A, \\ 0 & \text{if } \omega \notin A. \end{cases} \quad (3.15)$$

The random variable $X = 1_A$ takes the value 1 with probability $P(X = 1) = P(A)$ and the value 0 with probability $P(X = 0) = P(\bar{A}) = 1 - P(A)$. Therefore, $E[X] = 0 \times P(X = 0) + 1 \times P(X = 1) = P(X = 1) = P(A)$, that is to say,

$$E[1_A] = P(A). \quad (3.16)$$

In particular, $E[1] = 1$.

Sometimes, one needs to define $E[g(X)]$ for a complex function $g : \mathbb{R} \rightarrow \mathbb{C}$, that is, $g(x) = g_R(x) + i g_I(x)$ where g_R and g_I take real values. The definition of expectation is now

$$E[g(X)] = E[g_R(X)] + i E[g_I(X)], \quad (3.17)$$

provided that the expectations on the right-hand side are finite.

The triangle inequality

$$|E[g(X)]| \leq E[|g(X)|] \quad (3.18)$$

is useful, and its proof follows from the analogous inequality for the Stieltjes–Lebesgue integral:

$$\left| \int_{-\infty}^{+\infty} g(x) dF(x) \right| \leq \int_{-\infty}^{+\infty} |g(x)| dF(x).$$

Definition 3.2. *Mean and Variance*

The mean m and variance σ^2 of a real random variable X are defined by

$$m = E[X] = \int_{-\infty}^{+\infty} x dF(x),$$

$$\sigma^2 = E[(X - m)^2] = \int_{-\infty}^{+\infty} (x - m)^2 dF(x).$$

Of course, the integrals involved must be well defined. The variance is also denoted by $\text{Var}(X)$. From the linearity of expectation, it follows that $E[(X - m)^2] = E[X^2] - 2mE[X] + m^2$, that is,

$$\text{Var}(X) = E[X^2] - m^2. \quad (3.19)$$

Exercise 3.1. *Max*

Let X_1, X_2, \dots, X_n be independent random variables uniformly distributed on $[0, 1]$, that is to say, with the probability density $f(x) = 1_{[0,1]}(x)$. Compute the expectation of $Z = \max(X_1, \dots, X_n)$.

Solution. $P(\max(X_1, \dots, X_n) \leq z) = P(X_1 \leq z, \dots, X_n \leq z) = \prod_{k=1}^n P(X_k \leq z) = z^n$ for $z \in [0, 1]$. Therefore, Z has the probability density $\frac{d}{dz} z^n = n z^{n-1}$ for $z \in [0, 1]$, and 0 otherwise, and

$$E[Z] = \int_0^1 z(nz^{n-1}) dz = \frac{n}{n+1}.$$

The following is a very important formula to be used repeatedly in the sequel.

Theorem 3.1. *Telescope Formula*

For a random variable X taking its values in \mathbb{N} ,

$$E[X] = \sum_{n=1}^{\infty} P(X \geq n). \quad (3.20)$$

Proof.

$$\begin{aligned} E[X] &= 1 \cdot P(X = 1) + 2 \cdot P(X = 2) + 3 \cdot P(X = 3) + \dots \\ &= P(X = 1) + P(X = 2) + P(X = 3) + \dots \\ &\quad + P(X = 2) + P(X = 3) + \dots \\ &\quad + P(X = 3) + \dots \\ &= P(X \geq 1) + P(X \geq 2) + P(X \geq 3) + \dots \quad \square \end{aligned}$$

Theorem 3.2. *Wald's Lemma*

Let $\{X_n\}_{n \geq 1}$ be a sequence of integrable random variables such that $E[X_n] = E[X_1]$ for all $n \geq 1$. Let T be an integer-valued random variable such that for all $n \geq 1$, the event $\{T \geq n\}$ is independent of X_n . Then

$$E\left[\sum_{n=1}^T X_n\right] = E[X_1]E[T]. \quad (3.21)$$

Proof. Let $S = \sum_{n=1}^T X_n$. Then by dominated convergence (see Theorem 3.2 and Example 3.3 of the Appendix)

$$E[S] = E\left[\sum_{n=1}^{\infty} X_n 1_{\{n \leq T\}}\right] = \sum_{n=1}^{\infty} E[X_n 1_{\{n \leq T\}}].$$

But $E[X_n 1_{\{n \leq T\}}] = E[X_n]E[1_{\{n \leq T\}}] = E[X_1]P(n \leq T)$. The result then follows from the telescope formula. \square

3.3 Famous Random Variables

Consider a sequence $\{X_n\}_{n \geq 1}$ of random variables taking their values in $\{0, 1\}$, with the same probability distribution

$$P(X_n = 1) = p,$$

where $p \in (0, 1)$. Suppose, moreover, that the X_n are independent. Since $P(X_j = a_j) = p$ or $1 - p$ depending upon whether $a_j = 1$ or 0 , and since there are exactly $\sum_{j=1}^k a_j$ numbers among a_1, \dots, a_k that are equal to 1 ,

$$P(X_1 = a_1, \dots, X_k = a_k) = p^{\sum_{j=1}^k a_j} (1 - p)^{k - \sum_{j=1}^k a_j}. \quad (3.22)$$

Comparing with Examples 1.3 and 2.2, we see that we have modeled a game of heads and tails, with a biased coin if $p \neq \frac{1}{2}$. The framework of heads and tails gives rise to three famous discrete random variables: the binomial random variable, the geometric random variable, and the Poisson random variable.

Binomial. Define

$$S_n = X_1 + \dots + X_n.$$

This random variable takes the values $0, 1, \dots, n$. To obtain $S_n = i$ where $i \in [0, n]$, one must have $X_1 = a_1, \dots, X_n = a_n$ with $\sum_{j=1}^n a_j = i$. There are $\binom{n}{i}$ distinct ways of having this, and each occurs with probability $p^i (1 - p)^{n-i}$. Therefore, for $i \in [0, n]$,

$$P(S_n = i) = \binom{n}{i} p^i (1 - p)^{n-i}. \quad (3.23)$$

A random variable X with the above distribution is called a *binomial random variable of size n and parameter $p \in (0, 1)$* . A direct computation gives for the mean and the variance of the binomial random variable

$$E[X] = np, \quad \text{Var}(X) = np(1 - p). \quad (3.24)$$

(Do Problem 1.5.1 for an alternative to the direct computation via generating functions.)

Geometric. Define the random variable T to be the first time of occurrence of 1 in the sequence X_1, X_2, \dots , that is,

$$T = \inf \{n \geq 1; X_n = 1\},$$

with the convention that if $X_n = 0$ for all $n \geq 1$, then $T = \infty$ (actually, we shall show that this event has probability 0). The event $\{T = k\}$ is exactly $\{X_1 = 0, \dots, X_{k-1} = 0, X_k = 1\}$, and therefore, $P(T = k) = P(X_1 = 0) \cdots P(X_{k-1} = 0)P(X_k = 1)$, that is, for $k \geq 1$,

$$P(T = k) = (1 - p)^{k-1} p. \quad (3.25)$$

In particular, $P(T < \infty) = \sum_{k=1}^{\infty} P(T = k) = \sum_{k=1}^{\infty} p(1 - p)^{k-1} = 1$ since $p \in (0, 1)$.

A variable with the distribution (3.25) is called a *geometric random variable with parameter p* . A direct computation (or Problem 1.5.1) gives

$$E[T] = \frac{1}{p}. \quad (3.26)$$

Poisson. Suppose you play heads and tails for a large number of turns N with a coin such that

$$P(X_n = 1) = \frac{\alpha}{N}.$$

The number S_N of tails you observe is therefore distributed according to

$$P(S_N = k) = \binom{N}{k} \left(\frac{\alpha}{N}\right)^k \left(1 - \frac{\alpha}{N}\right)^{N-k}.$$

From (3.24), the average number of tails is $N \cdot \frac{\alpha}{N} = \alpha$, a constant. It is interesting to see the limit of the distribution of S_N as $N \rightarrow \infty$. Denoting by $p_N(k)$ the probability of $S_N = k$, we have that

$$p_N(0) = \left(1 - \frac{\alpha}{N}\right)^N$$

goes to $e^{-\alpha}$ as N goes to ∞ . Also,

$$\frac{p_N(k+1)}{p_N(k)} = \frac{\frac{N-k}{k+1} \frac{\alpha}{N}}{1 - \frac{\alpha}{N}}$$

goes to $\frac{\alpha}{k+1}$. Therefore, for all $k \geq 0$,

$$\lim_{n \uparrow \infty} p_N(k) = e^{-\alpha} \frac{\alpha^k}{k!},$$

with the convention $0! = 1$. By definition, a *Poisson random variable with parameter $\theta > 0$* is a random variable such that for all $k \geq 0$,

$$P(X = k) = e^{-\theta} \frac{\theta^k}{k!}. \quad (3.27)$$

Therefore, the probability distribution of S_N tends as $N \rightarrow \infty$ to the probability distribution of a Poisson r.v. of parameter α . This phenomenon is called the *Poisson law of rare events*, because as $N \rightarrow \infty$, obtaining tails is an increasingly rare event, of probability $\frac{\alpha}{N}$.

A direct computation (or Problem 1.5.1) gives for the Poisson random variable

$$E[X] = \theta, \quad \text{Var}(X) = \theta. \quad (3.28)$$

Exercise 3.2. Sum of Independent Poisson Variables

Let X_1 and X_2 be two independent Poisson random variables with means $\theta_1 > 0$ and $\theta_2 > 0$, respectively. Show that $X = X_1 + X_2$ is a Poisson random variable with mean $\theta = \theta_1 + \theta_2$.

Solution. For $k \geq 0$,

$$\begin{aligned} P(X = k) &= P(X_1 + X_2 = k) = P\left(\sum_{i=0}^k \{X_1 = i, X_2 = k - i\}\right) \\ &= \sum_{i=0}^k P(X_1 = i, X_2 = k - i) = \sum_{i=0}^k P(X_1 = i)P(X_2 = k - i) \\ &= \sum_{i=0}^k e^{-\theta_1} \frac{\theta_1^i}{i!} e^{-\theta_2} \frac{\theta_2^{k-i}}{(k-i)!} = \frac{e^{-(\theta_1 + \theta_2)}}{k!} (\theta_1 + \theta_2)^k. \end{aligned}$$

Exercise 3.3. Poisson Sums of i.i.d. Variables

Let $\{X_n\}_{n \geq 1}$ be independent random variables taking the values 0 and 1 with probability $q = 1 - p$ and p , respectively, where $p \in (0, 1)$. Let T be a Poisson random variable with mean $\theta > 0$, independent of $\{X_n\}_{n \geq 1}$. Define

$$S = X_1 + \cdots + X_T$$

(that is, $S(\omega) = X_1(\omega) + \cdots + X_n(\omega)$ if $T(\omega) = n$). Show that S is a Poisson random variable with mean $p\theta$.

Solution.

$$\begin{aligned} P(S = k) &= P(X_1 + \cdots + X_T = k) = P\left(\sum_{n=k}^{\infty} \{X_1 + \cdots + X_n = k, T = n\}\right) \\ &= \sum_{n=k}^{\infty} P(X_1 + \cdots + X_n = k, T = n) = \sum_{n=k}^{\infty} P(X_1 + \cdots + X_n = k)P(T = n) \\ &= \sum_{n=k}^{\infty} \frac{n!}{k!(n-k)!} p^k q^{n-k} e^{-\theta} \frac{\theta^n}{n!} = e^{-\theta} \frac{(p\theta)^k}{k!} \sum_{n=k}^{\infty} \frac{(q\theta)^{n-k}}{(n-k)!} \\ &= e^{-\theta} \frac{(p\theta)^k}{k!} \sum_{i=0}^{\infty} \frac{(q\theta)^i}{i!} = e^{-\theta} \frac{(p\theta)^k}{k!} e^{q\theta} = e^{-p\theta} \frac{(p\theta)^k}{k!}. \end{aligned}$$

Thus, if one “thins out” with thinning probability $1 - p$ a population sample of Poissonian size, the remaining sample has also a Poissonian size, with the obvious mean that is p times that of the original sample. \square

Exercise 3.4. Bernoulli–Poisson Eggs

Imagine a bird laying T eggs, each egg blue or pink, with probability p of laying a blue egg. The above exercise showed that if the number of eggs is Poisson with mean θ , then the

number of blue eggs is Poisson with mean θp and the number of pink eggs is Poisson with mean θq . Show that the number of blue eggs and the number of pink eggs are independent random variables.

Solution. If S is the number of blue eggs, $T - S$ is the number of pink eggs. One must show that for any integers $k \geq 0, \ell \geq 0$,

$$\begin{aligned} P(S = k, T - S = \ell) &= P(S = k)P(T - S = \ell) \\ &= e^{-\theta p} \frac{(\theta p)^k}{k!} e^{-\theta q} \frac{(\theta q)^\ell}{\ell!}. \end{aligned}$$

But

$$\begin{aligned} P(S = k, T - S = \ell) &= P(S = k, T = k + \ell) \\ &= P(X_1 + \cdots + X_{k+\ell} = k, T = k + \ell) \\ &= P(X_1 + \cdots + X_{k+\ell} = k)P(T = k + \ell) \\ &= \frac{(k + \ell)!}{k! \ell!} p^k q^\ell e^{-\theta} \frac{\theta^{k+\ell}}{(k + \ell)!} = e^{-p\theta} \frac{(p\theta)^k}{k!} e^{-q\theta} \frac{(q\theta)^\ell}{\ell!}. \quad \square \end{aligned}$$

Our small list of random variables continues with absolutely continuous ones. An *absolutely continuous* random variable is by definition a real r.v. with a probability density, i.e.,

$$P(X \leq x) = \int_{-\infty}^x f(x) dx, \quad (3.29)$$

where $f(x) \geq 0$, and since X is real,

$$\int_{-\infty}^{+\infty} f(x) dx = 1. \quad (3.30)$$

Uniform. A random variable X with probability density function

$$f(x) = \frac{1}{b - a} \quad (3.31)$$

for $x \in [a, b]$ and 0 otherwise is called a *uniform random variable* on $[a, b]$. Its mean and variance are

$$E[X] = \frac{a + b}{2}, \quad \text{Var}(X) = \frac{(b - a)^2}{12}, \quad (3.32)$$

as simple computations reveal.

Gaussian. The random variable X with p.d.f.

$$f(x) = \frac{1}{\sigma \sqrt{2\pi}} e^{-\frac{1}{2} \frac{(x - \mu)^2}{\sigma^2}}, \quad (3.33)$$




Figure 1.4.1. Exponential random variable

where $m \in \mathbb{R}$ and $\sigma > 0$, is called a *Gaussian random variable*. One can check that $E[X] = m$ and $\text{Var}(X) = \sigma^2$.

Exponential. The random variable X with c.d.f.

$$F(x) = (1 - e^{-\lambda x})1_{\{x \geq 0\}}$$

and p.d.f.

$$f(x) = \lambda e^{-\lambda x} 1_{\{x \geq 0\}} \quad (3.34)$$

(see Fig. 1.4.1) is called an *exponential random variable with parameter λ* . Its mean is

$$E[X] = 1/\lambda. \quad (3.35)$$

4 Random Vectors

4.1 Absolutely Continuous Random Vectors

A random vector of dimension n is a collection of n real random variables

$$X = (X_1, \dots, X_n). \quad (4.1)$$

Each of the random variables X_1, \dots, X_n can be characterized from a probabilistic point of view by its cumulative distribution function.

However, the c.d.f. of each coordinate of a random vector does not completely describe the probabilistic behavior of the whole vector. For instance, if U_1 and U_2 are two independent random variables with the same c.d.f. $G(x)$, the vectors $X = (X_1, X_2)$ defined respectively by $X_1 = U_1, X_2 = U_2$ and $X_1 = U_1, X_2 = U_1$ have each of their coordinates with the same c.d.f., and they are quite different.

The c.d.f. of the vector $X = (X_1, \dots, X_n)$ is the function $F : \mathbb{R}^n \rightarrow [0, 1]$ defined by

$$F(x_1, \dots, x_n) = P(X_1 \leq x_1, \dots, X_n \leq x_n), \quad (4.2)$$

and it contains all the probabilistic information about the random vector X .

An absolutely continuous random vector is one that admits a *probability density function*, that is,

$$F(x_1, \dots, x_n) = \int_{-\infty}^{x_1} \cdots \int_{-\infty}^{x_n} f(y_1, \dots, y_n) dy_1 \cdots dy_n \quad (4.3)$$

for some nonnegative function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ such that

$$\int_{-\infty}^{+\infty} \cdots \int_{-\infty}^{+\infty} f(y_1, \dots, y_n) dy_1 \cdots dy_n = 1. \quad (4.4)$$

For a function $g : \mathbb{R}^n \rightarrow \mathbb{R}$, the *expectation* of $g(X)$ is by definition

$$E[g(X)] = \int_{-\infty}^{+\infty} \cdots \int_{-\infty}^{+\infty} g(x_1, \dots, x_n) f(x_1, \dots, x_n) dx_1 \cdots dx_n, \quad (4.5)$$

where it is required that $E[|g(X)|] < \infty$ in order to make the integral in (4.5) meaningful.

Expectation so defined enjoys, *mutatis mutandis*, the properties mentioned for the scalar case: linearity (see (3.13)), monotonicity (see (3.14)), and the triangle inequality (see (3.18)).

Consider a 2-dimensional vector $X = (X_1, X_2)$ with p.d.f. $f_{X_1, X_2}(x_1, x_2)$. The p.d.f. of X_1 is obtained by integrating out x_2 :

$$f_{X_1}(x_1) = \int_{-\infty}^{+\infty} f_{X_1, X_2}(x_1, x_2) dx_2. \quad (4.6)$$

Indeed,

$$\begin{aligned} P(X_1 \leq a) &= P((X_1, X_2) \in (-\infty, a] \times \mathbb{R}) \\ &= \int_{-\infty}^a \int_{-\infty}^{+\infty} f_{X_1, X_2}(x_1, x_2) dx_1 dx_2 \\ &= \int_{-\infty}^a \left(\int_{-\infty}^{+\infty} f_{X_1, X_2}(x_1, x_2) dx_2 \right) dx_1. \end{aligned}$$

Formula (4.6) extends in an obvious manner to vectors.

If X_1, \dots, X_n are absolutely continuous random variables of respective probability density functions f_1, \dots, f_n , and if, moreover, X_1, \dots, X_n are independent, then

$$\begin{aligned} P(X_1 \leq x_1, \dots, X_n \leq x_n) &= P(X_1 \leq x_1) \cdots P(X_n \leq x_n) \\ &= \left(\int_{-\infty}^{x_1} f_1(y_1) dy_1 \right) \cdots \left(\int_{-\infty}^{x_n} f_n(y_n) dy_n \right) \\ &= \int_{-\infty}^{x_1} \cdots \int_{-\infty}^{x_n} f_1(y_1) \cdots f_n(y_n) dy_1 \cdots dy_n. \end{aligned}$$

The p.d.f. of the vector (X_1, \dots, X_n) is therefore the product of the p.d.f. of its coordinates

$$f(x_1, \dots, x_n) = f_1(x_1) \cdots f_n(x_n). \quad (4.7)$$

Conversely, if the vector X has a p.d.f. that factors as in (4.7), where f_1, \dots, f_n are p.d.f., then X_1, \dots, X_n are independent r.v.s with respective p.d.f. f_1, \dots, f_n . Indeed,

$$P(X_1 \in A_1, \dots, X_n \in A_n) = \int_{A_1} \cdots \int_{A_n} f_1(y_1) \cdots f_n(y_n) dy_1 \cdots dy_n;$$

that is, by Fubini's theorem,

$$P(X_1 \in A_1, \dots, X_n \in A_n) = \left(\int_{A_1} f_1(y_1) dy_1 \right) \times \cdots \times \left(\int_{A_n} f_n(y_n) dy_n \right).$$

Letting $A_2 = \cdots = A_n = \mathbb{R}$ in the last identity yields

$$P(X_1 \in A_1) = \int_{A_1} f_1(y_1) dy_1,$$

which proves that X_1 has the p.d.f. f_1 . Similarly, $P(X_i \in A_i) = \int_{A_i} f_i(y_i) dy_i$, and therefore

$$P(X_1 \in A_1, \dots, X_n \in A_n) = P(X_1 \in A_1) \cdots P(X_n \in A_n),$$

which proves independence, since the A_i 's are arbitrary.

Exercise 4.1.

Prove that $P(X = Y) = 0$ for any two independent absolutely continuous random variables X and Y .

Solution.

$$P(X = Y) = E[1_{\{X=Y\}}] = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} g(x, y) dx dy,$$

where $g(x, y) = 1_{\{x=y\}} f_X(x) f_Y(y)$ is null outside the diagonal. Since the diagonal has null Lebesgue measure (area), the integral is zero. \square

Exercise 4.2. Freezing a Random Variable

Let X_1, \dots, X_n be independent random variables with respective p.d.f. f_1, \dots, f_n . Show that

$$E[g(X_1, \dots, X_n)] = \int_{-\infty}^{+\infty} E[g(y, X_2, \dots, X_n)] f_1(y) dy$$

and that

$$P(X_1 \leq X_2, \dots, X_1 \leq X_n, X_1 \leq x) = \int_{-\infty}^x P(X_2 \geq y) \cdots P(X_n \geq y) f_1(y) dy.$$

Solution. We do the case $n = 2$ for simplicity. We have

$$\begin{aligned} E[g(X_1, X_2)] &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} g(x_1, x_2) f_1(x_1) f_2(x_2) dx_1 dx_2 \\ &= \int_{-\infty}^{+\infty} f_1(x_1) \left\{ \int_{-\infty}^{+\infty} g(x_1, x_2) f_2(x_2) dx_2 \right\} dx_1 \\ &= \int_{-\infty}^{+\infty} f_1(x) E[g(x, X_2)] dx. \end{aligned}$$

The second equality in the statement of the present exercise is obtained from the first one by letting (case $n = 3$ this time) $g(X_1, X_2, X_3) = 1_{\{X_1 \leq X_2\}} 1_{\{X_1 \leq X_3\}} 1_{\{X_1 \leq X_3\}}$ and observing that

$$\begin{aligned} E[g(y, X_2, X_3)] &= E[1_{\{y \leq X\}} 1_{\{X_2 \geq y\}} 1_{\{X_3 \geq y\}}] \\ &= 1_{\{y \leq x\}} P(X_2 \geq y, X_3 \geq y) = 1_{\{y \leq x\}} P(X_2 \geq y) P(X_3 \geq y). \quad \square \end{aligned}$$

Exercise 4.3. Sum of Independent Absolutely Continuous Variables

Show that the probability density function of the random variable $Z = X + Y$, where X and Y are independent random variables with respective probability densities f_X and f_Y is given by the *convolution formula*

$$f_Z(z) = \int_{-\infty}^{+\infty} f_Y(z - y) f_X(y) dy. \quad (4.8)$$

Solution. The p.d.f. of vector (X, Y) is $f_X(x) f_Y(y)$, and therefore, for all $a \in \mathbb{R}$,

$$\begin{aligned} P(Z \leq a) &= P(X + Y \leq a) = E[1_{\{X+Y \leq a\}}] \\ &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} 1_{\{x+y \leq a\}} f_X(x) f_Y(y) dx dy. \end{aligned}$$

The latter integral can be written, by Fubini's theorem,

$$\int_{-\infty}^{+\infty} \left\{ \int_{-\infty}^{+\infty} 1_{\{y \leq a-x\}} f_Y(y) dy \right\} f_X(x) dx = \int_{-\infty}^{+\infty} \left\{ \int_{-\infty}^{a-x} f_Y(y) dy \right\} f_X(x) dx;$$

that is, after a change of variable,

$$P(Z \leq a) = \int_{-\infty}^a \left\{ \int_{-\infty}^{+\infty} f_Y(z-x) f_X(x) dx \right\} dz. \quad \square$$

4.2 Discrete Random Vectors

In what follows, only the case of a discrete random vector $X = (X_1, \dots, X_n)$ where all the random variables X_i take their values in the *same* denumerable space E will be considered. This restriction is not essential, but it simplifies the notation.

The statistical behavior of X is described by its distribution ($p(u); u \in E^n$) where

$$p(u) = P(X_1 = u_1, \dots, X_n = u_n),$$

and the expectation of $g(X)$ is defined by

$$E[g(X)] = \sum_{u \in E^n} g(u)p(u) \quad (4.9)$$

as long as g is nonnegative or the sum in (4.9) is absolutely convergent.

Here again, as for the scalar case, the expectation so defined has the linearity and monotonicity properties (see (3.13) and (3.14)), and the triangle inequality is true (see (3.18)).

Example 4.1. *The Multinomial Distribution*

Consider k boxes B_1, \dots, B_K and n balls, to be placed in the boxes independently of one another, with the probability p_i for a given ball to be assigned to box B_i . Of course,

$$\sum_{i=1}^K p_i = 1. \quad (4.10)$$

After placing all the balls in the boxes, there are X_i balls in box B_i , where

$$\sum_{i=1}^K X_i = n. \quad (4.11)$$

The probability distribution of the random vector $X = (X_1, \dots, X_n)$ is the *multinomial distribution* of size (n, K) and parameters p_1, \dots, p_K

$$P(X_1 = m_1, \dots, X_K = m_K) = \frac{n!}{\prod_{i=1}^K (m_i)!} \prod_{i=1}^K p_i^{m_i}, \quad (4.12)$$

where $m_1 + \dots + m_K = n$.

To prove this, observe that there are $n! / \prod_{i=1}^K (m_i)!$ distinct ways of placing n balls in K boxes in such a manner that m_1 balls are in box B_1 , m_2 are in B_2 , etc., and that each of these distinct ways occurs with the same probability, $\prod_{i=1}^K p_i^{m_i}$. \diamond

Exercise 4.4. *First Coordinate of a Multinomial Vector*

This is the continuation of Example 4.1. Show that X_1 is a binomial random variable of size n and parameter p_1 .

Solution. Put together the boxes B_2, \dots, B_K so that they form a single box. The whole process is now one of placing n balls, independently of one another. A given ball is put in the first box B_1 with probability p_1 , in box $B_2 + \dots + B_K$ with probability $p_2 + \dots + p_K = 1 - p_1$. But when $K = 2$, and therefore $m_2 = n - m_1$, $p_2 = 1 - p_1$, (4.12) reduces to

$$P(X_1 = m_1, X_2 = m_2) = P(X_1 = m_1) = \frac{n!}{m_1!(n - m_1)!} p_1^{m_1} (1 - p_1)^{n - m_1}. \quad \square$$

4.3 Product Formula for Expectation

Let Y and Z be independent random vectors with the respective probability densities f_Y and f_Z . Arguments similar to those leading to formula (4.7) give for the vector $X = (Y, Z)$ the probability density

$$f_{Y,Z}(y, z) = f_Y(y)f_Z(z). \quad (4.13)$$

In particular, if $g_1 : \mathbb{R}^p \rightarrow \mathbb{R}$ and $g_2 : \mathbb{R}^q \rightarrow \mathbb{R}$ are such that

$$E[|g_1(Y)|] < \infty, E[|g_2(Z)|] < \infty, \quad (4.14)$$

then

$$E[g_1(Y)g_2(Z)] = E[g_1(Y)]E[g_2(Z)], \quad (4.15)$$

as follows from Fubini's theorem.

The analogous result holds when X and Z are discrete random vectors. It turns out that (4.15) is true for *any* random vectors X and Z , as long as the integrability conditions (4.14) are satisfied. Formula (4.15) is also true when the integrability conditions are replaced by the hypothesis that g_1 and g_2 are nonnegative.

5 Transforms of Probability Distributions

5.1 Generating Functions

Let $\bar{D}(0; 1)$ be the complex closed unit disk centered at 0.

Definition 5.1. *Generating Function*

Let X be an \mathbb{N} -valued random variable. Its generating function (g.f.) is the function $g_X : \bar{D}(0; 1) \rightarrow \mathbb{C}$ defined by

$$g_X(z) = E[z^X] = \sum_{k=0}^{\infty} P(X = k)z^k. \quad (5.1)$$

Example 5.1. *Binomial Random Variable.*

For the binomial r.v. of size n and parameter p ,

$$\sum_{k=0}^n P(X = k)z^k = \sum_{k=0}^n \binom{n}{k} (zp)^k (1-p)^{n-k},$$

and therefore

$$g_X(z) = (1 - p + pz)^n. \quad (5.2)$$

◇

Example 5.2. *Poisson Random Variable*

For the Poisson r.v. of mean θ ,

$$\sum_{k=0}^{\infty} P(X = k)z^k = e^{-\theta} \sum_{k=0}^{\infty} \frac{(\theta z)^k}{k!},$$

and therefore

$$g_X(z) = e^{\theta(z-1)}. \quad (5.3)$$

◇

Exercise 5.1. *Random Sums of i.i.d Random Variables*

Let $\{Y_n\}_{n \geq 1}$ be an i.i.d sequence of integer-valued random variables with the common generating function g_Y . Let T be another integer-valued random variable, independent of the sequence $\{Y_n\}_{n \geq 1}$, and let g_T be its generating function. Compute the generating function of

$$X = \sum_{n=1}^T Y_n,$$

where by convention $\sum_{n=1}^0 = 0$.

Solution.

$$z^X = z^{\sum_{n=1}^T Y_n} = \sum_{k=0}^{\infty} \left\{ \left(z^{\sum_{n=1}^T Y_n} \right) 1_{\{T=k\}} \right\} = \sum_{k=0}^{\infty} \left(z^{\sum_{n=1}^k Y_n} \right) 1_{\{T=k\}}.$$

Therefore,

$$E[z^X] = \sum_{k=0}^{\infty} E \left[1_{\{T=k\}} \left(z^{\sum_{n=1}^k Y_n} \right) \right] = \sum_{k=0}^{\infty} E[1_{\{T=k\}}] E[z^{\sum_{n=1}^k Y_n}],$$

where we have used independence of T and $\{Y_n\}_{n \geq 1}$. Now, $E[1_{\{T=k\}}] = P(T = k)$, and $E[z^{\sum_{n=1}^k Y_n}] = g_Y(z)^k$, and therefore $E[z^X] = \sum_{k=0}^{\infty} P(T = k)g_Y(z)^k$. That is, finally,

$$g_X(z) = g_T(g_Y(z)). \quad (5.4)$$

□

The power series associated with the sequence $\{P(X = n)\}_{n \geq 0}$ has a radius of convergence $R \geq 1$, since $\sum_{n=0}^{\infty} P(X = n) = 1 < \infty$. Therefore, the domain of definition of g_X contains the open unit disk. Inside this open disk, differentiation term by term is possible, for instance,

$$g'_X(z) = \sum_{n=1}^{\infty} nP(X = n)z^{n-1}, \quad (5.5)$$

$$g_X''(z) = \sum_{n=2}^{\infty} n(n-1)P(X=n)z^{n-2}. \quad (5.6)$$

The right-hand side of (5.5) is well-defined at $z = 1$, being equal to $\sum_{n=1}^{\infty} nP(X=n)$, a nonnegative quantity, possibly infinite. The left-hand side of (5.5) is, however, not always defined by formula (5.1) for $z = 1$. So we *define* $g'(1) = \sum_{n=1}^{\infty} nP(X=n)$, that is,

$$g'(1) = E[X]. \quad (5.7)$$

By *Abel's theorem* (see Theorem 1.2 of the Appendix), the limit as the *real* variable x increases to 1 of $\sum_{n=1}^{\infty} nP(X=n)x^{n-1}$ is $\sum_{n=1}^{\infty} nP(X=n)$, and therefore g_X , as a function on the real interval $[0, 1)$, can be extended to $[0, 1]$ by (5.7), and this extension preserves continuity. This is used as follows: Suppose $E[X]$ is not known but that you have an expression of $g_X'(x)$ in $[0, 1)$ for which you can compute $\lim_{x \rightarrow 1} g_X'(x)$ as x increases through real values to 1. Then you know that this limit equals $E[X]$.

Similarly, starting from (5.6), and defining

$$g_X''(1) = E[X(X-1)], \quad (5.8)$$

we make $g_X''(x)$ a continuous function on $[0, 1]$.

Example 5.3.

This is a continuation of Exercise 5.1. The generating function of $X = \sum_{n=1}^T Y_n$ was found to be $g_X(z) = g_T(g_Y(z))$. Therefore $g_X'(x) = g_T'(g_Y(x))g_Y'(x)$ for all $x \in [0, 1)$. Letting $x \rightarrow 1$, $g_X'(1) = g_T'(g_Y(1))g_Y'(1) = g_T'(1)g_Y'(1)$, and therefore $E[X] = E[T]E[Y]$. \diamond

The generating function characterizes the distribution of a random variable. This means the following. Suppose that, without knowing the distribution of X , you have been able to compute its generating function $g(z)$, and that, moreover, you are able to give its power series expansion in a neighborhood of the origin:

$$g(z) = \sum_{n=0}^{\infty} a_n z^n.$$

Since $g(z)$ is the generating function of Z ,

$$g(z) = \sum_{n=0}^{\infty} P(X=n)z^n$$

and since the power series expansion around the origin is unique, the distribution of X is identified as

$$P(X=n) = a_n$$

for all $n \geq 0$. Similarly, if two \mathbb{N} -valued random variables have the same g.f., they have the same distribution.

Exercise 5.2. *The Lottery*

Let $X_1, X_2, X_3, X_4, X_5,$ and X_6 be independent random variables uniformly distributed over $\{0, 1, \dots, 9\}$. Compute the generating function of $Y = 27 + X_1 + X_2 + X_3 - X_4 - X_5 - X_6$. Use this to compute the probability that in a 6-digit lottery the sum of the first three digits equals the sum of the last three digits.

Solution.

$$\begin{aligned} E[z^{X_i}] &= \frac{1}{10}(1 + z + \dots + z^9) = \frac{1}{10} \frac{1 - z^{10}}{1 - z}, \\ E[z^{-X_i}] &= \frac{1}{10} \left(1 + \frac{1}{z} + \dots + \frac{1}{z^9}\right) = \frac{1}{10} \frac{1 - z^{-10}}{1 - z^{-1}} = \frac{1}{10} \frac{1}{z^9} \frac{1 - z^{10}}{1 - z}, \\ E[z^Y] &= E\left[z^{27 + \sum_{i=1}^3 X_i - \sum_{i=4}^6 X_i}\right] = E\left[z^{27} \prod_{i=1}^3 z^{X_i} \prod_{i=4}^6 z^{-X_i}\right] \\ &= z^{27} \prod_{i=1}^3 E[z^{X_i}] \prod_{i=4}^6 E[z^{-X_i}]. \end{aligned}$$

Therefore,

$$g_Y(z) = \frac{1}{10^6} \frac{(1 - z^{10})^6}{(1 - z)^6}.$$

But $P(X_1 + X_2 + X_3 = X_4 + X_5 + X_6) = P(Y = 27)$ is the factor of z^{27} in the power series expansion of $g_Y(z)$. Since

$$(1 - z^{10})^6 = 1 - \binom{6}{1} z^{10} + \binom{6}{2} z^{20} + \dots$$

and

$$(1 - z)^{-6} = 1 + \binom{6}{5} z + \binom{7}{5} z^2 + \binom{8}{5} z^3 + \dots,$$

we find that

$$P(Y = 27) = \frac{1}{10^6} \left(\binom{32}{5} - \binom{6}{1} \binom{22}{5} + \binom{6}{2} \binom{12}{5} \right).$$

Theorem 5.1. *Graph of a Generating Function*

(α) Let $g : [0, 1] \rightarrow \mathbb{R}$ be defined by $g(x) = E[x^X]$, where X is a nonnegative integer-valued random variable. Then g is nondecreasing and convex. Moreover, if $P(X = 0) < 1$, then g is strictly increasing, and if $P(X \leq 1) < 1$, it is strictly convex \cup .

(β) Suppose $P(X \leq 1) < 1$. If $E[X] \leq 1$, the equation $x = g(x)$ has a unique solution $x \in [0, 1]$, namely $x = 1$. If $E[X] > 1$, it has two solutions in $[0, 1]$, $x = 1$ and $x = x_0 \in (0, 1)$.

Proof. Just observe that for $x \in [0, 1]$,

$$g'(x) = \sum_{n=1}^{\infty} nP(X = n)x^{n-1} \geq 0,$$

and therefore g is nondecreasing, and

$$g''(x) = \sum_{n=2}^{\infty} n(n-1)P(X = n)x^{n-2} \geq 0,$$

and therefore g is convex. For $g'(x)$ to be null for some $x \in (0, 1)$, it is necessary to have $P(X = n) = 0$ for all $n \geq 1$, and therefore $P(X = 0) = 1$. For $g''(x)$ to be null for some $x \in (0, 1)$, one must have $P(X = n) = 0$ for all $n \geq 2$, and therefore $P(X = 0) + P(X = 1) = 1$.




Figure 1.5.1. Two aspects of the generating function

The graph of $g : [0, 1] \rightarrow \mathbb{R}$ has, in the strictly increasing strictly convex case $P(X \leq 1)$, the general shape shown in Figure 1.5.1, where we distinguish two cases: $E[X] = g'(1) \leq 1$, and $E[X] = g'(1) > 1$. The rest of the proof is then easy. \square

5.2 Characteristic Functions

Definition 5.2. Characteristic Function

The characteristic function (c.f.) $\psi_X : \mathbb{R}^n \rightarrow \mathbb{C}$ of a real random vector $X = (X_1, \dots, X_n)$ is defined by

$$\psi_X(u) = E[e^{iu^T X}]. \tag{5.9}$$

Example 5.4. Gaussian and Exponential

One can check that the following formulas give the characteristic functions of the two main continuous random variables:

(i) Gaussian

$$f(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2} \frac{(x-\mu)^2}{\sigma^2}} \leftrightarrow \psi(u) = e^{i\mu u - \frac{1}{2} \sigma^2 u^2}.$$

(ii) Exponential

$$f(x) = \lambda e^{-\lambda x} 1_{x>0} \leftrightarrow \psi(u) = \frac{\lambda}{\lambda - iu}. \quad \diamond$$

The characteristic function determines the distribution of a random vector, just as generating functions determine the distribution of integer-valued random variables. This will be admitted without proof. However, for continuous random vectors it is an easy consequence of the classical Fourier theory. Indeed, suppose that f_X and f_Y are the probability distribution functions of two random vectors X and Y on \mathbb{R}^n and suppose that their characteristic functions are the same. Then $f_X - f_Y$ admits the Fourier transform $\psi_X - \psi_Y \equiv 0$, and therefore, by the Fourier inversion theorem, $\int_C f_X(x) dx = \int_C f_Y(y) dy$ for all intervals $C \subset \mathbb{R}^n$.

The Fourier inversion formula gives, when $\psi_X(u)$ is integrable,

$$f_X(x) = \frac{1}{(2\pi)^n} \int_{\mathbb{R}^n} e^{-iu^T x} \psi_X(u) du$$

for the p.d.f. of X .

Exercise 5.3. Cauchy

Show that the characteristic function of the Cauchy random variable, with the p.d.f.

$$f(x) = \frac{1}{\pi} \frac{1}{1+x^2}$$

is $\psi_X(u) = e^{-|u|}$.

Solution. The Fourier transform of $u \rightarrow e^{-|u|}$ is $x \rightarrow \int_{-\infty}^{+\infty} e^{-|u|} e^{-iux} du = \frac{2}{1+x^2} = g(x)$. Since $g(x)$ is integrable, the Fourier inversion formula applies, and we get

$$e^{-|u|} = \frac{1}{2\pi} \int_{-\infty}^{+\infty} g(x) e^{iux} dx = \int_{-\infty}^{+\infty} \frac{1}{\pi} \frac{1}{1+x^2} e^{iux} dx.$$

Theorem 5.2. Characteristic Function Independence Criterion

Suppose that Y and Z are two random vectors of respective dimensions p and q , and that for all $v \in \mathbb{R}^p$, $w \in \mathbb{R}^q$, it holds that

$$E[e^{i(v^T Y + w^T Z)}] = \psi_1(v) \psi_2(w), \quad (5.10)$$

where $\psi_1(v)$ and $\psi_2(w)$ are the characteristic functions of some random vectors \tilde{Y} and \tilde{Z} of appropriate dimensions. Then Y and Z are independent, Y has the same distribution as \tilde{Y} , and Z has the same distribution as \tilde{Z} .

Proof. Define $X = (Y, Z)$ and $u = (v, w)$, so that (5.10) reads

$$E[e^{iu^T X}] = \psi(u) = \psi_1(v) \psi_2(w).$$

If one can find a vector $\hat{X} = (\hat{Y}, \hat{Z})$ such that $\psi_{\hat{X}}(u) = \psi(u)$, then, since the c.f. characterizes the c.d.f., $X = (Y, Z)$ has the same c.d.f. as $\hat{X} = (\hat{Y}, \hat{Z})$. Take \hat{Y} distributed as \tilde{Y} , \hat{Z} distributed as \tilde{Z} , and \hat{Y} and \hat{Z} independent. Then, using the product formula,

$$E[e^{iu^T \hat{X}}] = E[e^{iv^T \hat{Y}} e^{iw^T \hat{Z}}] = E[e^{iv^T \hat{Y}}] E[e^{iw^T \hat{Z}}] = \psi_1(v) \psi_2(w).$$

Therefore, (Y, Z) has the same distribution as (\hat{Y}, \hat{Z}) and in particular, Y and Z are independent. \square

Theorem 5.3. *Independence of an Event and of a Random Variable*

Let A be an event and X a random variable such that for all $u \in \mathbb{R}$,

$$E[1_A e^{iuX}] = P(A) E[e^{iuX}].$$

Then A and X are independent; that is, 1_A and X are independent.

Proof. We must check that for all $u, v \in \mathbb{R}$,

$$E[e^{iuX} e^{ivY}] = E[e^{iuX}] E[e^{ivY}].$$

But

$$e^{ivY} = 1 - 1_A + 1_A e^{iv} = 1 + 1_A(e^{iv} - 1),$$

and therefore

$$E[e^{iuX} e^{ivY}] = E[e^{iuX}] + (e^{iv} - 1) E[1_A e^{iuX}]$$

and

$$E[e^{iuX}] E[e^{ivY}] = E[e^{iuX}] + (e^{iv} - 1) P(A) E[e^{iuX}].$$

The hypothesis concludes the proof. \square

Let A be some event of positive probability, and let P_A denote the probability P conditioned by A , that is,

$$P_A(\cdot) = P(\cdot | A).$$

Definition 5.3.

The random variables X and Y are said to be conditionally independent given A if they are independent with respect to probability P_A .

By Theorem 5.2, a necessary and sufficient condition for this is that for all $u, v \in \mathbb{R}$,

$$E_A[e^{iuX} e^{ivY}] = E_A[e^{iuX}] E_A[e^{ivY}].$$

Observe that for an integrable or nonnegative random variable Z ,

$$P(A) E_A[Z] = E[Z 1_A].$$

The following result is then immediate. It is recorded for future reference.

Theorem 5.4.

Let A be an event of positive probability. The random variables X and Y are conditionally independent given A if and only if for all $u, v \in \mathbb{R}$,

$$P(A) E[e^{iuX} e^{ivY} 1_A] = E[e^{iuX} 1_A] E[e^{ivY} 1_A].$$

6 Transformations of Random Vectors

6.1 Smooth Change of Variables

Let $X = (X_1, \dots, X_n)$ be a random vector with the p.d.f. f_X , and define the random vector

$$Y = g(X) \quad (6.1)$$

where $g : \mathbb{R}^n \rightarrow \mathbb{R}^n$. More explicitly,

$$\begin{cases} Y_1 = g_1(X_1, \dots, X_n), \\ \vdots \\ Y_n = g_n(X_1, \dots, X_n). \end{cases} \quad (6.2)$$

Under smoothness assumptions on g , the random vector Y is absolutely continuous, and its p.d.f. can be explicitly computed from g and the p.d.f. and f_X . The precise result is recalled below.

The function g is from an open set $U \subset \mathbb{R}^n$ into \mathbb{R}^n ,

$$g : U \rightarrow \mathbb{R}^n, \quad (6.3)$$

and the coordinate functions g_i are continuously differentiable on U . Denote the Jacobian matrix of the function g by

$$J_g(x_1, \dots, x_n) = \left\{ \frac{\partial g_i}{\partial x_j}(x_1, \dots, x_n) \right\}_{1 \leq i, j \leq n} \quad (6.4)$$

and assume that on U ,

$$|\det J_g(x_1, \dots, x_n)| > 0. \quad (6.5)$$

Then $V = g(U)$ is an open subset of \mathbb{R}^n , and there exists an inverse $g^{-1} : V \rightarrow \mathbb{R}^n$ of g with the same properties as g . In particular, on V ,

$$|\det J_{g^{-1}}(y_1, \dots, y_n)| > 0.$$

Also,

$$J_{g^{-1}}(y) = J_g(g^{-1}(y))^{-1}.$$

Theorem 6.1. *Smooth Change of Variables*

Under the conditions just stated for X , g , and U , and if moreover

$$P(X \in U) = 1, \quad (6.6)$$

then Y admits the density

$$f_Y(y) = f_X(g^{-1}(y)) |\det J_g(g^{-1}(y))|^{-1} 1_V(y). \quad (6.7)$$

Proof. The proof consists in checking that for any bounded function $h : \mathbb{R} \rightarrow \mathbb{R}$,

$$E[h(Y)] = \int_{\mathbb{R}^n} h(y)\psi(y)dy, \quad (6.8)$$

where ψ is the function on the right-hand side of (6.7). Indeed, taking $h(y) = 1_{\{y \leq a\}} = 1_{\{y_1 \leq a_1\}} \cdots 1_{\{y_n \leq a_n\}}$, (6.8) reads

$$P(Y_1 \leq a_1, \dots, Y_n \leq a_n) = \int_{-\infty}^{a_1} \cdots \int_{-\infty}^{a_n} \psi(y_1, \dots, y_n) dy_1 \cdots dy_n.$$

To prove that (6.8) holds with the appropriate ψ , one just uses the basic rule of change of variables of calculus (see Rudin, 1976, §10.9):

$$\int_U u(x)dx = \int_{g(U)} u(g^{-1}(y))|\det J_{g^{-1}}(y)|dy.$$

Indeed, using this rule

$$\begin{aligned} E[h(Y)] &= E[h(g(X))] \\ &= \int_U h(g(x))f_X(x)dx \\ &= \int_V h(y)f_X(g^{-1}(y))|\det J_{g^{-1}}(y)|dy. \quad \square \end{aligned}$$

Example 6.1. Invertible Affine Transformations.

Here $U = \mathbb{R}^n$ and g is an affine function $g(x) = Ax + b$, where A is a $n \times n$ invertible matrix and $b \in \mathbb{R}^n$. Then $|\det J_{g^{-1}}(y)| = \frac{1}{|\det A|}$. Therefore, the random vector

$$Y = AX + B \quad (6.9)$$

admits the density

$$f_Y(y) = f_X(A^{-1}(y - b))\frac{1}{|\det A|}. \quad (6.10)$$

◇

6.2 Order Statistics

Let X_1, \dots, X_n be independent random variables with the same p.d.f. $f(x)$. It was shown (see Exercise 4.1) that there is zero probability that two or more among X_1, \dots, X_n take the same value. Thus one can define unambiguously the random variables Z_1, \dots, Z_n obtained by reordering X_1, \dots, X_n in increasing order:

$$\begin{cases} Z_i \in \{X_1, \dots, X_n\}, \\ Z_1 < Z_2 < \cdots < Z_n. \end{cases} \quad (6.11)$$

In particular, $Z_1 = \min(X_1, \dots, X_n)$ and $Z_n = \max(X_1, \dots, X_n)$.

Theorem 6.2. *Order Statistics Probability Density*

The probability density of $Z = (Z_1, \dots, Z_n)$ is

$$f_Z(z_1, \dots, z_n) = n! \left\{ \prod_{j=1}^n f(z_j) \right\} 1_C(z_1, \dots, z_n), \quad (6.12)$$

where

$$C = \{(z_1, \dots, z_n) \in \mathbb{R}^n; z_1 < z_2 < \dots < z_n\}. \quad (6.13)$$

Proof. Let σ be the permutation of $\{1, \dots, n\}$ that orders X_1, \dots, X_n in ascending order, i.e.,

$$X_{\sigma(i)} = Z_i$$

(note that σ is a *random* permutation). For any set $A \subset \mathbb{R}^n$,

$$P(Z \in A) = P(Z \in A \cap C) = P(X_{\sigma} \in A \cap C) = \sum_{\sigma_o} P(X_{\sigma_o} \in A \cap C, \sigma = \sigma_o),$$

where the sum is over all permutations of $\{1, \dots, n\}$. Observing that $X_{\sigma_o} \in A \cap C$ implies $\sigma = \sigma_o$,

$$P(X_{\sigma_o} \in A \cap C, \sigma = \sigma_o) = P(X_{\sigma_o} \in A \cap C)$$

and therefore since the probability distribution of X_{σ_o} does not depend upon a fixed permutation σ_o (here we use independence and equidistribution of the X_i 's),

$$P(X_{\sigma_o} \in A \cap C) = P(X \in A \cap C).$$

Therefore,

$$\begin{aligned} P(Z \in A) &= \sum_{\sigma_o} P(X \in A \cap C) = n! P(X \in A \cap C) \\ &= n! \int_{A \cap C} f_X(x) dx = \int_A n! f_X(x) 1_C(x) dx. \end{aligned} \quad \square$$

Example 6.2. *Ordering Uniforms*

Applying formula (6.12) to the situation where the X_i 's are uniformly distributed over $[a, b]$ gives

$$f_Z(z_1, \dots, z_n) = \frac{n!}{(b-a)^n} 1_{[a,b]^n}(z_1, \dots, z_n) 1_C(z_1, \dots, z_n). \quad (6.14)$$

In particular, since $\int_{\mathbb{R}^n} f_Z(z) dz = 1$,

$$\int_a^b \cdots \int_a^b 1_C(z_1, \dots, z_n) dz_1 \cdots dz_n = \frac{(b-a)^n}{n!}. \quad (6.15)$$

7 Conditional Expectation of Discrete Variables

7.1 Definition and Basic Properties

This section introduces the notion of conditional expectation for *discrete* random variables and gives the results that are needed in the study of Markov chains with a countable state space.

Definition 7.1. Conditional Expectation

Let X and Y be two discrete random variables taking their values in the denumerable sets F and G , respectively, and let $g : F \times G \rightarrow \mathbb{R}$ be a nonnegative function. One defines the function $\psi : G \rightarrow \mathbb{R}$ by

$$\psi(y) = \sum_{x \in F} g(x, y)P(X = x | Y = y). \quad (7.1)$$

For each $y \in G$, $\psi(y)$ is called the conditional expectation of $g(X, Y)$ given $Y = y$, and is denoted by $E^{Y=y}[g(X, Y)]$, or $E[g(X, Y) | Y = y]$:

$$E^{Y=y}[g(X, Y)] = \psi(y). \quad (7.2)$$

The random variable $\psi(Y)$ is called the conditional expectation of $g(X, Y)$ given Y , and is denoted by $E^Y[g(X, Y)]$ (or $E[g(X, Y) | Y]$):

$$E^Y[g(X, Y)] = \psi(Y). \quad (7.3)$$

Note that if $g \geq 0$ and $E[g(X, Y)] < \infty$, then $\psi(y) < \infty$ for all $y \in G$ such that $P(Y = y) > 0$; that is to say, $\psi(Y) < \infty$ almost surely. This follows from

$$\sum_{y \in G} \psi(y)P(Y = y) = E[g(X, Y)] < \infty.$$

Therefore, if $g : F \times G \rightarrow \mathbb{R}$ is a function of arbitrary sign such that $E[|g(X, Y)|] < \infty$, one can define $E^{Y=y}[g^\pm(X, Y)]$, since both terms on the right-hand side are finite, in view of the previous remark and of the fact that $E[g^\pm(X, Y)] < \infty$. One may then define

$$E^{Y=y}[g(X, Y)] = \sum_{x \in F} g(x, y)P(X = x | Y = y)$$

(call this $\psi(y)$). The conditional expectation of $g(X, Y)$ given Y is then defined by (7.3) as for the nonnegative case.

Finally, for any event A , we have that $P(A | Y)$ or $P^Y(A)$, the probability of A given Y , is equal by definition to $E^Y[1_A]$.

Example 7.1.

Let X_1 and X_2 be two independent Poisson random variables with respective means $\theta_1 > 0$ and $\theta_2 > 0$. We seek to compute $E^{X_1+X_2}[X_1]$, or $E^Y[X]$ where $X = X_1$, $Y = X_1 + X_2$. Following the instructions of Definition 7.1, we must first compute

$$\begin{aligned} P(X = x | Y = y) &= \frac{P(X = x, Y = y)}{P(Y = y)} = \frac{P(X_1 = x, X_1 + X_2 = y)}{P(X_1 + X_2 = y)} \\ &= \binom{y}{x} \left(\frac{\theta_1}{\theta_1 + \theta_2} \right)^x \left(\frac{\theta_2}{\theta_1 + \theta_2} \right)^{y-x} 1_{\{y \geq x\}}. \end{aligned}$$

Therefore, setting $\alpha = \frac{\theta_1}{\theta_1 + \theta_2}$,

$$\psi(y) = E^{Y=y}[X] = \sum_{x=0}^y x \binom{y}{x} \alpha^x (1 - \alpha)^{y-x} = \alpha y.$$

Finally, $E^Y[X] = \psi(Y) = \alpha Y$, that is,

$$E^{X_1+X_2}[X_1] = \frac{\theta_1}{\theta_1 + \theta_2} (X_1 + X_2). \quad \diamond$$

The first property of conditional expectation, linearity, is obvious from the definition: For all $\lambda_1, \lambda_2 \in \mathbb{R}$,

$$E^Y[\lambda_1 g_1(X, Y) + \lambda_2 g_2(X, Y)] = \lambda_1 E^Y[g_1(X, Y)] + \lambda_2 E^Y[g_2(X, Y)] \quad (7.4)$$

whenever the conditional expectations thereof are well-defined and do not produce $\infty - \infty$ forms.

Monotonicity is equally obvious: if $g_1(x, y) \leq g_2(x, y)$, then

$$E^Y[g_1(X, Y)] \leq E^Y[g_2(X, Y)]. \quad (7.5)$$

Next, we have

$$E[E^Y[g(X, Y)]] = E[g(X, Y)], \quad (7.6)$$

since the left-hand side is $\sum_{y \in G} \psi(y) P(Y = y) = \sum_{y \in G} \sum_{x \in F} g(x, y) P(X = x | Y = y) P(Y = y) = \sum_x \sum_y g(x, y) P(X = x, Y = y)$.

Also,

$$E^Y[w(Y)] = w(Y), \quad (7.7)$$

and more generally,

$$E^Y[w(Y)h(X, Y)] = w(Y)E^Y[h(X, Y)]. \quad (7.8)$$

This again follows from the definitions. Of course, we assume that the left-hand sides of (7.7) and (7.8) are well-defined.

If X and Y are independent and if $v : F \rightarrow \mathbb{R}$ is such that $E[|v(X)|] < \infty$, then

$$E^Y[v(X)] = E[v(X)]. \quad (7.9)$$

Indeed, $E^{Y=y}[v(X)] = \sum_{x \in F} v(x) P(X = x | Y = y) = \sum_{x \in F} v(x) P(X = x)$.

7.2 Successive Conditioning

In the definition of conditional expectation, we assumed G to be denumerable but otherwise arbitrary. Take G of the product form $G = G_1 \times G_2$, that is, $Y = (Y_1, Y_2)$, where Y_1 and Y_2 take their values in the denumerable sets G_1 and G_2 , respectively. In this situation, we use the more developed notation

$$E^Y[g(X, Y)] = E^{Y_1, Y_2}[g(X, Y_1, Y_2)]. \quad (7.10)$$

Theorem 7.1. Successive Conditioning

Under either one of the following conditions

- (i) $g(X, Y) \geq 0$, a.s.,
- (ii) $E[|g(X, Y)|] < \infty$,

we have

$$E^{Y_2}[E^{Y_1, Y_2}[g(X, Y_1, Y_2)]] = E^{Y_2}[g(X, Y_1, Y_2)]. \quad (7.11)$$

Proof. Set

$$\psi(Y_1, Y_2) = E^{Y_1, Y_2}[g(X, Y_1, Y_2)].$$

We must show that

$$E^{Y_2}[\psi(Y_1, Y_2)] = E^{Y_2}[g(X, Y_1, Y_2)].$$

But

$$\psi(y_1, y_2) = \sum_x g(x, y_1, y_2)P(X = x | Y_1 = y_1, Y_2 = y_2)$$

and

$$E^{Y_2=y_2}[\psi(Y_1, Y_2)] = \sum_{y_1} \psi(y_1, y_2)P(Y_1 = y_1 | Y_2 = y_2),$$

that is,

$$\begin{aligned} E^{Y_2=y_2}[\psi(Y_1, Y_2)] &= \sum_{y_1} \sum_x g(x, y_1, y_2)P(X = x | Y_1 = y_1, Y_2 = y_2)P(Y_1 = y_1 | Y_2 = y_2) \\ &= \sum_{y_1} \sum_x g(x, y_1, y_2)P(X = x, Y_1 = y_1 | Y_2 = y_2) \\ &= E^{Y_2=y_2}[g(X, Y_1, Y_2)]. \quad \square \end{aligned}$$

Exercise 7.1.

Let X_1 and X_2 be two independent identically distributed random variables with values in the denumerable set E . Assume that $E[|X_1|] < \infty$. Show that

$$E^{X_1+X_2}[X_1] = \frac{X_1 + X_2}{2}.$$

Solution. $E^{X_1+X_2}[X_1] = \psi(X_1 + X_2)$, and by symmetry, $E^{X_1+X_2}[X_2] = \psi(X_1 + X_2)$. Therefore,

$$2\psi(X_1 + X_2) = E^{X_1+X_2}[X_1] + E^{X_1+X_2}[X_2] = E^{X_1+X_2}[X_1 + X_2] = X_1 + X_2.$$

8 The Strong Law of Large Numbers

8.1 Borel–Cantelli Lemma

Consider a sequence of events $\{A_n\}_{n \geq 1}$ where the index n can be, if one wishes, interpreted as time. One is interested in the probability that A_n occurs infinitely often, that is, the probability of the event

$$\{\omega; \omega \in A_n \text{ for an infinity of indices } n\}, \quad (8.1)$$

denoted by $\{A_n \text{ i.o.}\}$, where *i.o.* means *infinitely often*.

Lemma 8.1. *Borel–Cantelli Lemma*

Let $\{A_n\}_{n \geq 1}$ be a sequence of events such that

$$\sum_{n=1}^{\infty} P(A_n) < \infty. \quad (8.2)$$

Then

$$P(A_n \text{ i.o.}) = 0. \quad (8.3)$$

Proof. We first give a manageable expression for the event in (8.1). We have

$$\{A_n \text{ i.o.}\} = \bigcap_{n=1}^{\infty} \bigcup_{k \geq n} A_k.$$

Indeed, if ω belongs to the set on the right-hand side, then for *all* $n \geq 1$, ω belongs to at least one among A_n, A_{n+1}, \dots , which implies that ω is in A_n for an infinite number of indices n . Conversely, if ω is in A_n for an infinite number of indices n , it is for *all* $n \geq 1$ in at least one of the sets A_n, A_{n+1}, \dots

Let now B_n be the set $\bigcup_{k \geq n} A_k$, which decreases as n increases, so that by the sequential continuity property of probability,

$$P(A_n \text{ i.o.}) = \lim_{n \uparrow \infty} P\left(\bigcup_{k \geq n} A_k\right).$$

But by sub- σ -additivity,

$$P\left(\bigcup_{k \geq n} A_k\right) \leq \sum_{k \geq n} P(A_k),$$

and by hypothesis (8.2), the right-hand side of this inequality goes to 0 as n goes to ∞ . \square

8.2 Almost-Sure Convergence

Definition 8.1. Almost-Sure Convergence

A sequence $\{Z_n\}_{n \geq 1}$ of real random variables is said to converge P-almost surely (P-a.s.) to the real random variable Z if

$$P(\lim_{n \uparrow \infty} Z_n = Z) = 1. \quad (8.4)$$

(Paraphrasing: For all ω outside a set N of null probability, $\lim_{n \uparrow \infty} Z_n(\omega) = Z(\omega)$).

The following is a first example of the usefulness of the Borel–Cantelli lemma.

Example 8.1.

Let $\{X_n\}_{n \geq 1}$ be an independent family of $\{0, 1\}$ -valued random variables such that

$$\sum_{n=1}^{\infty} P(X_n = 1) < \infty. \quad (8.5)$$

Then, in view of the Borel–Cantelli lemma, $P(X_n = 1 \text{ i.o.}) = 0$. But the event $\overline{\{X_n = 1 \text{ i.o.}\}}$ is exactly $\{\omega; \lim_{n \uparrow \infty} X_n(\omega) = 0\}$, since the process takes only the values 0 or 1. Therefore, (8.5) implies that $\lim_{n \uparrow \infty} X_n = 0$ almost surely. \diamond

We now give conditions for almost-sure convergence, using the Borel–Cantelli lemma.

Let $\{\epsilon_n\}_{n \geq 1}$ be a sequence of positive numbers converging to 0. If for a given ω , $|Z_n(\omega) - Z(\omega)| < \epsilon_n$ for all but a finite number of indices n , then $\lim_{n \uparrow \infty} Z_n(\omega) = Z(\omega)$. Therefore, by the Borel–Cantelli lemma, we have the following theorem.

Theorem 8.1.

Let $\{Z_n\}_{n \geq 1}$ and Z be random variables. If

$$\sum_{n \geq 1} P(|Z_n - Z| \geq \epsilon_n) < \infty \quad (8.6)$$

for some sequence of positive numbers $\{\epsilon_n\}_{n \geq 1}$ converging to 0, then the sequence $\{Z_n\}_{n \geq 1}$ converges P-a.s. to Z .

The next result is very similar, but this time we have a necessary and sufficient condition.

Theorem 8.2. Fundamental Criterion of Almost-Sure Convergence

The sequence $\{Z_n\}_{n \geq 1}$ of real random variables converges P-a.s. to the real random variable Z if and only if for all $\epsilon > 0$,

$$P(|Z_n - Z| \geq \epsilon \text{ i.o.}) = 0. \quad (8.7)$$

Proof. For the necessity, observe that

$$\{|Z_n - Z| \geq \epsilon \text{ i.o.}\} \subset \overline{\{\omega; \lim_{n \uparrow \infty} Z_n(\omega) = Z(\omega)\}},$$

and therefore

$$P(|Z_n - Z| \geq \epsilon \text{ i.o.}) \leq 1 - P(\lim_{n \uparrow \infty} Z_n = Z) = 0.$$

For the sufficiency, let N_k be the last index n such that $|Z_n - Z| \geq \frac{1}{k}$ (set $N_k = \infty$ if $|Z_n - Z| \geq \frac{1}{k}$ for all $n \geq 1$). By (8.7) with $\epsilon = \frac{1}{k}$, we have $P(N_k < \infty) = 0$. By sub- σ -additivity, $P(\cup_{k \geq 1} \{N_k < \infty\}) = 0$. Equivalently, $P(N_k = \infty, \text{ for all } k \geq 1) = 1$, which implies $P(\lim_{n \uparrow \infty} Z_n = Z) = 1$. \square

A law of large numbers is a statement about the almost sure convergence of the empirical average

$$\frac{S_n}{n} = \frac{X_1 + \cdots + X_n}{n},$$

where $\{X_n\}_{n \geq 1}$ is an i.i.d. (independent and identically distributed) sequence of random variables.

Theorem 8.3. *Kolmogorov's SLLN*

Let $\{X_n\}_{n \geq 1}$ be an i.i.d. sequence of random variables such that

$$E[|X_1|] < \infty. \tag{8.8}$$

Then, P-a.s.

$$\lim_{n \uparrow \infty} \frac{S_n}{n} = E[X_1]. \tag{8.9}$$

In other words, there exists a set N of probability 0 such that if ω is not in N , then

$$\lim_{n \uparrow \infty} \frac{S_n(\omega)}{n} = E[X_1].$$

The empirical average is asymptotically equal to the probabilistic average. Such is the “physical” content of Kolmogorov’s strong law of large numbers (1933). Emile Borel proved the strong law of large numbers (SLLN) in 1909 in the special case where $X_n = 0$ or 1, with $P(X_n = 1) = p$, thus showing in particular that the average fraction of heads in a fair game of coins should tend to $\frac{1}{2}$ as the number of tosses increases indefinitely.

There are numerous versions of the SLLN extending Kolmogorov’s result to situations where independence of the sequence $\{X_n\}_{n \geq 1}$ is not required. As a matter of fact, the ergodic theorem for irreducible positive recurrent Markov chains (see Chapter 3) is one of these extensions.

We shall give the proof of Kolmogorov’s SLLN in the last subsection. Before that we shall prove Borel’s SLLN. It is based on the Borel–Cantelli lemma and on Markov’s inequality.

8.3 Markov's Inequality

Theorem 8.4. Markov's inequality

Let X be a random variable with values in \mathbb{R} , $f: \mathbb{R} \rightarrow \mathbb{R}_+$, and $a > 0$. We then have

$$P(f(X) \geq a) \leq \frac{E[f(X)]}{a}. \quad (8.10)$$

Proof. From the inequality

$$f(X) \geq a 1_{\{f(X) \geq a\}}$$

it follows by taking expectations that

$$E[f(X)] \geq aE[1_{\{f(X) \geq a\}}] = aP(f(X) \geq a). \quad \square$$

Specializing Markov's inequality to $f(x) = (x - m)^2$, $a = \epsilon^2 > 0$, we obtain *Chebyshev's inequality*

$$P(|X - m| \geq \epsilon) \leq \frac{\sigma^2}{\epsilon^2}. \quad (8.11)$$

Example 8.2. Borel's SLLN

This is Kolmogorov's SLLN when the X_n 's are bounded, say without loss of generality by 1. We shall bound the probability that $\left|\frac{S_n}{n} - m\right|$ exceeds some $\epsilon > 0$ where $m = E[X_1]$. For this we apply Markov's inequality

$$P\left(\left|\frac{S_n}{n} - m\right| \geq \epsilon\right) = P\left(\left(\frac{S_n}{n} - m\right)^4 \geq \epsilon^4\right) \leq \frac{E\left[\left(\frac{S_n}{n} - m\right)^4\right]}{\epsilon^4}.$$

Now,

$$\left(\frac{S_n}{n} - m\right)^4 = \frac{\left(\sum_{i=1}^n (X_i - m)\right)^4}{n^4},$$

and therefore

$$P\left(\left|\frac{S_n}{n} - m\right| \geq \epsilon\right) \leq \frac{E\left[\left(\sum_{i=1}^n (X_i - m)\right)^4\right]}{n^4 \epsilon^4}.$$

If we can prove that

$$E\left[\left(\sum_{i=1}^n (X_i - m)\right)^4\right] \leq Kn^2 \quad (8.12)$$

for some finite K , then

$$P\left(\left|\frac{S_n}{n} - m\right| \geq \epsilon\right) \leq \frac{K}{n^2 \epsilon^4},$$

and in particular, with $\epsilon = n^{-\frac{1}{8}}$,

$$P\left(\left|\frac{S_n}{n} - m\right| \geq n^{-\frac{1}{8}}\right) \leq \frac{K}{n^{\frac{3}{2}}},$$

from which it follows that

$$\sum_{n=1}^{\infty} P\left(\left|\frac{S_n}{n} - m\right| \geq n^{-\frac{1}{3}}\right) < \infty.$$

Therefore, by Theorem 8.1, $\left|\frac{S_n}{n} - m\right|$ converges P-a.s. to 0.

It remains to prove (8.12).

To simplify notation, call Y_i the r.v. $X_i - m$, and remember that $E[Y_i] = 0$. Also, in view of the independence hypothesis, $E[Y_1 Y_2 Y_3 Y_4] = E[Y_1]E[Y_2]E[Y_3]E[Y_4] = 0$, $E[Y_1 Y_2^3] = E[Y_1]E[Y_2^3] = 0$, and the like. Finally, in the development

$$E\left[\left(\sum_{i=1}^n Y_i\right)^4\right] = \sum_{i,j,k,\ell=1}^n E[Y_i Y_j Y_k Y_\ell],$$

only the terms of the form $E[Y_i^4]$ and $E[Y_i^2 Y_j^2]$ ($i \neq j$) remain. There are n terms of the first type and $3n(n-1)$ terms of the second type. Therefore,

$$E\left[\left(\sum_{i=1}^n Y_i\right)^4\right] = nE[Y_1^4] + 3n(n-1)E[Y_1^2 Y_2^2],$$

which is smaller than Kn^2 for some finite K . ◇

Example 8.3. Chernoff's Bound

Let X be a real-valued random variable such that for all $t \in \mathbb{R}$, $E[e^{tX}] < \infty$. Set $\psi(t) = \log E[e^{tX}]$. Observing that for $t \geq 0$, $a \in \mathbb{R}$, $X > a$ is equivalent to $e^{tX} > e^{ta}$, and using Markov's inequality, we find that

$$P(X > a) = P(e^{tX} > e^{ta}) \leq \frac{E[e^{tX}]}{e^{ta}} = \exp\{-(at - \psi(t))\}.$$

This being true for all $t \geq 0$, it follows that

$$P(X > a) \leq e^{-h(a)}, \tag{8.13}$$

where

$$h(a) = \sup_{t \geq 0} \{at - \psi(t)\}. \tag{8.14}$$

◇

8.4 Proof of Kolmogorov's SLLN

We shall assume without loss of generality that $E[X_1] = 0$.

Lemma 8.2. Kolmogorov's Inequality

Let X_1, \dots, X_n be independent random variables such that for all $i \in [1, n]$,

$$E[|X_i|^2] < \infty, \quad E[X_i] = 0. \quad (8.15)$$

Then for all $\lambda > 0$,

$$P(\max_{1 \leq k \leq n} |S_k| \geq \lambda) \leq \frac{E[S_n^2]}{\lambda^2}, \quad (8.16)$$

where $S_k = X_1 + \dots + X_k$.

Proof. Let T be the first (random) index $k \in [1, n]$, such that $|S_k| \geq \lambda$, with $T = \infty$ if $\max_{1 \leq k \leq n} |S_k| < \lambda$. For $k \leq n$,

$$\begin{aligned} E[S_n^2 1_{\{T=k\}}] &= E[1_{\{T=k\}} \{(S_n - S_k)^2 + 2S_k(S_n - S_k) + S_k^2\}] \\ &= E[1_{\{T=k\}} \{(S_n - S_k)^2 + S_k^2\}] \geq E[1_{\{T=k\}} S_k^2], \end{aligned}$$

where we used the fact that $1_{\{T=k\}} S_k$ is a function of X_1, \dots, X_k and therefore independent of $S_n - S_k$, so that $E[1_{\{T=k\}} S_k(S_n - S_k)] = E[1_{\{T=k\}} S_k] E[S_n - S_k] = 0$. Therefore,

$$\begin{aligned} E[|S_n|^2] &\geq \sum_{k=1}^n E[1_{\{T=k\}} S_k^2] \geq \sum_{k=1}^n E[1_{\{T=k\}} \lambda^2] = \lambda^2 \sum_{k=1}^n P(T = k) \\ &= \lambda^2 P(T \leq n) = \lambda^2 P(\max_{1 \leq k \leq n} |S_k| \geq \lambda). \quad \square \end{aligned}$$

The following corollary contains a proof of the SLLN when $E[|X_n|^2] < \infty$.

Corollary 8.1.

Let $\{X_n\}_{n \geq 1}$ be a sequence of independent random variables such that for all $n \geq 1$,

$$E[|X_n|^2] < \infty, \quad E[X_n] = 0. \quad (8.17)$$

If

$$\sum_{n \geq 1} \frac{E[X_n^2]}{n^2} \leq \infty, \quad (8.18)$$

then

$$\lim_{n \uparrow \infty} \frac{1}{n} \sum_{k=1}^n X_k = 0, \quad \text{P-a.s.} \quad (8.19)$$

Proof. If $2^{k-1} \leq n \leq 2^k$, then $\frac{|S_n|}{n} \geq \epsilon$ implies $\frac{|S_n|}{2^{k-1}} \geq \epsilon$. Therefore, for all $\epsilon > 0$, and all $k \geq 1$,

$$\begin{aligned} P\left(\frac{|S_n|}{n} \geq \epsilon \text{ for some } n \in [2^{k-1}, 2^k]\right) &\leq P(|S_n| \geq \epsilon 2^{k-1} \text{ for some } n \in [2^{k-1}, 2^k]) \\ &\leq P(|S_n| \geq \epsilon 2^{k-1} \text{ for some } n \in [1, 2^k]) \\ &= P\left(\max_{1 \leq n \leq 2^k} |S_n| \geq \epsilon 2^{k-1}\right) \leq \frac{4}{\epsilon^2} \frac{1}{(2^k)^2} \sum_{n=1}^{2^k} E[X_n^2], \end{aligned}$$

where the last inequality follows from Kolmogorov's inequality. But defining $m = m(n)$ by $2^{m-1} \leq n < 2^m$, we have

$$\sum_{k=1}^{\infty} \frac{1}{(2^k)^2} \sum_{n=1}^{2^k} E[X_n^2] = \sum_{n=1}^{\infty} E[X_n^2] \sum_{j=m}^{\infty} \frac{1}{(2^j)^2},$$

which is bounded by

$$\sum_{n=1}^{\infty} E[X_n^2] \frac{K}{(2^m)^2} \leq \sum_{n=1}^{\infty} E[X_n^2] \frac{K}{n^2}$$

for some finite K . Therefore, by (8.18),

$$\sum_{k=1}^{\infty} P\left(\frac{|S_n|}{n} \geq \epsilon \text{ for some } n; 2^{k-1} < n \leq 2^k\right) < \infty.$$

and by the Borel–Cantelli lemma,

$$P\left(\frac{|S_n|}{n} \geq \epsilon \text{ i.o.}\right) = 0.$$

The result then follows from Theorem 8.2. □

Having proved the corollary, it now remains to get rid of the assumption $E[|X_n|^2] < \infty$, and the natural technique for this is truncation. Define

$$\tilde{X}_n = \begin{cases} X_n & \text{if } |X_n| \leq n, \\ 0 & \text{otherwise.} \end{cases}$$

Since $E[|X_1|] < \infty$, we have by dominated convergence (see Theorem 3.2 of the Appendix) that $\lim_{n \uparrow \infty} E[X_1 1_{\{|X_1| \leq n\}}] = E[X_1] = 0$. Since X_n has the same distribution as X_1 ,

$$\lim_{n \uparrow \infty} E[\tilde{X}_n] = \lim_{n \uparrow \infty} E[X_n 1_{\{|X_n| \leq n\}}] = \lim_{n \uparrow \infty} E[X_1 1_{\{|X_1| \leq n\}}] = E[X_1] = 0.$$

In particular (Cesaro's Lemma, Lemma 1.5 of the Appendix),

$$\lim_{n \uparrow \infty} \frac{1}{n} \sum_{k=1}^n E[\tilde{X}_k] = 0.$$

Also,

$$\begin{aligned} E[|X_1|] &= \sum_{n=0}^{\infty} E[|X_1|1_{(n,n+1]}(|X_1|)] \geq \sum_{n=0}^{\infty} E[n1_{(n,n+1]}(|X_1|)] \\ &= \sum_{n=0}^{\infty} n(P(|X_1| > n) - P(|X_1| > n+1)) = \sum_{n=1}^{\infty} P(|X_1| > n). \end{aligned}$$

In particular,

$$\sum_{n=1}^{\infty} P(|X_n| > n) = \sum_{n=1}^{\infty} P(|X_1| > n) \leq E[|X_1|] < \infty,$$

and therefore, by the Borel–Cantelli lemma,

$$P(\tilde{X}_n \neq X_n \text{ i.o.}) = P(|X_n| > n \text{ i.o.}) = 0.$$

Therefore, to prove the SLLN it suffices to show that

$$\lim_{n \uparrow \infty} \frac{1}{n} \sum_{k=1}^n (\tilde{X}_k - E[\tilde{X}_k]) = 0.$$

In view of the preceding corollary, it suffices to prove that

$$\sum_{n=1}^{\infty} \frac{E[(\tilde{X}_n - E[\tilde{X}_n])^2]}{n^2} < \infty.$$

But

$$E[(\tilde{X}_n - E[\tilde{X}_n])^2] \leq E[\tilde{X}_n^2] = E[X_1^2 1_{(|X_1| \leq n)}].$$

It is therefore enough to show that

$$\sum_{n=1}^{\infty} \frac{E[X_1^2 1_{(|X_1| \leq n)}]}{n^2} < \infty.$$

The left-hand side of the above inequality is equal to

$$\sum_{n=1}^{\infty} \frac{1}{n^2} \sum_{k=1}^n E[X_1^2 1_{\{k-1 < |X_1| \leq k\}}] = \sum_{k=1}^{\infty} \left(\sum_{n=k}^{\infty} \frac{1}{n^2} \right) E[X_1^2 1_{\{k-1 < |X_1| \leq k\}}].$$

Using the fact that

$$\sum_{n=k}^{\infty} \frac{1}{n^2} \leq \frac{1}{k^2} + \int_k^{\infty} \frac{1}{x^2} dx = \frac{1}{k^2} + \frac{1}{k} \leq \frac{2}{k}$$

(draw the graph of $x \rightarrow x^{-2}$), this is less than or equal to

$$\begin{aligned} \sum_{k=1}^{\infty} \frac{2}{k} E[X_1^2 1_{\{k-1 < |X_1| \leq k\}}] &= 2 \sum_{k=1}^{\infty} E\left[\frac{X_1^2}{k} 1_{\{k-1 < |X_1| \leq k\}}\right] \\ &\leq 2 \sum_{k=1}^{\infty} E[|X_1| 1_{\{k-1 < |X_1| \leq k\}}] \\ &= 2E[|X_1|] < \infty. \end{aligned}$$

□

Problems

1.1.1. Prove and generalize the following identities:

$$P(A \cup B) = P(A) + P(B) - P(A \cap B);$$

$$P(A \cup B \cup C) = P(A) + P(B) + P(C) - P(A \cap B) - P(B \cap C) - P(C \cap A) + P(A \cap B \cap C).$$

1.1.2. Prove that $P(A \cup B) = 1 - P(\bar{A} \cap \bar{B})$.

1.1.3. Give a probability model for three successive tosses of an unbiased die. What is the probability that one of these tosses results in a number that is the sum of the two other numbers?

1.1.4. In the coin model (Examples 1.3, 1.6, and 1.9) show that $\bigcup_{n=1}^{\infty} \bigcap_{k=n}^{\infty} \{X_k = 1\}$ is the event “after some finite random time, *all* tosses result in heads” (recall that heads = 1).

1.2.1. Consider the following probabilistic model: $\Omega = [0, 1]^2$, $P(A) = \text{area of } A$. Thus $\omega = (x, y)$, where $x, y \in [0, 1]$. Define $X(\omega) = x$, $Y(\omega) = y$. Show that X and Y are independent random variables.

1.2.2. Consider the following probability model: $\Omega = \{(x, y) \in \mathbb{R}^2, x^2 + y^2 \leq 1\}$, $P(A) = \frac{1}{\pi} \times (\text{area of } A)$. Thus $\omega = (x, y) \in \Omega$ is a point uniformly distributed inside the unit disk. Defining $X(\omega) = x$ and $Y(\omega) = y$, show that X and Y are *not* independent random variables.

1.2.3. Give a simple example of a probability space (Ω, \mathcal{F}, P) with three events A_1, A_2, A_3 that are pairwise independent, but *not* globally independent (i.e., the family $\{A_1, A_2, A_3\}$ is not independent).

1.2.4. Show that if $\{A_i\}_{i \in I}$ is an independent collection of events, then $\{\tilde{A}_i\}_{i \in I}$ is also, where, independently for each $i \in I$, $\tilde{A}_i = A_i$ or \bar{A}_i (for instance, with $I = \mathbb{N}$, $\tilde{A}_0 = A_0$, $\tilde{A}_1 = \bar{A}_1$, $\tilde{A}_2 = A_2$, $\tilde{A}_3 = \bar{A}_3, \dots$).

1.2.5. There are 3 cards. The first one has both faces red, the second one has both faces white, and the third one is white on one face, red on the other. A card is drawn at random, and the color of a randomly selected face of this card is announced. What is the winning strategy if you must bet on the color of the other face?

1.2.6. In the School for Intellectual Apartheid, students have been separated into three groups for pedagogical purposes. In group A, one finds students who individually have a probability of passing equal to 0.95. In group B this probability is 0.75, and in group C only 0.65. What is the probability that a student passing the course comes from group A? B? C?

1.3.1. A given insect of a specific breed has the probability θ of being a male. An entomologist seeks to collect exactly $M > 1$ males, and therefore stops hunting as soon as she captures M males. She has to capture an insect in order to determine its gender. What is the distribution of X , the number of insects she must catch to collect *exactly* M males?

1.3.2. Find the c.d.f. and the p.d.f. of $Y = X^2$, where X is a Gaussian random variable with mean 0 and variance $\sigma^2 > 0$.

1.3.3. Let T be a geometric random variable. Show that for any integers $k, k_0 \geq 1$, we have $P(T = k + k_0 | T > k_0) = P(T = k)$. Let X be an exponential random variable with mean $1/\lambda$. Show that for all $t, t_0 \in \mathbb{R}_+$, we have $P(X \geq t_0 + t | X \geq t_0) = P(X \geq t)$.

1.3.4. Let X be a real valued random variable with a p.d.f. $f(x)$ and let Y be an integer valued random variable with the distribution $P(X = k) = p_k, k \geq 0$. Find the p.d.f. of $Z = X + Y$ when X and Y are independent.

1.4.1. Let $X_i, i \in [0, n]$, be independent exponential random variables with the respective parameters $\lambda_i, i \in [0, n]$. Define $Z = \inf(X_1, \dots, X_n)$ and let J be the (random) index such that $X_J = Z$ (J is for almost all $\omega \in \Omega$ unambiguously defined in view of Exercise 4.1). Show that Z and J are independent, and give their respective distributions.

1.4.2. Let $X = (X_1, \dots, X_k)$ and $Y = (Y_1, \dots, Y_k)$ be two independent multinomial random vectors (see Example 4.1) of sizes (n, K) and (m, K) , respectively, and with the same parameters p_1, \dots, p_K . What is the distribution of $Z = X + Y$?

1.4.3. Let $\{X_n\}_{n \geq 1}$ be an i.i.d sequence of exponential random variables with mean $1/\theta$, where $\theta \in (0, \infty)$. Let $\{Y_n\}_{n \geq 1}$ be an i.i.d sequence with $P(Y_1 = 1) = 1 - P(Y_1 = 0) = p \in (0, 1)$, independent of $\{X_n\}_{n \geq 1}$. Let U_1, U_2, \dots be the successive indices n for which $Y_n = 1$, and define $S_1 = X_1 + \dots + X_{U_1}, S_2 = X_{U_1+1} + \dots + X_{U_2}, S_3 = X_{U_2+1} + \dots + X_{U_3}$, etc. Show that $\{S_n\}_{n \geq 1}$ is an i.i.d sequence of exponential random variables with mean $1/(p\theta)$.

1.5.1. Compute the mean and variance of the binomial random variable from its generating function. Do the same for the Poisson and the geometric random variables.

1.6.1. Find the probability distribution of the random variable Z_i , the i th smallest among X_1, \dots, X_n , when the X_i 's are independent $[0, 1]$ -uniform random variables.

1.6.2. Let X and Y be two independent random variables with a common exponential distribution of mean θ^{-1} . Give the p.d.f. of the vector $(X/Y, Y)$ and of the variable X/Y .

1.6.3. Consider a random segment of length U uniformly distributed on $[0, 1]$, and make of it two random pieces of length $X = UV, Y = U(1 - V)$, respectively, where V is uniformly distributed on $[0, 1]$ and independent of U . Compute the p.d.f. of the length UV of the first piece.

1.7.1. Let X_1 and X_2 be two independent random variables taking their values in $\{1, 2, \dots, N\}$, and uniformly distributed, that is, $P(X_1 = k) = P(X_2 = k) = \frac{1}{N}$ ($1 \leq k \leq N$). Compute $E^{\max(X_1, X_2)}\{X_1\}$.

1.7.2. Let X and Y be two discrete random variables with values in \mathbb{N} , and let $h : \mathbb{N} \rightarrow \mathbb{N}$ be one-to-one and onto. Show that for all $v : \mathbb{N} \rightarrow \mathbb{R}$ such that $E[|v(X)|] < \infty$, $E^Y[v(X)] = E^Z[v(X)]$, where $Z = h(Y)$.

1.7.3. Let (X_1, \dots, X_k) be a multinomial random vector (see (4.12)). Compute $E^{X_1}[X_2 + \dots + X_{k-1}]$ and $E^{X_1}[X_2]$.

1.8.1. Let $\{S_n\}_{n \geq 1}$ be an i.i.d. sequence of real random variables such that $P(S_1 \in (0, \infty)) = 1$ and $E[S_1] < \infty$, and let for each $t \geq 0$, $N(t) = \sum_{n \geq 1} 1_{(0,t]}(T_n)$, where $T_n = S_1 + \dots + S_n$. Prove that $\lim_{t \rightarrow \infty} \frac{N(t)}{t} = \frac{1}{E[S_1]}$.

1.8.2. Let X_n denote the position at time n of a “particle” moving on $E = \{0, 1, 2, \dots\}$. Suppose that at each time $n \geq 0$ there is a probability of eventually visiting “site” $0 \in E$ after time n that is bounded from below by $r > 0$. Show that the “process” $\{X_n\}_{n \geq 0}$ almost surely visits site 0 infinitely often.

1.8.3. Let X_1, \dots, X_n be i.i.d random variables satisfying the same condition as X in Example 8.3 (Chernoff’s bound). Show that

$$P\left(\frac{X_1 + \dots + X_n}{n} > a\right) \leq e^{-nh(a)}.$$

Use this result to show that $\frac{S_n}{n}$ converges to 0 almost surely.

Discrete-Time Markov Models

1 The Transition Matrix

1.1 Markov Property

Sequences of independent and identically distributed random variables *are* stochastic processes, but they are not always interesting as stochastic models because they behave more or less in the same way. In order to introduce more variability, one can allow for some dependence on the past, in the manner of deterministic recurrence equations. Discrete-time homogeneous Markov chains possess the required feature, since they can always be represented—at least distributionwise—by a stochastic recurrence equation $X_{n+1} = f(X_n, Z_{n+1})$, where $\{Z_n\}_{n \geq 1}$ is an i.i.d sequence, independent of the initial state X_0 .

The probabilistic dependence on the past is only through the previous state, but this limited amount of memory suffices to produce a great diversity of behaviors. For this reason Markov chains have found applications in many domains, including biology, physics, sociology, operations research, and engineering, where they provide qualitative and quantitative answers as well as precious insights for systems design.

This chapter gives the basic definitions concerning discrete-time homogeneous Markov chains and the classical examples that will illustrate the theory of the next chapters.

A sequence $\{X_n\}_{n \geq 0}$ of random variables with values in a set E is called a *discrete-time stochastic process* with *state space* E . In this book, the state space is countable, and its elements are denoted by i, j, k, \dots . If $X_n = i$, the process is said to be in state i at time n , or to visit state i at time n .

Definition 1.1. *Homogeneous Markov Chain*

Let $\{X_n\}_{n \geq 0}$ be a discrete-time stochastic process with countable state space E . If for all

integers $n \geq 0$ and all states $i_0, i_1, \dots, i_{n-1}, i, j$,

$$P(X_{n+1} = j \mid X_n = i, X_{n-1} = i_{n-1}, \dots, X_0 = i_0) = P(X_{n+1} = j \mid X_n = i) \quad (1.1)$$

whenever both sides are well-defined, this stochastic process is called a *Markov chain*. It is called a *homogeneous Markov chain (HMC)* if in addition, the right-hand side of (1.1) is independent of n .

Property (1.1) is the *Markov property*. The matrix $\mathbf{P} = \{p_{ij}\}_{i,j \in E}$, where

$$p_{ij} = P(X_{n+1} = j \mid X_n = i), \quad (1.2)$$

is the *transition matrix* of the HMC. Since its entries are probabilities, and since a transition from any state i must be to *some* state, it follows that

$$p_{ij} \geq 0, \quad \sum_{k \in E} p_{ik} = 1$$

for all states i, j . A matrix \mathbf{P} indexed by E and satisfying the above properties is called a *stochastic matrix*. The state space may be infinite, and therefore such a matrix is in general not of the kind studied in linear algebra. However, the basic operations of addition and multiplication will be defined by the same formal rules. For instance, with $A = \{a_{ij}\}_{i,j \in E}$ and $B = \{b_{ij}\}_{i,j \in E}$, the product $C = AB$ is the matrix $\{c_{ij}\}_{i,j \in E}$, where $c_{ij} = \sum_{k \in E} a_{ik}b_{kj}$. The notation $x = \{x_i\}_{i \in E}$ formally represents a *column* vector, and x^T is a row vector, the transpose of x . For instance, $y = \{y_i\}_{i \in E}$ given by $y^T = x^T A$ is defined by $y_i = \sum_{k \in E} x_k a_{ki}$. Similarly, $z = \{z_i\}_{i \in E}$ given by $z = Ax$ is defined by $z_i = \sum_{k \in E} a_{ik}x_k$.

Proving the Markov property is not, in general, a difficult task, and Theorems 2.1 and 2.2 below will suffice in most situations. However, there are cases outside their scope, and the following one is quite important, both in theory and in applications.

Example 1.1. Machine Replacement

Let $\{U_n\}_{n \geq 1}$ be a sequence of i.i.d random variables taking their values in $\{1, 2, \dots, +\infty\}$. The random variable U_n can be interpreted as the lifetime of some machine, the n th one, which is replaced by the $(n + 1)$ st one upon failure. Thus at time 0, machine 1 is put in service until it breaks down at time U_1 , whereupon it is immediately replaced by machine 2, which breaks down at time $U_1 + U_2$, and so on. The elapsed time in service of the current machine at time n is denoted by X_n . Thus, the process $\{X_n\}_{n \geq 0}$ takes its values in $E = \mathbb{N}$ and increases linearly from 0 at time $R_k = \sum_{i=1}^k U_i$ to $U_{k+1} - 1$ at time $R_{k+1} - 1$.

The sequence $\{R_k\}_{k \geq 0}$ defined in this way, with $R_0 = 0$, is called a *renewal sequence*, and X_n is called the *backward recurrence time* at time n (see Fig. 2.1.1). There is a rich and useful theory associated with renewal sequences, the so-called *renewal theory*. It will be developed in Chapter 4.

The process $\{X_n\}_{n \geq 0}$ is an HMC with state space $E = \mathbb{N}$, and the nonnull entries of its transition matrix are of the form $p_{i,i+1}$ and $p_{i,0} = 1 - p_{i,i+1}$, where

$$p_{i,i+1} = \frac{P(U_1 > i + 1)}{P(U_1 > i)}. \quad (1.3)$$




Figure 2.1.1. Backward recurrence time

To prove this, one must first verify (1.1), that is, writing $B = \{X_0 = i_0, \dots, X_{n-1} = i_{n-1}\}$,

$$\frac{P(X_{n+1} = j, X_n = i, B)}{P(X_n = i, B)} = \frac{P(X_{n+1} = j, X_n = i)}{P(X_n = i)}$$

for sequences $i_0, \dots, i_{n-1}, i, j$ such that $P(B, X_n = i, X_{n+1} = j) > 0$. In particular, $j = i + 1$ or 0 , and $i_{n-1} = i - 1, \dots, i_{n-i} = 0$.

Let $\nu(n)$ be the number of renewal times R_k in the interval $[1, n]$. Taking, for instance, $j = i + 1$, and writing $D = \{X_{n-i-1} = i_{n-i-1}, \dots, X_0 = i_0\}$, we have

$$\begin{aligned} P(X_{n+1} = j, X_n = i, B) &= P(X_{n+1} = i + 1, X_n = i, X_{n-1} = i - 1, \dots, X_{n-i} = 0, D) \\ &= \sum_{k=0}^{\infty} P(X_{n+1} = i + 1, X_n = i, X_{n-1} = i - 1, \dots, X_{n-i} = 0, D, \nu(n) = k). \end{aligned}$$

The general term in the latter sum equals $P(U_{k+1} > i + 1, R_k = n - i, D) = P(U_{k+1} > i + 1)P(R_k = n - i, D) = P(U_1 > i + 1)P(R_k = n - i, D)$. The independence of $\{U_n\}_{n \geq 1}$ has been used for the first equality, and the identity of the distributions of U_{k+1} and U_1 for the second one. Therefore,

$$P(X_{n+1} = i + 1, X_n = i, B) = P(U_1 > i + 1) \left(\sum_{k=0}^{\infty} P(R_k = n - i, D) \right).$$

Similar computations yield

$$P(X_n = i, B) = P(U_1 > i) \left(\sum_{k=0}^{\infty} P(R_k = n - i, D) \right),$$

so that

$$P(X_{n+1} = i + 1 | X_n = i, B) = \frac{P(U_1 > i + 1)}{P(U_1 > i)}.$$

The same calculations lead to the same evaluation for $P(X_{n+1} = i + 1 | X_n = i)$. This proves the announced results. □

The above example is atypical. Proving the Markov property and computing the transition probabilities are usually much easier. Most of the time, a representation of the state process in terms of a recurrence equation makes things easy (see Theorem 2.1 below). Nevertheless, there are a few tough cases.

Transition Graph

A transition matrix \mathbf{P} is sometimes represented by its *transition graph* G , a graph having for nodes (or vertices) the states of E . This graph has an oriented edge from i to j if and only if $p_{ij} > 0$, in which case this edge is adorned with the label p_{ij} .

The transition graph of the Markov chain of Example 1.1 is shown in Figure 2.1.2, where

$$p_i = \frac{P(U_1 = i + 1)}{P(U_1 > i)}.$$




Figure 2.1.2. Transition graph of the backward recurrence chain

1.2 Distribution of an HMC

The random variable X_0 is called the *initial state*, and its probability distribution ν ,

$$\nu(i) = P(X_0 = i), \tag{1.4}$$

is the *initial distribution*. From Bayes's sequential rule, $P(X_0 = i_0, X_1 = i_1, \dots, X_k = i_k) = P(X_0 = i_0)P(X_1 = i_1 | X_0 = i_0) \cdots P(X_k = i_k | X_{k-1} = i_{k-1}, \dots, X_0 = i_0)$, and

therefore, in view of the homogeneous Markov property and the definition of the transition matrix,

$$P(X_0 = i_0, X_1 = i_1, \dots, X_k = i_k) = v(i_0)p_{i_0i_1} \cdots p_{i_{k-1}i_k}. \quad (1.5)$$

The data (1.5) for all $k \geq 0$, all states i_0, i_1, \dots, i_k , constitute the *probability law*, or *distribution* of the HMC. Therefore we have the following result.

Theorem 1.1. *Distribution of an HMC*

The distribution of a discrete-time HMC is determined by its initial distribution and its transition matrix.

The distribution at time n of the chain is the vector v_n , where

$$v_n(i) = P(X_n = i). \quad (1.6)$$

From Bayes's rule of exclusive and exhaustive causes, $v_{n+1}(j) = \sum_{i \in E} v_n(i)p_{ij}$, that is, in matrix form, $v_{n+1}^T = v_n^T \mathbf{P}$. Iteration of this equality yields

$$v_n^T = v_0^T \mathbf{P}^n. \quad (1.7)$$

The matrix \mathbf{P}^n is called the *n-step transition matrix* because its general term is

$$p_{ij}(m) = P(X_{n+m} = j | X_n = i). \quad (1.8)$$

Indeed, using Bayes's sequential rule and the Markov property, one finds for the right-hand side of the latter equality

$$\sum_{i_1, \dots, i_{m-1} \in E} p_{ii_1} p_{i_1 i_2} \cdots p_{i_{m-1} j},$$

and this is the general term of the m th power of \mathbf{P} .

The Markov property (1.1) extends to

$$\begin{aligned} P(X_{n+1} = j_1, \dots, X_{n+k} = j_k | X_n = i, X_{n-1} = i_{n-1}, \dots, X_0 = i_0) \\ = P(X_{n+1} = j_1, \dots, X_{n+k} = j_k | X_n = i) \end{aligned}$$

for all $i_0, \dots, i_{n-1}, i, j_1, \dots, j_k$ such that both sides of the equality are defined (Problem 2.1.2). Writing

$$A = \{X_{n+1} = j_1, \dots, X_{n+k} = j_k\}, B = \{X_0 = i_0, \dots, X_{n-1} = i_{n-1}\},$$

the last equality reads $P(A | X_n = i, B) = P(A | X_n = i)$, which is in turn equivalent to

$$P(A \cap B | X_n = i) = P(A | X_n = i)P(B | X_n = i). \quad (1.9)$$

In words: The future at time n and the past at time n are conditionally independent given the present state $X_n = i$. This shows in particular that the Markov property is independent of the direction of time.

Notation We shall abbreviate $P(A | X_0 = i)$ as $P_i(A)$. If μ is a probability distribution on E , then $P_\mu(A) = \sum_{i \in E} \mu(i)P_i(A)$ is the probability of A given that the initial state is distributed according to μ .

2 Markov Recurrences

2.1 A Canonical Representation

Many HMCs receive a natural description in terms of a recurrence equation driven by white noise.

Theorem 2.1. *HMCs Driven by White Noise.*

Let $\{Z_n\}_{n \geq 1}$ be an i.i.d sequence of random variables with values in an arbitrary space F . Let E be a countable space, and $f : E \times F \rightarrow E$ be some function. Let X_0 be a random variable with values in E , independent of $\{Z_n\}_{n \geq 1}$. The recurrence equation

$$X_{n+1} = f(X_n, Z_{n+1}) \quad (2.1)$$

then defines an HMC.

(The phrase *white noise* comes from signal theory and refers to the driving sequence $\{Z_n\}_{n \geq 1}$.)

Proof. Iteration of recurrence (2.1) shows that for all $n \geq 1$, there is a function g_n such that $X_n = g_n(X_0, Z_1, \dots, Z_n)$, and therefore $P(X_{n+1} = j \mid X_n = i, X_{n-1} = i_{n-1}, \dots, X_0 = i_0) = P(f(i, Z_{n+1}) = j \mid X_n = i, X_{n-1} = i_{n-1}, \dots, X_0 = i_0) = P(f(i, Z_{n+1}) = j)$, since the event $\{X_0 = i_0, \dots, X_{n-1} = i_{n-1}, X_n = i\}$ is expressible in terms of X_0, Z_1, \dots, Z_n and is therefore independent of Z_{n+1} . Similarly, $P(X_{n+1} = j \mid X_n = i) = P(f(i, Z_{n+1}) = j)$. We therefore have a Markov chain, and it is homogeneous, since the right-hand side of the last equality does not depend on n . Explicitly

$$p_{ij} = P(f(i, Z_1) = j). \quad (2.2)$$

□

Not all homogeneous Markov chains are naturally described by the model of Theorem 2.1. A slight modification of Theorem 2.1, however, considerably enlarges its scope.

Theorem 2.2.

Let things be as in Theorem 2.1 except for the statistics of X_0, Z_1, Z_2, \dots . Suppose instead that for all $n \geq 0$, Z_{n+1} is conditionally independent of $Z_n, \dots, Z_1, X_{n-1}, \dots, X_0$ given X_n , that is, for all $k, k_1, \dots, k_n \in F, i_0, i_1, \dots, i_{n-1}, i \in E$,

$$\begin{aligned} P(Z_{n+1} = k \mid X_n = i, X_{n-1} = i_{n-1}, \dots, X_0 = i_0, Z_n = k_n, \dots, Z_1 = k_1) \\ = P(Z_{n+1} = k \mid X_n = i), \end{aligned}$$

where the latter quantity is independent of n . Then $\{X_n\}_{n \geq 0}$ is an HMC, with transition matrix \mathbf{P} given by

$$p_{ij} = P(f(i, Z_1) = j \mid X_0 = i).$$

Proof. The proof is analogous to that of Theorem 2.1 (Problem 2.2.6). \square

Remark 2.1. Not all homogeneous Markov chains receive a “natural” description of the type featured in Theorems 2.1 and 2.2, as Example 1.1 (machine replacement) shows. However, for any transition matrix \mathbf{P} on E , there exists a homogeneous Markov chain $\{X_n\}_{n \geq 0}$ with this transition matrix and with a representation such as in Theorem 2.1, namely,

$$X_{n+1} = j \text{ if } Z_{n+1} \in \left[\sum_{k=0}^{j-1} p_{X_n k}, \sum_{k=0}^j p_{X_n k} \right],$$

where $\{Z_n\}_{n \geq 1}$ is i.i.d. uniform on $[0, 1]$. We can apply Theorem 2.1, and check that this HMC has the announced transition matrix. This artificial representation is useful for simulating small Markov chains and can also be helpful for the theory. \diamond

2.2 A Few Famous Examples

The examples below will often be used to illustrate the theory.

Example 2.1. 1-D Random Walk

Let X_0 be a random variable with values in \mathbb{Z} . Let $\{Z_n\}_{n \geq 1}$ be a sequence of i.i.d. random variables, independent of X_0 , taking the values $+1$ or -1 , and with the probability distribution

$$P(Z_n = +1) = p,$$

where $p \in (0, 1)$. The process $\{X_n\}_{n \geq 1}$ defined by

$$X_{n+1} = X_n + Z_{n+1} \tag{2.3}$$

is, in view of Theorem 2.1, an HMC, called the *random walk* on \mathbb{Z} . \diamond

Example 2.2. Repair Shop

During day n , Z_{n+1} machines break down, and they enter the repair shop on day $n + 1$. Every day one machine among those waiting for service is repaired. Therefore, denoting by X_n the number of machines in the shop on day n ,

$$X_{n+1} = (X_n - 1)^+ + Z_{n+1}, \tag{2.4}$$

where $a^+ = \max(a, 0)$. In particular, if $\{Z_n\}_{n \geq 1}$ is an i.i.d. sequence independent of the initial state X_0 , then $\{X_n\}_{n \geq 0}$ is a homogeneous Markov chain. In terms of the probability distribution

$$P(Z_1 = k) = a_k, \quad k \geq 0, \tag{2.5}$$

its transition matrix is

$$\mathbf{P} = \begin{pmatrix} a_0 & a_1 & a_2 & a_3 & \cdots \\ a_0 & a_1 & a_2 & a_3 & \cdots \\ 0 & a_0 & a_1 & a_2 & \cdots \\ 0 & 0 & a_0 & a_1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \tag{2.6}$$

Indeed, from (2.2) and (2.5),

$$p_{ij} = P((i-1)^+ + Z_1 = j) = P(Z_1 = j - (i-1)^+) = a_{j-(i-1)^+}. \quad \diamond$$

Example 2.3. Inventory

A given commodity is stocked in order to satisfy a continuing demand. The aggregated demand between time n and time $n+1$ is Z_{n+1} units, and it is assumed that $\{Z_n\}_{n \geq 1}$ is i.i.d, and independent of the initial value X_0 of the stock. Replenishment of the stock takes place at times $n+0$ (that is to say, immediately after time n) for all $n \geq 1$.




Figure 2.2.1. A sample path of the inventory Markov chain

A popular management strategy is the so-called (s, S) -strategy, where s and S are integers such that $0 < s < S$. Under this inventory policy, if the level of the stock at time n is found not larger than s , then it is brought to level S at time $n+0$. Otherwise, nothing is done. The initial stock X_0 is assumed not greater than S , and therefore $\{X_n\}_{n \geq 1}$ takes its value in $E = \{S, S-1, S-2, \dots\}$. (See Fig. 2.2.1.) Negative values of the stock are allowed, with the interpretation that an unfilled demand is immediately satisfied upon restocking. With the above rules of operation, the evolution of the stock is governed by the dynamic equation

$$X_{n+1} = \begin{cases} X_n - Z_{n+1} & \text{if } s < X_n \leq S, \\ S - Z_{n+1} & \text{if } X_n \leq s. \end{cases} \quad (2.7)$$

In view of (2.7) and Theorem 2.1, $\{X_n\}_{n \geq 1}$ is a homogeneous Markov chain. \(\diamond\)

Example 2.4. Branching Process

In this model $Z_n = (Z_n^{(1)}, Z_n^{(2)}, \dots)$, where the random variables $\{Z_n^{(j)}\}_{n \geq 1, j \geq 1}$ are i.i.d and integer valued. The recurrence equation

$$X_{n+1} = \sum_{k=1}^{X_n} Z_{n+1}^{(k)}, \quad (2.8)$$

with the convention $X_{n+1} = 0$ if $X_n = 0$, receives the following interpretation: X_n is the number of individuals in the n th generation of a given population (humans, particles, etc.). Individual number k of the n th generation gives birth to $Z_{n+1}^{(k)}$ descendants, and this accounts for (2.8).

If the number X_0 of ancestors is independent of $\{Z_n\}_{n \geq 1}$, then according to Theorem 2.1, $\{X_n\}_{n \geq 0}$ is a homogeneous Markov chain, called a branching process because of the genealogical tree that it generates (see Figure 2.2.2). The branching process is also known as the *Galton–Watson process*.




Figure 2.2.2. Sample tree of a branching process

We seek to obtain the probability of extinction of the population. For this we shall introduce g , the common generating function of the variables $Z_n^{(k)}$. The recurrence equation (2.8) provides a recurrent equation for the generating function of the number of individuals in the n th generation,

$$\psi_n(z) = E[z^{X_n}]. \quad (2.9)$$

Indeed,

$$\psi_{n+1}(z) = E[z^{X_{n+1}}] = E[z^{\sum_{k=1}^{X_n} Z^{(k)}}],$$

where $Z_{n+1}^{(k)}$ was denoted by $Z^{(k)}$ for simplicity. Since X_n is a functional of X_0, Z_1, \dots, Z_n , it is independent of Z_{n+1} , and therefore, in the latter equality, X_n is independent of $Z^{(k)}$. From a computation made in Chapter 1, Exercise 5.1,

$$E \left[z^{\sum_{k=1}^{X_n} Z^{(k)}} \right] = \psi_n(g(z)).$$

Therefore,

$$\psi_{n+1}(z) = \psi_n(g(z)).$$

Iterating this equality, we obtain

$$\psi_{n+1}(z) = \psi_0(g^{(n+1)}(z)),$$

where $g^{(n)}$ is the n th iterate of g . If there is only *one ancestor*, then $\psi_0(z) = z$, and therefore $\psi_{n+1}(z) = g^{(n+1)}(z) = g(g^{(n)}(z))$, that is,

$$\psi_{n+1}(z) = g(\psi_n(z)). \quad (2.10)$$

In particular, since $\psi_n(0) = P(X_n = 0)$, we have

$$P(X_{n+1} = 0) = g(P(X_n = 0)), \quad (2.11)$$

an equality that we shall use to discuss extinction. The event \mathcal{E} = "an extinction occurs" is just "at least one generation is empty," that is,

$$\mathcal{E} = \cup_{n=1}^{\infty} \{X_n = 0\}.$$

Also, since $X_n = 0$ implies $X_{n+1} = 0$, the family $\{X_n = 0\}$ is nondecreasing, and by monotone sequential continuity,

$$P(\mathcal{E}) = \lim_{n \rightarrow \infty} P(X_n = 0). \quad (2.12)$$

The generating function g is continuous, and therefore from (2.11) and (2.12), the probability of extinction necessarily satisfies equation

$$P(\mathcal{E}) = g(P(\mathcal{E})). \quad (2.13)$$

Let Z be any of the random variables $Z_n^{(k)}$. It was shown in Chapter 1, Theorem 5.1, that excluding the trivial cases where $P(Z = 0) = 1$ or $P(Z \geq 2) = 0$,

(α) if $E[Z] \leq 1$, the only solution of $x = g(x)$ in $[0, 1]$ is 1, and therefore $P(\mathcal{E}) = 1$. The branching process eventually becomes extinct.

(β) if $E[Z] > 1$, there are two solutions of $x = g(x)$ in $[0, 1]$, 1 and x_0 such that $0 < x_0 < 1$. From the strict convexity of $f : [0, 1] \rightarrow [0, 1]$, it follows that the sequence $y_n = P(X_n = 0)$ that satisfies $y_0 = 0$ and $y_{n+1} = g(y_n)$ converges to x_0 . Therefore, when the mean number of descendants $E[Z]$ is strictly larger than 1, then $P(\mathcal{E}) \in (0, 1)$, and in particular, there is a nonnull probability of extinction.

Here is a small example. Suppose that

$$g(z) = \frac{1}{4} + \frac{1}{4}z + \frac{1}{2}z^2,$$

that is, the probabilities of having 0, 1, or 2 sons are respectively $\frac{1}{4}$, $\frac{1}{4}$, and $\frac{1}{2}$. In particular, $E[Z] = 1.25 > 1$, and $P(\mathcal{E})$ is the solution strictly between 0 and 1 of

$$x = \frac{1}{4} + \frac{1}{4}x + \frac{1}{2}x^2,$$

that is, $P(\mathcal{E}) = \frac{1}{2}$. ◇

Example 2.5. Stochastic Automata

A finite automaton (E, \mathcal{A}, f) can read sequences of letters from a finite alphabet \mathcal{A} written on some infinite tape. It can be in any state of a finite set E , and its evolution is governed by a function $f : E \times \mathcal{A} \rightarrow E$, as follows. When the automaton is in state $i \in E$ and reads letter $a \in \mathcal{A}$, it switches from state i to state $j = f(i, a)$ and then reads on the tape the next letter to the right.




Figure 2.2.3. The automaton: the recognition process and the Markov chain

An automaton can be represented by its transition graph G having for nodes the states of E . There is an oriented edge from the node (state) i to the node j if and only if there exists $a \in \mathcal{A}$ such that $j = f(i, a)$, and this edge then receives label a . If $j = f(i, a_1) = f(i, a_2)$ for $a_1 \neq a_2$, then there are two edges from i to j with labels a_1 and a_2 , or, more economically, one such edge with label (a_1, a_2) . More generally, a given oriented edge can have multiple labels of any order.

Consider, for instance, the automaton with alphabet $\mathcal{A} = \{0, 1\}$ corresponding to the transition graph of Figure 2.2.3a. As the automaton, initialized in state 0, reads the sequence of Figure 2.2.3b from left to right, it passes successively through the states (including the initial state 0)

$$0\ 1\ 0\ 0\ 1\ 2\ 3\ 1\ 0\ 0\ 1\ 2\ 3\ 1\ 2\ 3\ 0\ 1\ 0.$$

Rewriting the sequence of states below the sequence of letters, it appears that the automaton is in state 3 after it has seen three consecutive 1's. This automaton is therefore able to recognize and count such blocks of 1's. However, it does not take into account overlapping blocks (see Fig. 2.2.3b).

If the sequence of letters read by the automaton is $\{Z_n\}_{n \geq 1}$, the sequence of states $\{X_n\}_{n \geq 0}$ is then given by the recurrence equation $X_{n+1} = f(X_n, Z_{n+1})$ and therefore, if $\{Z_n\}_{n \geq 1}$ is i.i.d and independent of the initial state X_0 , then $\{X_n\}_{n \geq 1}$ is, according to Theorem 2.1 an HMC. \diamond

Example 2.6. The Urn of Ehrenfest

This simplified model of diffusion through a porous membrane was proposed in 1907 by the Austrian physicists Tatiana and Paul Ehrenfest to describe in terms of statistical mechanics the exchange of heat between two systems at different temperatures. Their model also considerably helped our understanding of thermodynamic irreversibility (we shall discuss this in Section 2.3 of Chapter 4).

There are N particles that can be either in compartment A or in compartment B . Suppose that at time $n \geq 0$, $X_n = i$ particles are in A . One then chooses a particle at random, and this particle is moved at time $n + 1$ from where it is to the other compartment. Thus, the next state X_{n+1} is either $i - 1$ (the displaced particle was found in compartment A) with probability $\frac{i}{N}$, or $i + 1$ (it was found in B) with probability $\frac{N-i}{N}$.

This model pertains to Theorem 2.2. For all $n \geq 0$,

$$X_{n+1} = X_n + Z_{n+1}, \quad (2.14)$$

where $Z_n \in \{-1, +1\}$ and $P(Z_{n+1} = -1 | X_n = i) = \frac{i}{N}$. The nonzero entries of the transition matrix are therefore

$$p_{i,i+1} = \frac{N-i}{N}, \quad p_{i,i-1} = \frac{i}{N}. \quad (2.15)$$

\diamond

3 First-Step Analysis

3.1 Absorption Probability

Many functionals of homogeneous Markov chains, in particular probabilities of absorption by a closed set (A is called *closed* if $\sum_{j \in A} p_{ij} = 1$ for all $i \in A$) and average times before absorption, can be evaluated by a technique called *first-step analysis*. This technique, which is the motor of most computations in Markov chain theory, is best illustrated by the following example.

Example 3.1. Gambler's Ruin

Two players A and B play “heads or tails”, where heads occur with probability $p \in (0, 1)$, and the successive outcomes form an i.i.d sequence. Calling X_n the fortune in dollars of player A at time n , then $X_{n+1} = X_n + Z_{n+1}$, where $Z_{n+1} = +1$ (resp., -1) with probability p (resp., $q = 1 - p$), and $\{Z_n\}_{n \geq 1}$ is i.i.d. In other words, A bets \$1 on heads at each toss, and B bets \$1 on tails. The respective initial fortunes of A and B are a and b . The game ends when a player is ruined, and therefore the process $\{X_n\}_{n \geq 1}$ is a random walk as described in Example 2.1, except that it is restricted to $E = \{0, \dots, a, a + 1, \dots, a + b = c\}$. The duration of the game is T , the first time n at which $X_n = 0$ or c , and the probability of winning for A is $u(a) = P(X_T = c \mid X_0 = a)$.




Figure 2.3.1. The basic random walk and the gambler's ruin

Instead of computing $u(a)$ alone, first-step analysis computes

$$u(i) = P(X_T = c \mid X_0 = i)$$

for all states $i \in [0, c]$, and for this, it first generates a recurrence equation for the $u(i)$'s by breaking down event “ A wins” according to what can happen after the first step (the first toss) and using the rule of exclusive and exhaustive causes. If $X_0 = i \in [1, c - 1]$, then $X_1 = i + 1$ (resp., $X_1 = i - 1$) with probability p (resp., q), and the probability of ruin of

B starting with A 's initial fortune $i + 1$ (resp., $i - 1$) is $u(i + 1)$ (resp., $u(i - 1)$). Therefore, for $i \in [1, c - 1]$ (see, however, a rigorous proof at the close of the example),

$$u(i) = pu(i + 1) + qu(i - 1), \quad (3.1)$$

with the boundary conditions

$$u(0) = 0, u(c) = 1.$$

The characteristic equation associated with this linear recurrence equation is $pr^2 - r + q = 0$. It has two distinct roots, $r_1 = 1$ and $r_2 = \frac{q}{p}$, if $p \neq q$, and a double root, $r_1 = 1$, if $p = q = \frac{1}{2}$. Therefore, the general solution is $u(i) = \lambda r_1^i + \mu r_2^i = \lambda + \mu \left(\frac{q}{p}\right)^i$ when $p \neq q$, and $u(i) = \lambda r_1^i + \mu i r_1^i = \lambda + \mu i$ when $p = q = \frac{1}{2}$. Taking into account the boundary conditions, one can determine the values of λ and μ . The result is, for $p \neq q$,

$$u(i) = \frac{1 - \left(\frac{q}{p}\right)^i}{1 - \left(\frac{q}{p}\right)^c}, \quad (3.2)$$

and for $p = q = \frac{1}{2}$,

$$u(i) = \frac{i}{c}. \quad (3.3)$$

In the case $p = q = \frac{1}{2}$, the probability $v(i)$ that B wins when the initial fortune of B is $c - i$ is obtained by replacing i by $c - i$ in expression (3.3): $v(i) = \frac{c-i}{c} = 1 - \frac{i}{c}$. One checks that $u(i) + v(i) = 1$, which means in particular that the probability that the game lasts forever is null. The reader is invited to check that the same is true in the case $p \neq q$. \diamond

A justification of the use of the rule of exclusive and exhaustive causes in the obtention of the recurrence equation (3.1) will be given, since we are not exactly in the standard framework of application. The same kind of proof can be performed for every instance of first-step analysis. However, needless to say, one should not go through it every time.

Let $Y = \{Y_n\}_{n \geq 0}$ denote the Markov chain obtained by delaying $X = \{X_n\}_{n \geq 0}$ by one time unit: $Y_n = X_{n+1}$. When $X_0 \in [1, c - 1]$, the events " X is absorbed by 0" and " Y is absorbed by 0" are identical, and therefore

$$\begin{aligned} P(X \text{ is absorbed by } 0, X_1 = i \pm 1, X_0 = i) \\ = P(Y \text{ is absorbed by } 0, X_1 = i \pm 1, X_0 = i). \end{aligned}$$

Since $\{Y_n\}_{n \geq 0}$ and X_0 are independent given X_1 , the right-hand side of the above equality is equal to

$$P(X_0 = i, X_1 = i \pm 1)P(Y \text{ is absorbed by } 0 \mid Y_0 = i \pm 1).$$

The two chains have the same transition matrix, and therefore when they have the same initial state, they have the same distributions. Hence

$$P(Y \text{ is absorbed by } 0 \mid Y_0 = i \pm 1) = u(i \pm 1).$$

The rest of the proof is left for the reader.

Example 3.2. *Cat Eats Mouse Eats Cheese*

A merry mouse moves in a maze. If it is at time n in a room with k adjacent rooms, it will be at time $n + 1$ in one of the k adjacent rooms, choosing one at random, each with probability $\frac{1}{k}$. A fat lazy cat remains all the time in a given room, and a piece of cheese waits for the mouse in another room. (See Fig. 2.3.2.) The cat is not completely lazy: If the mouse enters the room inhabited by the cat, the cat will eat it. What is the probability that the mouse ever gets to eat the cheese when starting from room 1, the cat and the cheese being in rooms 3 and 5, respectively?

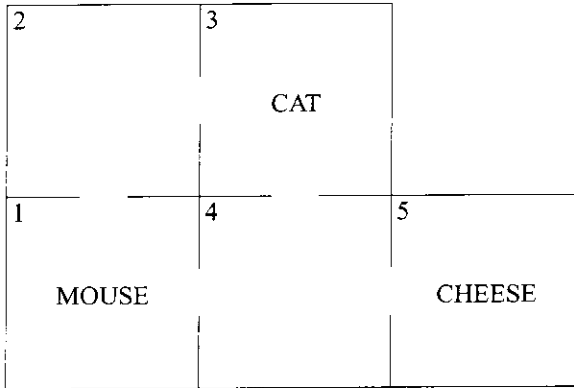


Figure 2.3.2. Maze, mouse, and murder

To apply first-step analysis, call $u(i)$ the probability that the mouse initially in room i reaches the cheese without being murdered by the cat. The boundary conditions

$$u(3) = 0 \text{ and } u(5) = 1$$

are clear. If the mouse is in room 1, its first move will take it to room 2 (resp., 4) with probability $\frac{1}{2}$, and its chance of tasting the cheese will then be $u(2)$ (resp., $u(4)$). Therefore,

$$u(1) = \frac{1}{2}u(2) + \frac{1}{2}u(4).$$

Similarly,

$$\begin{aligned} u(2) &= \frac{1}{2}u(1) + \frac{1}{2}u(3) = \frac{1}{2}u(1), \\ u(4) &= \frac{1}{3}u(1) + \frac{1}{3}u(3) + \frac{1}{3}u(5) = \frac{1}{3}u(1) + \frac{1}{3}. \end{aligned}$$

The solution is $u(1) = \frac{2}{7}$.

◇

Example 3.3. *Tennis*

Ignoring tie-breaks, a game in tennis can be modeled as a Markov chain with the transition graph of Figure 2.3.3a, where p is the probability that the server A wins the point, and $q = 1 - p$.

We want to compute the probability that B wins the game. Clearly, from Figure 2.3.3a, a game can be decomposed into two stages: it first reaches the five upper states of the graph, and it then evolves in the upper states until absorption in “game A ” or “game B .” With an appropriate change of the labels of the upper states, one obtains the transition graph of the Markov chain corresponding to the second stage (Fig. 2.3.3.b).

This is a familiar picture already encountered in the gambler’s ruin problem of Example 3.1. However, a direct first-step analysis gives for b_i , the probability that B wins given that the game starts from upper state i ,

$$b_1 = 1, b_5 = 0$$

and

$$\begin{aligned} b_2 &= q + pb_3, \\ b_3 &= qb_2 + pb_4, \\ b_4 &= qb_3. \end{aligned}$$

For $p \neq q$,

$$(b_1, b_2, b_3, b_4, b_5) = \left(1, q \frac{1-pq}{1-2pq}, \frac{q^2}{1-2pq}, \frac{q^3}{1-2pq}, 0 \right).$$

Starting from state $0 - 0$, the probability that B wins is $\sum_{i=1}^5 p(i)b_i$, where $p(i)$ is the probability that the first upper state reached is i . A simple enumeration of the paths from $0 - 0$ to upper state 1 gives $p(1) = q^4 + q^3pq + q^2pq^2 + qpq^3 + pq^4$, that is,

$$p(1) = q^4(1 + 4p).$$

Similar calculations give

$$p(2) = 4q^3p^2, p(3) = 6p^2q^2, p(4) = 4p^3q^2, p(5) = p^4(1 + 4q).$$

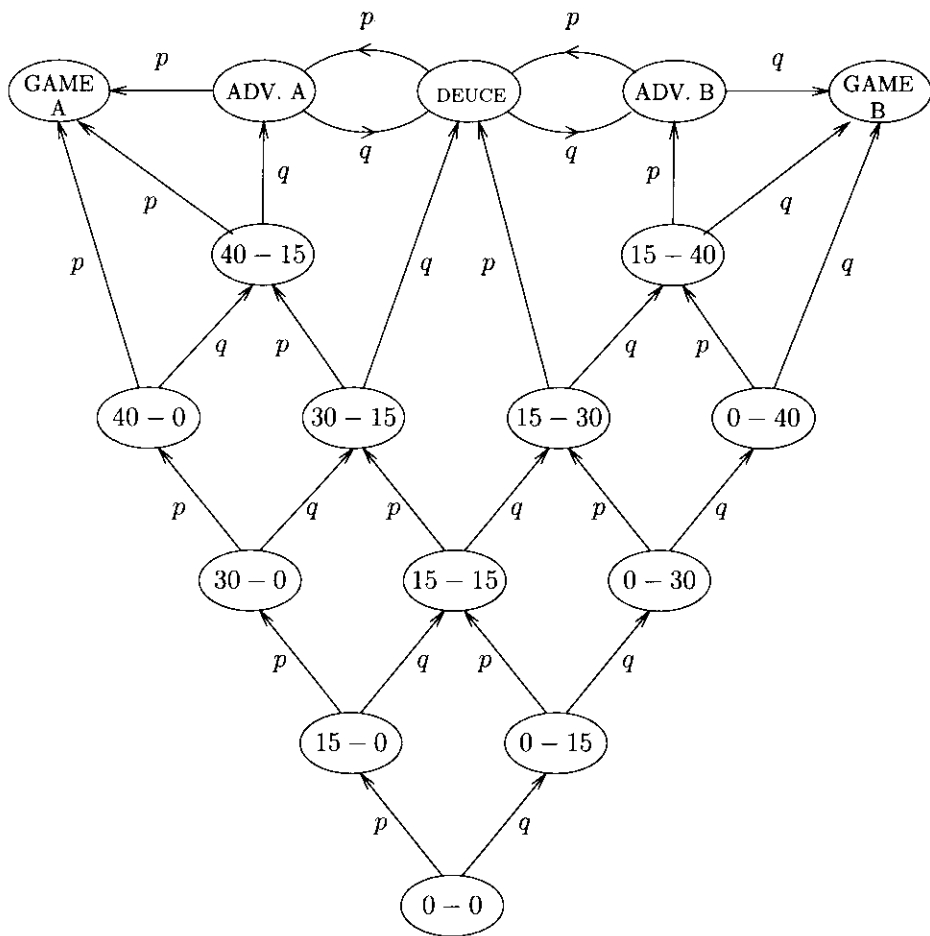
Putting everything together, we find that the probability for B to win is

$$\frac{q^4(1+2p)(1+4p^2)}{1-2pq}.$$


◇

3.2 Mean Time to Absorption

Examples 3.1, 3.2, and 3.3 deal with basically the same problem, that of computing the probability of reaching a state before another state is visited. First-step analysis can also be used to compute average times before absorption.



a



b

Figure 2.3.3. Is this tennis?

Example 3.4. Gambler's Ruin

This example continues Example 3.1. The average duration $m(i) = E[T \mid X_0 = i]$ of the game when the initial fortune of player A is i satisfies the recurrence equation

$$m(i) = 1 + pm(i+1) + qm(i-1) \quad (3.4)$$

for $i \in [1, c-1]$. Indeed, the coin will be tossed at least once, and then with probability p (resp., q) the fortune of player A will be $i+1$ (resp., $i-1$), and therefore $m(i+1)$ (resp., $m(i-1)$) more tosses will be needed on average before one of the players goes broke. The boundary conditions are

$$m(0) = 0, m(c) = 0. \quad (3.5)$$

In order to solve (3.4) with the boundary conditions (3.5), write (3.4) in the form $-1 = p(m(i+1) - m(i)) - q(m(i) - m(i-1))$. Defining

$$y_i = m(i) - m(i-1),$$

we have, for $i \in [1, c-1]$,

$$-1 = py_{i+1} - qy_i \quad (3.6)$$

and

$$m(i) = y_1 + y_2 + \cdots + y_i. \quad (3.7)$$

We now solve (3.6) with $p = q = \frac{1}{2}$. From (3.6),

$$\begin{aligned} -1 &= \frac{1}{2}y_2 - \frac{1}{2}y_1, \\ -1 &= \frac{1}{2}y_3 - \frac{1}{2}y_2, \\ &\vdots \\ -1 &= \frac{1}{2}y_i - \frac{1}{2}y_{i-1}, \end{aligned}$$

and therefore, summing up,

$$-(i-1) = \frac{1}{2}y_i - \frac{1}{2}y_1,$$

that is, for $i \in [1, c]$,

$$y_i = y_1 - 2(i-1).$$

Reporting this expression in (3.7), and observing that $y_1 = m(1)$, we obtain

$$m(i) = im(1) - 2[1 + 2 + \cdots + (i-1)] = im(1) - i(i-1).$$

The boundary condition $m(c) = 0$ gives $cm(1) = c(c-1)$ and therefore, finally,

$$m(i) = i(c-i). \quad (3.8)$$

◇

First-step analysis leads to necessary conditions in the form of a system of linear equations. In the above examples, it turns out that the system in question has a unique solution, a situation that prevails when the state space is finite but that is not the general case with an infinite state space. The issue of uniqueness, and of which solution to choose in case of nonuniqueness, is addressed in Chapter 4, where absorption is studied in more detail.

4 Topology of the Transition Matrix

4.1 Communication

All the properties defined in the present section are *topological* in the sense that they concern only the *naked* transition graph (without the labels).

Definition 4.1. Communication

State j is said to be *accessible* from state i if there exists $M \geq 0$ such that $p_{ij}(M) > 0$. In particular, a state i is always accessible from itself, since $p_{ii}(0) = 1$. States i and j are said to *communicate* if i is accessible from j and j is accessible from i , and this is denoted by $i \leftrightarrow j$.

For $M \geq 1$, $p_{ij}(M) = \sum_{i_1, \dots, i_{M-1}} p_{ii_1} \cdots p_{i_{M-1}j}$, and therefore $p_{ij}(M) > 0$ if and only if there exists at least one path $i, i_1, \dots, i_{M-1}, j$ from i to j such that

$$p_{ii_1} p_{i_1 i_2} \cdots p_{i_{M-1}j} > 0,$$

or, equivalently, if there is an oriented path from i to j in the transition graph G . Clearly,

$$\begin{aligned} i &\leftrightarrow i && \text{(reflexivity),} \\ i &\leftrightarrow j \Rightarrow j &\leftrightarrow i & \text{(symmetry),} \\ i &\leftrightarrow j, j &\leftrightarrow k \Rightarrow i &\leftrightarrow k && \text{(transitivity).} \end{aligned}$$

Therefore, the communication relation (\leftrightarrow) is an equivalence relation, and it generates a partition of the state space E into disjoint equivalence classes called *communication classes*.

Definition 4.2. Closed Sets

A state i such that $p_{ii} = 1$ is called *closed*. More generally, a set C of states such that for all $i \in C$, $\sum_{j \in C} p_{ij} = 1$ is called *closed*.

Example 4.1.

The transition graph of Figure 2.4.1 has 3 communication classes: $\{1, 2, 3, 4\}$, $\{5, 7, 8\}$, and $\{6\}$. State 6 is closed. The communication class $\{5, 7, 8\}$ is not closed, but the set $\{5, 6, 7, 8\}$ is. \diamond

Observe in this example that there may exist oriented edges linking two different communication classes E_k and E_ℓ . However, all the oriented edges between two communication classes have the same orientation (all from E_k to E_ℓ or all from E_ℓ to E_k). Why?




Figure 2.4.1. A transition graph with 3 communication classes

Definition 4.3. *Irreducibility*

If there exists only one communication class, then the chain, its transition matrix, and its transition graph, are said to be *irreducible*.

4.2 Period

Consider the random walk on \mathbb{Z} (Example 2.1). Since $p \in (0, 1)$, it is irreducible. Observe that $E = C_0 + C_1$, where C_0 and C_1 , the set of even and odd relative integers respectively, have the following property. If you start from $i \in C_0$ (resp., C_1), then in one step you can go only to a state $j \in C_1$ (resp., C_0). The chain $\{X_n\}$ passes alternately from one class to the other. In this sense, the chain has a periodic behavior, corresponding to the period 2. More generally, we have the following.

Theorem 4.1. *Cyclic Structure*

For any *irreducible* Markov chain, one can find a *unique partition* of E into d classes C_0, C_1, \dots, C_{d-1} such that for all $k, i \in C_k$,

$$\sum_{j \in C_{k+1}} p_{ij} = 1,$$

where by convention $C_d = C_0$, and where d is maximal (that is, there is no other such partition $C'_0, C'_1, \dots, C'_{d'-1}$ with $d' > d$).

Proof. A direct consequence of Theorem 4.3 below. □

The number $d \geq 1$ is called the *period* of the chain (resp., of the transition matrix, of the transition graph). The classes C_0, C_1, \dots, C_{d-1} are called the *cyclic classes*.

The chain therefore moves from one class to the other at each transition, and this cyclically, as shown in Figure 2.4.2.




Figure 2.4.2. Cycles

Example 4.2.

The chain with the transition graph depicted in Figure 2.4.3 is irreducible and has period $d = 3$, with the cyclic classes $C_0 = \{1, 2\}$, $C_1 = \{4, 7\}$, $C_3 = \{3, 5, 6\}$. \diamond




Figure 2.4.3. An irreducible transition graph with period 3

Consider now an irreducible chain of period d with the cyclic classes C_0, C_1, \dots, C_d . Renumbering the states of E if necessary, the transition matrix has the block structure below (where $d = 4$, to be explicit),

$$\mathbf{P} = \begin{matrix} & C_0 & C_1 & C_2 & C_3 \\ \begin{matrix} C_0 \\ C_1 \\ C_2 \\ C_3 \end{matrix} & \begin{pmatrix} 0 & A_0 & 0 & \\ 0 & 0 & A_1 & 0 \\ 0 & 0 & 0 & A_2 \\ A_3 & 0 & 0 & 0 \end{pmatrix} & , \end{matrix}$$

and therefore \mathbf{P}^2 , \mathbf{P}^3 , and \mathbf{P}^4 also have a block structure corresponding to C_0, C_1, C_2, C_3 :

$$\mathbf{P}^2 = \begin{pmatrix} 0 & 0 & B_0 & 0 \\ 0 & 0 & 0 & B_1 \\ B_2 & 0 & 0 & 0 \\ 0 & B_3 & 0 & 0 \end{pmatrix}, \quad \mathbf{P}^3 = \begin{pmatrix} 0 & 0 & 0 & D_0 \\ D_1 & 0 & 0 & 0 \\ 0 & D_2 & 0 & 0 \\ 0 & 0 & D_3 & 0 \end{pmatrix}, \quad \mathbf{P}^4 = \begin{pmatrix} E_0 & 0 & 0 & 0 \\ 0 & E_1 & 0 & 0 \\ 0 & 0 & E_2 & 0 \\ 0 & 0 & 0 & E_3 \end{pmatrix}.$$

We observe two phenomena: block-shifting and the block-diagonal form of \mathbf{P}^d . This is, of course, general: \mathbf{P}^d has a block-diagonal form corresponding to the cyclic classes C_0, C_1, \dots, C_{d-1} :

$$\mathbf{P}^d = \begin{matrix} & C_0 & C_1 & \cdots & C_{d-1} \\ \begin{matrix} C_0 \\ C_1 \\ \vdots \\ C_{d-1} \end{matrix} & \begin{pmatrix} E_0 & & & \\ & E_1 & & \\ & & \ddots & \\ 0 & & & E_{d-1} \end{pmatrix} & & & \end{matrix}. \quad (4.1)$$

The d -step transition matrix \mathbf{P}^d is also a stochastic matrix, and obviously the \mathbf{P} -cyclic classes C_0, C_1, \dots, C_{d-1} are all in different \mathbf{P}^d -communication classes, as the diagonal block structure shows.

The question is this: Is there, in C_0 for instance, more than one \mathbf{P}^d -communication class? The answer is no, and therefore matrix E_0 in (4.1) is an irreducible stochastic matrix. To see this, take two different states $i, j \in C_0$. Since they \mathbf{P} -communicate (\mathbf{P} was assumed irreducible), there exist $m > 0$ and $n > 0$ such that $p_{ij}(m) > 0$ and $p_{ji}(n) > 0$. But since \mathbf{P} has period d , necessarily $m = Md, n = Nd$ for some $M > 0$ and $N > 0$. Therefore, $p_{ij}(Md) > 0$ and $p_{ji}(Nd) > 0$. But $p_{ij}(Md)$ is the (i, j) term of $(\mathbf{P}^d)^M$, and similarly for $p_{ji}(Nd)$. We therefore have proven that i and j \mathbf{P}^d -communicate.

Also (Problem 2.4.7), the restriction of \mathbf{P}^d to any cyclic class is aperiodic.

For an arbitrary transition matrix, not necessarily irreducible, the formal notion of *period* is the following.

Definition 4.4. *Arithmetic Definition of Period*

The period d_i of state $i \in E$ is, by definition,

$$d_i = \gcd\{n \geq 1; p_{ii}(n) > 0\}, \quad (4.2)$$

with the convention $d_i = +\infty$ if there is no $n \geq 1$ with $p_{ii}(n) > 0$. If $d_i = 1$, the state i is called *aperiodic*.

Theorem 4.2. *Period Is a Class Property*

If states i and j communicate they have the same period.

Proof. As i and j communicate, there exist integers N and M such that $p_{ij}(M) > 0$ and $p_{ji}(N) > 0$. For any $k \geq 1$,

$$p_{ii}(M + nk + N) \geq p_{ij}(M)(p_{jj}(k))^n p_{ji}(N)$$

(indeed, the path $X_0 = i, X_M = j, X_{M+k} = j, \dots, X_{M+nk} = j, X_{M+nk+N} = i$ is just one way of going from i to i in $M + nk + N$ steps).

Therefore, for any $k \geq 1$ such that $p_{jj}(k) > 0$, we have $p_{ii}(M + nk + N) > 0$ for all $n \geq 1$. Therefore, d_i divides $M + nk + N$ for all $n \geq 1$, and in particular, d_i divides k . We have therefore shown that d_i divides all k such that $p_{jj}(k) > 0$, and in particular, d_i divides d_j . By symmetry, d_j divides d_i , and therefore, finally, $d_i = d_j$. \square

We can therefore speak of the period of a communication class or of an irreducible chain.

The important result concerning periodicity is the following.

Theorem 4.3. Lattice Theorem

Let \mathbf{P} be an irreducible stochastic matrix with period d . Then for all states i, j there exist $m \geq 0$ and $n_0 \geq 0$ (m and n_0 possibly depending on i, j) such that

$$p_{ij}(m + nd) > 0, \forall n \geq n_0. \quad (4.3)$$

Proof. First observe that it suffices to prove this for $i = j$. Indeed, there exists m such that $p_{ij}(m) > 0$, because j is accessible from i , the chain being irreducible, and therefore, if for some $n_0 \geq 0$ we have $p_{jj}(nd) > 0$ for all $n \geq n_0$, then $p_{ij}(m + nd) \geq p_{ij}(m)p_{jj}(nd) > 0$ for all $n \geq n_0$. The gcd of the set $A = \{k \geq 1; p_{jj}(k) > 0\}$ is d , and A is closed under addition. The set A therefore contains all but a finite number of the positive multiples of d (see Theorem 1.1 of the Appendix). In other words, there exists n_0 such that $n > n_0$ implies $p_{jj}(nd) > 0$. \square

5 Steady State

5.1 Stationarity

We now introduce the central notion of the stability theory of discrete-time HMCs.

Definition 5.1. Stationary Distribution

A probability distribution π satisfying

$$\pi^T = \pi^T \mathbf{P} \quad (5.1)$$

is called a *stationary distribution* of the transition matrix \mathbf{P} , or of the corresponding HMC.

The *global balance equation* (5.1) says that for all states i ,

$$\pi(i) = \sum_{j \in E} \pi(j)p_{ji}. \quad (5.2)$$

Iteration of (5.1) gives $\pi^T = \pi^T \mathbf{P}^n$ for all $n \geq 0$, and therefore, in view of (1.7), if the initial distribution $\nu = \pi$, then $\nu_n = \pi$ for all $n \geq 0$. Thus, if a chain is started with a stationary distribution, it keeps the same distribution forever. But there is more, because then,

$$P(X_n = i_0, X_{n+1} = i_1, \dots, X_{n+k} = i_k) = \pi(i_0)p_{i_0 i_1} \cdots p_{i_{k-1} i_k} \quad (5.3)$$

does not depend on n . In this sense the chain is *stationary*. One also says that the chain is in a *stationary regime*, or in *equilibrium*, or in *steady state*. In summary:

Theorem 5.1. Steady State

A chain started with a stationary distribution is stationary.

Remark 5.1. The balance equation $\pi^T \mathbf{P} = \pi^T$, together with the requirement that π be a probability vector, i.e., $\pi^T \mathbf{1} = 1$ (where $\mathbf{1}$ is a column vector with all its entries equal to 1), constitute when E is finite, $|E| + 1$ equations for $|E|$ unknown variables. One of the $|E|$ equations in $\pi^T \mathbf{P} = \pi^T$ is superfluous given the constraint $\pi^T \mathbf{1} = 1$. Indeed, summing up all equalities of $\pi^T \mathbf{P} = \pi^T$ yields the equality $\pi^T \mathbf{P} \mathbf{1} = \pi^T \mathbf{1}$, that is, $\pi^T \mathbf{1} = 1$. \diamond

5.2 Examples

Example 5.1. Two-State Markov Chain

Take $E = \{1, 2\}$ and define the transition matrix

$$\mathbf{P} = \frac{1}{2} \begin{pmatrix} 1 & 2 \\ 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix},$$

where $\alpha, \beta \in (0, 1)$. The global balance equations are

$$\begin{aligned} \pi(1) &= \pi(1)(1 - \alpha) + \pi(2)\beta, \\ \pi(2) &= \pi(1)\alpha + \pi(2)(1 - \beta). \end{aligned}$$

This is a dependent system which reduces to the single equation $\pi(1)\alpha = \pi(2)\beta$, to which must be added $\pi(1) + \pi(2) = 1$ expressing that π is a probability vector. We obtain

$$\pi(1) = \frac{\beta}{\alpha + \beta}, \quad \pi(2) = \frac{\alpha}{\alpha + \beta}. \quad \diamond$$

Example 5.2. The Urn of Ehrenfest

The corresponding HMC was described in Example 2.6. The global balance equations are, for $i \in [1, N - 1]$,

$$\pi(i) = \pi(i - 1) \left(1 - \frac{i - 1}{N} \right) + \pi(i + 1) \frac{i + 1}{N}$$

and, for the boundary states,

$$\pi(0) = \pi(1)\frac{1}{N}, \quad \pi(N) = \pi(N-1)\frac{1}{N}.$$

Leaving $\pi(0)$ undetermined, one can solve the balance equations for $i = 0, 1, \dots, N$ successively, to obtain

$$\pi(i) = \pi(0)\binom{N}{i}.$$

The value of $\pi(0)$ is then determined by writing that π is a probability vector:

$$1 = \sum_{i=0}^N \pi(i) = \pi(0) \sum_{i=0}^N \binom{N}{i} = \pi(0)2^N.$$

This gives for π the binomial distribution of size N and parameter $\frac{1}{2}$:

$$\pi(i) = \frac{1}{2^N} \binom{N}{i}. \quad (5.4)$$

This is the distribution one would obtain by placing independently each particle in the compartments, with probability $\frac{1}{2}$ for each compartment. \diamond

Example 5.3. *Symmetric Random Walk*

A symmetric random walk on \mathbb{Z} cannot have a stationary distribution. Indeed, the solution of the balance equation

$$\pi(i) = \frac{1}{2}\pi(i-1) + \frac{1}{2}\pi(i+1)$$

for $i \geq 0$, with initial data $\pi(0)$ and $\pi(1)$, is

$$\pi(i) = \pi(0) + (\pi(1) - \pi(0))i.$$

Since $\pi(i) \in [0, 1]$, necessarily $\pi(1) - \pi(0) = 0$. Therefore, $\pi(i)$ is a constant, necessarily 0 because the total mass of π is finite. Thus for all $i \geq 0$, and therefore, in view of the global balance equation, for all i , $\pi(i) = 0$, a contradiction if we want π to be a probability distribution. \diamond

Example 5.4. *Stationary Distributions May Be Many*

Take the identity as transition matrix. Then any probability distribution on the state space is a stationary distribution. \diamond

Recurrence equations can be used to obtain the stationary distribution when the latter exists and is unique. Generating functions sometimes usefully exploit the dynamics.

Example 5.5. Repair Shop

This example continues Example 2.2. For any complex number z with modulus not larger than 1, it follows from the recurrence equation (2.4) that

$$\begin{aligned} z^{X_{n+1}+1} &= \left(z^{(X_n-1)^+ + 1} \right) z^{Z_{n+1}} \\ &= \left(z^{X_n} \mathbf{1}_{\{X_n > 0\}} + z \mathbf{1}_{\{X_n = 0\}} \right) z^{Z_{n+1}} \\ &= \left(z^{X_n} - \mathbf{1}_{\{X_n = 0\}} + z \mathbf{1}_{\{X_n = 0\}} \right) z^{Z_{n+1}}, \end{aligned}$$

and therefore

$$z z^{X_{n+1}} - z^{X_n} z^{Z_{n+1}} = (z - 1) \mathbf{1}_{\{X_n = 0\}} z^{Z_{n+1}}.$$

From the independence of X_n and Z_{n+1} , $E[z^{X_n} z^{Z_{n+1}}] = E[z^{X_n}] g_Z(z)$, where $g_Z(z)$ is the generating function of Z_{n+1} , and $E[\mathbf{1}_{\{X_n = 0\}} z^{Z_{n+1}}] = \pi(0) g_Z(z)$, where $\pi(0) = P(X_n = 0)$. Therefore,

$$z E[z^{X_{n+1}}] - g_Z(z) E[z^{X_n}] = (z - 1) \pi(0) g_Z(z).$$

But in steady state, $E[z^{X_{n+1}}] = E[z^{X_n}] = g_X(z)$, and therefore

$$g_X(z)(z - g_Z(z)) = \pi(0)(z - 1)g_Z(z). \quad (5.5)$$

This gives the generating function $g_X(z) = \sum_{i=0}^{\infty} \pi(i) z^i$, as long as $\pi(0)$ is available. To obtain $\pi(0)$, differentiate (5.5):

$$g'_X(z)(z - g_Z(z)) + g_X(z)(1 - g'_Z(z)) = \pi(0)(g_Z(z) + (z - 1)g'_Z(z)),$$

and let $z = 1$, to obtain, taking into account the equalities $g_X(1) = g_Z(1) = 1$ and $g'_Z(1) = E[Z]$,

$$\pi(0) = 1 - E[Z]. \quad (5.6)$$

Since $\pi(0)$ must be nonnegative, this immediately gives the necessary condition $E[Z] \leq 1$. Actually, one must have, if the trivial case $Z_{n+1} \equiv 1$ is excluded,

$$E[Z] < 1. \quad (5.7)$$

Indeed, if $E[Z] = 1$, implying $\pi(0) = 0$, it follows from (5.5) that

$$g_X(x)(x - g_Z(x)) = 0$$

for all $x \in [0, 1]$. But excluding the case $Z_{n+1} \equiv 1$ (that is, $g_Z(x) \equiv x$), the equation $x - g_Z(x) = 0$ has only $x = 1$ for a solution when $g'_Z(1) = E[Z] \leq 1$ (see Chapter 1, Theorem 5.1). Therefore, $g_X(x) \equiv 0$ for all $x \in [0, 1)$, and consequently $g_X(z) \equiv 0$ on $\{|z| < 1\}$. This leads to a contradiction, since the generating function of an integer-valued random variable cannot be identically null.

We shall prove later that $E[Z] < 1$ is also a sufficient condition for the existence of a steady state. For the time being, we learn from (5.5) and (5.6) that, if the stationary distribution exists, then its generating function is given by the formula

$$\sum_{i=0}^{\infty} \pi(i) z^i = (1 - E[Z]) \frac{(z - 1)g_Z(z)}{z - g_Z(z)}. \quad (5.8)$$

◇

The unknown $\pi(0)$ is obtained by $\sum_{i=0}^N \pi(i) = 1$, that is,

$$\pi(0) \left\{ 1 + \frac{1}{q_1} + \frac{p_1}{q_1 q_2} + \cdots + \frac{p_1 p_2 \cdots p_{N-1}}{q_1 q_2 \cdots q_{N-1} q_N} \right\} = 1. \quad (5.10)$$

◇

Example 5.7. Birth and Death with One Reflecting Barrier

The model is the same as above, except that the state space is $E = \mathbb{N}$, and therefore the upper barrier is at infinity. The same computations as above lead to the expression (5.9) for the general solution of $\pi^T \mathbf{P} = \pi^T$, which depends on the initial condition $\pi(0)$. For this solution to be a probability, we must have $\pi(0) > 0$. Also, writing $\sum_{i=1}^{\infty} \pi(i) = 1$,

$$\pi(0) \left\{ 1 + \frac{1}{q_1} + \sum_{j=1}^{\infty} \frac{p_1 p_2 \cdots p_j}{q_1 q_2 \cdots q_{j+1}} \right\} = 1. \quad (5.11)$$

Thus a stationary distribution exists if and only if

$$\sum_{j=1}^{\infty} \frac{p_1 p_2 \cdots p_j}{q_1 q_2 \cdots q_{j+1}} < \infty. \quad (5.12)$$

In this case $\pi(i)$ is given by the expressions in (5.9), where $\pi(0)$ is determined by (5.11). ◇

6 Time Reversal

6.1 Reversed Chain

The notions of time-reversal and time-reversibility are very productive, in particular in the theory of Markov chains, and especially in Monte Carlo simulation (Chapter 7) and queuing theory (Chapter 9).

Let $\{X_n\}_{n \geq 0}$ be an HMC with transition matrix \mathbf{P} and admitting a stationary distribution π such that

$$\pi(i) > 0 \quad (6.1)$$

for all states i . Define the matrix \mathbf{Q} , indexed by E , by

$$\pi(i) q_{ij} = \pi(j) p_{ji}. \quad (6.2)$$

This matrix is stochastic, since

$$\sum_{j \in E} q_{ij} = \sum_{j \in E} \frac{\pi(j)}{\pi(i)} p_{ji} = \frac{1}{\pi(i)} \sum_{j \in E} \pi(j) p_{ji} = \frac{\pi(i)}{\pi(i)} = 1,$$

where the third equality uses the balance equations. Its interpretation is the following: Suppose that the initial distribution of $\{X_n\}$ is π , in which case for all $n \geq 0$, all $i \in E$,

$$P(X_n = i) = \pi(i). \quad (6.3)$$

Then, from Bayes's retrodiction formula,

$$P(X_n = j | X_{n+1} = i) = \frac{P(X_{n+1} = i | X_n = j)P(X_n = j)}{P(X_{n+1} = i)},$$

that is, in view of (6.2) and (6.3),

$$P(X_n = j | X_{n+1} = i) = q_{ji}. \quad (6.4)$$

We see that \mathbf{Q} is the transition matrix of the initial chain when time is reversed.

The following is a very simple observation that will be promoted to the rank of a theorem in view of its usefulness and also for the sake of easy reference.

Theorem 6.1. *Reversal Test*

Let \mathbf{P} be a stochastic matrix indexed by a countable set E , and let π be a probability distribution on E . Let \mathbf{Q} be a stochastic matrix indexed by E such that for all $i, j \in E$,

$$\pi(i)q_{ij} = \pi(j)p_{ji}. \quad (6.5)$$

Then π is a stationary distribution of \mathbf{P} .

Proof. For fixed $i \in E$, sum equalities (6.5) with respect to $j \in E$ to obtain

$$\sum_{j \in E} \pi(i)q_{ij} = \sum_{j \in E} \pi(j)p_{ji}.$$

But the left-hand side is equal to $\pi(i) \sum_{j \in E} q_{ij} = \pi(i)$, and therefore, for all $i \in E$,

$$\pi(i) = \sum_{j \in E} \pi(j)p_{ji}. \quad \square$$

Example 6.1. *Extension of a Stationary Chain to Negative Times*

Time reversal can also be used to extend to negative times a chain $\{X_n\}_{n \geq 0}$ in steady state corresponding to a stationary distribution π such that $\pi(i) > 0$ for all $i \in E$. See Problem 2.6.2. \diamond

6.2 Time Reversibility

Definition 6.1. *Reversible Chain*

One calls *reversible* a stationary Markov chain with initial distribution π (a stationary distribution) assumed positive if for all $i, j \in E$,

$$\pi(i)p_{ij} = \pi(j)p_{ji}. \quad (6.6)$$

In this case, $q_{ij} = p_{ij}$, and therefore the chain and the time-reversed chain are statistically the same, since the distribution of a homogeneous Markov chain is entirely determined by its initial distribution and its transition matrix (Theorem 1.1). Equations (6.6) are called the *detailed balance equations*. The following is an immediate corollary of Theorem 6.1.

Corollary 6.1. *Detailed Balance Test*

Let \mathbf{P} be a transition matrix on the countable state space E , and let π be some probability distribution on E . If for all $i, j \in E$, the detailed balance equations (6.6) are satisfied, then π is a stationary distribution of \mathbf{P} .

Example 6.2. *The Urn of Ehrenfest*

This example continues Examples 2.6 and 5.2. Recall that we obtained the expression

$$\pi(i) = \frac{1}{2^N} \binom{N}{i}$$

for the stationary distribution. Checking the detailed balance equations

$$\pi(i)p_{i,i+1} = \pi(i+1)p_{i+1,i}$$

is immediate. ◇

Example 6.3. *The Generalized Mouse*

The reason for the title of this example is that it is an abstract form of the motion of a mouse in a maze; see Example 3.2. However, the professionals call this a *random walk on a graph*. Consider a finite nonoriented graph and call E the set of vertices, or nodes, of this graph. Call d_i the number of edges “adjacent” to node i . Transform this graph into an oriented graph by splitting each edge into two oriented edges of opposite directions, and make it a transition graph by associating to the oriented edge from i to j the transition probability $\frac{1}{d_i}$ (see Fig. 2.6.1).




Figure 2.6.1. A random walk on a graph

It will be assumed, as is the case in Figure 2.6.1, that $d_i > 0$ for all states i . A stationary distribution (in fact, *the* stationary distribution, as we shall see later, in Chapter 3) is given by

$$\pi(i) = \frac{d_i}{\sum_{j \in E} d_j}.$$

For this, we can use Corollary 6.1, making the insider's guess that the chain is reversible. We just have to check that

$$\pi(i) \frac{1}{d_i} = \pi(j) \frac{1}{d_j}.$$

Hence we have obtained the stationary distribution and proved the reversibility of the chain. \diamond

Example 6.4. Birth and Death

One verifies that for both Examples 5.6 and 5.7, when the stationary distribution π exists, the detailed balance equations $\pi(i)p_i = \pi(i+1)q_{i+1}$ hold for all $i \in E$. \diamond

7 Regeneration

7.1 Strong Markov Property

Definition 7.1. Stopping Times

A *stopping time* with respect to a stochastic process $\{X_n\}_{n \geq 0}$ is, by definition, a random variable τ taking its values in $\mathbb{N} \cup \{+\infty\}$ and such that for all integers $m \geq 0$, the event $\{\tau = m\}$ can be expressed in terms of X_0, X_1, \dots, X_m .

The latter property is symbolized by the notation

$$\{\tau = m\} \in \mathcal{X}_0^m. \tag{7.1}$$

When the state space is countable, (7.1) means that

$$1_{\{\tau=m\}} = \psi_m(X_0, \dots, X_m),$$

for some function ψ_m with values in $\{0, 1\}$.

Example 7.1. Return Times

In the theory of Markov chains, a typical and most important stopping time is the *return time* to state $i \in E$,

$$T_i = \inf \{n \geq 1; X_n = i\}, \tag{7.2}$$

where $T_i = \infty$ if $X_n \neq i$ for all $n \geq 1$. It is indeed a stopping time. Do a direct proof, or wait for Example 7.5. \diamond

Observe that $T_i \geq 1$, and in particular, $X_0 = i$ does *not* imply $T_i = 0$. This is why T_i is called the *return* time to i , and not the *hitting* time of i . The latter is $S_i = T_i$ if $X_0 \neq i$, and $S_i = 0$ if $X_0 = i$. It is also a stopping time.

Example 7.2. Deterministic Times

A constant time is a stopping time (check this). \diamond

Example 7.3. Delayed Stopping Times

If τ is a stopping time and n_0 a nonnegative deterministic time, then $\tau + n_0$ is a stopping time. Indeed, $\{\tau + n_0 = m\} \equiv \{\tau = m - n_0\}$ is expressible in terms of $X_0, X_1, \dots, X_{m-n_0}$. \diamond

For a given stopping time τ , one can decide whether $\tau = m$ just by observing X_0, X_1, \dots, X_m . This is why stopping times are said to be *nonanticipative*.

Example 7.4. Counterexample

The random time

$$\tau = \inf \{n \geq 0; X_{n+1} = i\},$$

where $\tau = \infty$ if $X_{n+1} \neq i$ for all $n \geq 0$, is anticipative because $\{\tau = m\} = \{X_1 \neq i, \dots, X_m \neq i, X_{m+1} = i\}$ for all $m \geq 0$. Knowledge of this random time provides information about the value of the process just after it. It is *not* a stopping time. \diamond

Example 7.5. Successive Returns

Let $\tau_1 = T_i, \tau_2, \dots$ be the successive return times to state i . If there are only r returns to state i , let $\tau_{r+1} = \tau_{r+2} = \dots = \infty$. These random times are stopping times with respect to $\{X_n\}_{n \geq 0}$, since for any $m \geq 1$,

$$\{\tau_k = m\} = \left\{ \sum_{n=1}^m 1_{\{X_n=i\}} = k, X_m = i \right\}$$

is indeed expressible in terms of X_0, \dots, X_m . \diamond

Let τ be a random time taking its values in $\mathbb{N} \cup \{+\infty\}$, and let $\{X_n\}_{n \geq 0}$ be a stochastic process with values in the countable set E . In order to define X_τ when $\tau = \infty$, one must decide how to define X_∞ . This is done by taking some arbitrary element Δ not in E , and setting

$$X_\infty = \Delta.$$

By definition, the “process $\{X_n\}$ after τ ” is the stochastic process

$$\{X_{n+\tau}\}_{n \geq 0}.$$

The “process $\{X_n\}$ before τ ,” or the “process $\{X_n\}$ stopped at time τ ,” is the process

$$\{X_{n \wedge \tau}\}_{n \geq 0},$$

which freezes at time τ at the value X_τ .

The main result of the present section roughly says that the Markov property, i.e. the independence of past and future given the present state, extends to the situation where the present time is a stopping time. More precisely:

Theorem 7.1. *Strong Markov Property*

Let $\{X_n\}_{n \geq 0}$ be an HMC with countable state space E and transition matrix \mathbf{P} . Let τ be a stopping time with respect to this chain. Then for any state $i \in E$, given that $X_\tau = i$ (in particular, $\tau < \infty$, since $i \neq \Delta$), the following hold:

(α) The process after τ and the process before τ are independent.

(β) The process after τ is an HMC with transition matrix \mathbf{P} .

Proof. (α) One must show that for all times $k \geq 1, n \geq 0$, and all states $i_0, \dots, i_n, i, j_1, \dots, j_k$,

$$\begin{aligned} P(X_{\tau+1} = j_1, \dots, X_{\tau+k} = j_k \mid X_\tau = i, X_{\tau \wedge 0} = i_0, \dots, X_{\tau \wedge n} = i_n) \\ = P(X_{\tau+1} = j_1, \dots, X_{\tau+k} = j_k \mid X_\tau = i). \end{aligned} \quad (7.3)$$

We shall prove a simplified version of the above equality, namely

$$P(X_{\tau+k} = j \mid X_\tau = i, X_{\tau \wedge n} = i_n) = P(X_{\tau+k} = j \mid X_\tau = i). \quad (7.4)$$

(The general case would be obtained by similar arguments.) The left-hand side of the above equality is equal to

$$\frac{P(X_{\tau+k} = j, X_\tau = i, X_{\tau \wedge n} = i_n)}{P(X_\tau = i, X_{\tau \wedge n} = i_n)}.$$

The numerator of the above expression can be developed as

$$\sum_{r \geq 0} P(\tau = r, X_{r+k} = j, X_r = i, X_{r \wedge n} = i_n). \quad (7.5)$$

But $P(\tau = r, X_{r+k} = j, X_r = i, X_{r \wedge n} = i_n) = P(X_{r+k} = j \mid X_r = i, X_{r \wedge n} = i_n, \tau = r) P(\tau = r, X_{r \wedge n} = i_n, X_r = i)$, and since $r \wedge n \leq r$ and $\{\tau = r\} \in X'_0$, the event $B = \{X_{r \wedge n} = i_n, \tau = r\}$ is in X'_0 . Therefore, by the Markov property, $P(X_{r+k} = j \mid X_r = i, X_{r \wedge n} = i_n, \tau = r) = P(X_{r+k} = j \mid X_r = i) = p_{ij}(k)$. Finally, expression (7.5) reduces to

$$\sum_{r \geq 0} p_{ij}(k) P(\tau = r, X_{r \wedge n} = i_n, X_r = i) = p_{ij}(k) P(X_{\tau=i}, X_{\tau \wedge n} = i_n).$$

Therefore, the left-hand side of (7.4) is just $p_{ij}(k)$. Similar computations show that the right-hand side of (7.4) is also $p_{ij}(k)$, so that (α) is proven.

(β) We must show that for all states $i, j, k, i_{n-1}, \dots, i_1$,

$$\begin{aligned} P(X_{\tau+n+1} = k \mid X_{\tau+n} = j, X_{\tau+n-1} = i_{n-1}, \dots, X_\tau = i) \\ = P(X_{\tau+n+1} = k \mid X_{\tau+n} = j) = p_{jk}. \end{aligned}$$

But the first equality follows from the fact proven in (α) that for the stopping time $\tau' = \tau + n$, the processes before and after τ' are independent given $X_{\tau'} = j$. The second equality is obtained by calculations similar to those in the proof of (α). \square

7.2 Regenerative Cycles

Let the number of visits to state i strictly after time 0 be denoted by

$$N_i = \sum_{n \geq 1} 1_{\{X_n = i\}}. \quad (7.6)$$

Theorem 7.2. Visits to a State

The distribution of N_i given $X_0 = j$ is

$$P_j(N_i = r) = \begin{cases} f_{ji} f_{ii}^{r-1} (1 - f_{ii}) & \text{for } r \geq 1, \\ 1 - f_{ji} & \text{for } r = 0, \end{cases} \quad (7.7)$$

where

$$f_{ji} = P_j(T_i < \infty) \quad (7.8)$$

and T_i is the return time to i .

Proof. For $r = 0$, this is just the definition of f_{ji} . Now let $r \geq 1$, and assume (7.7) to be true for all $k \in [1, r]$. In particular,

$$P_j(N_i > r) = 1 - \sum_{k=0}^r P_j(N_i = k) = f_{ji} f_{ii}^r.$$

Denoting by τ_r the r th return time to state i ,

$$\begin{aligned} P_j(N_i = r + 1) &= P_j(N_i = r + 1, X_{\tau_{r-1}} = i) \\ &= P_j(\tau_{r+2} - \tau_{r+1} = \infty, X_{\tau_{r+1}} = i) \\ &= P_j(\tau_{r+2} - \tau_{r+1} = \infty \mid X_{\tau_{r+1}} = i) P_j(X_{\tau_{r+1}} = i). \end{aligned}$$

But

$$\begin{aligned} P_j(\tau_{r+2} - \tau_{r+1} = \infty \mid X_{\tau_{r+1}} = i) &= P(\tau_{r+2} - \tau_{r+1} = \infty \mid X_{\tau_{r+1}} = i, X_0 = j) \\ &= P(\tau_{r+2} - \tau_{r+1} = \infty \mid X_{\tau_{r+1}} = i) \end{aligned}$$

by the strong Markov property. Since $\tau_{r+2} - \tau_{r+1}$ is the return time to i of the process after τ_{r+1} , the strong Markov property gives

$$P(\tau_{r+2} - \tau_{r+1} = \infty \mid X_{\tau_{r+1}} = i) = P(T_i = \infty \mid X_0 = i).$$

Also,

$$P_j(X_{\tau_{r-1}} = i) = P_j(N_i > r)$$

(if $N_i \leq r$, then $X_{\tau_{r+1}} = X_\infty = \Delta \notin E$). Therefore,

$$P_j(N_i = r + 1) = P_i(T_i = \infty) P_j(N_i > r) = (1 - f_{ii}) f_{ji} f_{ii}^r.$$

The result is therefore proven, by induction. \square

The distribution of N_i given $X_0 = j$ and given $N_i \geq 1$ is geometric (exercise). This has two main consequences. Firstly,

$$P_i(T_i < \infty) = 1 \Leftrightarrow P_i(N_i = \infty) = 1.$$

In words: if starting from i you almost surely return to i , then you will visit i infinitely often. Secondly, we have

$$E_i[N_i] = \sum_{r=1}^{\infty} r P_i(N_i = r) = \sum_{r=1}^{\infty} r f_{ii}^r (1 - f_{ii}) = \frac{f_{ii}}{1 - f_{ii}}.$$

In particular,

$$P_i(T_i < \infty) < 1 \Leftrightarrow E_i[N_i] < \infty.$$

We collect these results for future reference.

Theorem 7.3. Recurrence

For any state $i \in E$,

$$P_i(T_i < \infty) = 1 \Leftrightarrow P_i(N_i = \infty) = 1, \quad (7.9)$$

and

$$P_i(T_i < \infty) < 1 \Leftrightarrow P_i(N_i = \infty) = 0 \Leftrightarrow E_i[N_i] < \infty. \quad (7.10)$$

In particular, the event $\{N_i = \infty\}$ has P_i -probability 0 or 1.

Consider a Markov chain with a state conventionally denoted by 0 such that $P_0(T_0 < \infty) = 1$. In view of the last theorem, the chain starting from state 0 will return infinitely often to this state. Let $\tau_1 = T_0, \tau_2, \dots$ be the successive return times to 0, and set $\tau_0 \equiv 0$.

By the strong Markov property, for any $k \geq 1$, the process after τ_k is independent of the process before τ_k (observe that condition $X_{\tau_k} = 0$ is always satisfied), and the process after τ_k is a Markov chain with the same transition matrix as the original chain, and with initial state 0, by construction. Therefore, we have the following

Theorem 7.4. Regenerative Cycle Theorem

Let $\{X_n\}_{n \geq 0}$ be an HMC with an initial state 0 that is almost surely visited infinitely often. Denoting by $\tau_0 = 0, \tau_1, \tau_2, \dots$ the successive times of visit to 0, the pieces of trajectory

$$\{X_{\tau_k}, X_{\tau_k+1}, \dots, X_{\tau_{k+1}-1}\}, k \geq 0,$$

are independent and identically distributed.

Such pieces are called the *regenerative cycles* of the chain between visits to state 0. Each random time τ_k is a *regeneration time*, in the sense that $\{X_{\tau_k+n}\}_{n \geq 0}$ is independent of the past X_0, \dots, X_{τ_k-1} and has the same distribution as $\{X_n\}_{n \geq 0}$. In particular, the sequence $\{\tau_k - \tau_{k-1}\}_{k \geq 1}$ is i.i.d.

Example 7.6. *Returns to Zero of the 1-D Symmetric Walk*

Let $\tau_1 = T_0, \tau_2, \dots$ be the successive return times to state 0 of the random walk on \mathbb{Z} of Example 2.1 with $p = \frac{1}{2}$. We shall admit that $P_0(T_0 < \infty) = 1$, a fact that will be proven in the next chapter, and obtain the probability distribution of T_0 given $X_0 = 0$.

Observe that for $n \geq 1$,

$$P_0(X_{2n} = 0) = \sum_{k \geq 1} P_0(\tau_k = 2n),$$

and therefore, for all $z \in \mathbb{C}$ such that $|z| < 1$,

$$\sum_{n \geq 1} P_0(X_{2n} = 0)z^{2n} = \sum_{k \geq 1} \sum_{n \geq 1} P_0(\tau_k = 2n)z^{2n} = \sum_{k \geq 1} E_0[z^{\tau_k}].$$

But $\tau_k = \tau_1 + (\tau_2 - \tau_1) + \dots + (\tau_k - \tau_{k-1})$ and therefore, in view of Theorem 7.4, and since $\tau_1 = T_0$,

$$E_0[z^{\tau_k}] = (E_0[z^{T_0}])^k.$$

In particular,

$$\sum_{n \geq 0} P_0(X_{2n} = 0)z^{2n} = \frac{1}{1 - E_0[z^{T_0}]}$$

(note that the latter sum includes the term for $n = 0$, that is, 1). Direct evaluation of the left-hand side yields

$$\sum_{n \geq 0} \frac{1}{2^{2n}} \frac{(2n)!}{n!n!} z^{2n} = \frac{1}{\sqrt{1 - z^2}}.$$

Therefore, the generating function of the return time to 0 given $X_0 = 0$ is

$$E_0[z^{T_0}] = 1 - \sqrt{1 - z^2}.$$

Its first derivative

$$\frac{z}{\sqrt{1 - z^2}}$$

tends to ∞ as $z \rightarrow 1$ from below via real values. Therefore, by Abel's theorem (see Theorem 1.2 of the Appendix),

$$E_0[T_0] = \infty.$$

We see that although given $X_0 = 0$ the return time is almost surely finite, it has an infinite expectation. \diamond

Problems

2.1.1 An HMC $\{X_n\}_{n \geq 0}$ with state space $E = \{0, 1, 2\}$ has the following transition matrix:

$$\mathbf{P} = \begin{matrix} & \begin{matrix} 0 & 1 & 2 \end{matrix} \\ \begin{matrix} 0 \\ 1 \\ 2 \end{matrix} & \begin{pmatrix} 0.2 & 0.5 & 0.3 \\ 0.1 & 0.1 & 0.8 \\ 0.5 & 0.2 & 0.3 \end{pmatrix} \end{matrix}.$$

Draw its transition graph and “read” from it $P(X_3 = 1 | X_0 = 1)$ and $P(X_7 = 2 | X_4 = 0)$.

2.1.2 For an HMC $\{X_n\}_{n \geq 0}$ with state space E , prove that for all $n \in \mathbb{N}$, and all states $i_0, i_1, \dots, i_{n-1}, i, j_1, j_2 \in E$,

$$\begin{aligned} P(X_{n+2} = j_2, X_{n+1} = j_1 | X_n = i, X_{n-1} = i_{n-1}, \dots, X_0 = i_0) \\ = P(X_{n+2} = j_2, X_{n+1} = j_1 | X_n = i), \end{aligned}$$

whenever both sides are well-defined.

2.1.3 The Markov property does not imply that the past and the future are independent given any information concerning the present. Find a simple example of an HMC $\{X_n\}_{n \geq 0}$ with state space $E = \{1, 2, 3, 4, 5, 6\}$ such that

$$P(X_2 = 6 | X_1 \in \{3, 4\}, X_0 = 2) \neq P(X_2 = 6 | X_1 \in \{3, 4\}).$$

2.1.4 Consider N balls numbered from 1 to N and placed in two urns A and B . Suppose that at stage n , urn A contains X_n balls. One then chooses a ball among the N balls at random (we may suppose that the balls are numbered and that a lottery gives the number of the selected ball, which can be in either of the two urns), and then chooses an urn, A with probability p , B with probability $q = 1 - p$. The selected ball is then placed in the selected urn, and the number of balls in urn A is now X_{n+1} . Show that $\{X_n\}_{n \geq 0}$ is an HMC, and give its transition matrix.

2.1.5 Let $\{X_n\}_{n \geq 0}$ be an HMC with state space E and transition matrix \mathbf{P} . Let τ be the first time n for which $X_n \neq X_0$, where $\tau = +\infty$ if $X_n = X_0$ for all $n \geq 0$. Compute $E[\tau | X_0 = i]$ in terms of p_{ii} .

2.1.6 Let $\{X_n\}_{n \geq 0}$ be an HMC with state space E and transition matrix \mathbf{P} , and let $E = \sum_{k=1}^{\infty} A_k$ be a partition of E . Define the process $\{\hat{X}_n\}_{n \geq 0}$ with state space $\hat{E} = \{\hat{1}, \hat{2}, \dots\}$ by $\hat{X}_n = \hat{k}$ if and only if $X_n \in A_k$. Show that a necessary and sufficient condition for $\{\hat{X}_n\}_{n \geq 0}$ to be an HMC for any initial distribution μ of $\{X_n\}_{n \geq 0}$ is that $\sum_{j \in A_\ell} p_{ij}$ be independent of $i \in A_k$ for all k, ℓ , and that in this case, $\hat{p}_{\hat{k}\hat{\ell}} = \sum_{j \in A_\ell} p_{ij}$ (any $i \in A_k$) is the general entry of the transition matrix $\hat{\mathbf{P}}$ of $\{\hat{X}_n\}_{n \geq 0}$.

2.2.1 Give the transition matrix of the HMC in Example 2.3 (Inventory) with the parameters $s = 2$, $S = 5$, and the following probability distribution for the demand: $P(Z = i) = \frac{1}{5}$ for all $i \in [0, 4]$.

2.2.2 In Example 2.4 (Branching Process), let $g_Z(z)$ be the generating function of any one of the i.i.d random variables $\{Z_n^{(j)}\}_{n \geq 1, j \geq 1}$. Show that the (i, j) th entry p_{ij} of the transition matrix of $\{X_n\}_{n \geq 0}$ is the coefficient of z^j in $(g_Z(z))^i$.

2.2.3 Let $\{Z_n\}_{n \geq 1}$ be an i.i.d sequence of geometric random variables: For $k \geq 0$, $P(Z_n = k) = (1 - p)^k p$, where $p \in (0, 1)$. Let $X_n = \max(X_0, Z_1, \dots, Z_n)$ be the record value at time n , and suppose X_0 is an \mathbb{N} -valued random variable independent of the sequence $\{Z_n\}_{n \geq 1}$. Show that $\{X_n\}_{n \geq 0}$ is an HMC and give its transition matrix.

2.2.4 Let $\{X_n\}_{n \geq 0}$ be an HMC with state space E and transition matrix \mathbf{P} . Define for $L \geq 1$, $Y_n = (X_n, X_{n+1}, \dots, X_{n+L})$. The process $\{Y_n\}_{n \geq 0}$ takes its values in $F = \{(i_0, \dots, i_L) \in E^{L+1}; p_{i_0 i_1} p_{i_1 i_2} \cdots p_{i_{L-1} i_L} > 0\}$. Prove that $\{Y_n\}_{n \geq 0}$ is an HMC and give the general entry of its transition matrix. (The chain $\{Y_n\}_{n \geq 0}$ is called the *snake chain* of length $L + 1$ associated with $\{X_n\}_{n \geq 0}$.) This problem is continued in Problems 2.4.6 and 2.5.2.

2.2.5 In certain digital communication systems, a sequence of 0's and 1's (bits) is encoded into a sequence of 0's, +1's, and -1's as follows. If the input sequence contains a 0, the output sequence contains a 0 at the same place. If the input sequence contains a 1, then the output sequence will have a -1 or a +1. The choice between -1 and +1 is made in such a way that -1's and +1's alternate in the output sequence. The first 1 is encoded as +1. For instance, 011101 becomes 0, +1, -1, +1, 0, -1. Find an automaton with four states +1, -1, 0_+ , and 0_- for which the sequence of visited states, not counting the initial state 0_+ , is exactly the encoded sequence (where 0_+ and 0_- are rewritten as 0) when it is fed by the input sequence.

Suppose that the input sequence is i.i.d. with 0 and 1 equiprobable. The sequence of states visited by the automaton is then an HMC. Compute its transition matrix \mathbf{P} , its stationary distribution π , and its iterates \mathbf{P}^n . Call $\{Y_n\}_{n \geq 0}$ the output sequence (taking its values in $\{0, -1, +1\}$). Compute $\lim_{n \rightarrow \infty} \{E[Y_n Y_{n+k}] - E[Y_n]E[Y_{n+k}]\}$ for all $k \geq 0$.

2.2.6 Prove Theorem 2.2.

2.3.1 Rat and Cat move between two rooms, using different paths. Their motions are independent, governed by their respective transition matrices

$$\begin{array}{cc} & \begin{matrix} 1 & 2 \end{matrix} \\ \begin{matrix} 1 \\ 2 \end{matrix} & \begin{pmatrix} 0.2 & 0.8 \\ 0.8 & 0.2 \end{pmatrix}, \end{array} \quad \begin{array}{cc} & \begin{matrix} 1 & 2 \end{matrix} \\ \begin{matrix} 1 \\ 2 \end{matrix} & \begin{pmatrix} 0.3 & 0.7 \\ 0.6 & 0.4 \end{pmatrix}, \end{array}$$

Cat starts from room 1, Rat from room 2. If they are ever in the same room, Cat eats Rat. How long will Rat survive on the average?

2.3.2 Give a Markovian model of a *tie-break* between two tennis players with the respective probabilities α and β of winning a point on their own service. Compute the probability that A wins when he starts serving. How long will the tie-break last on average when A starts serving? (Recall the rules of a tie-break when A starts serving: A has one service then B has two; A has two, etc. The first player with 7 points and at least 2 points ahead wins. If a player reaches 7 points with only a 1-point advantage, the tie-break continues until one of the players makes a break of 2 points.)

2.3.3 Three characters, A , B , and C , armed with guns, suddenly meet at the corner of a Washington, D.C., street, whereupon they naturally start shooting at one another. Each street-gang kid shoots every tenth second, as long as he is still alive. The probability of a hit for A , B , and C are α , β , and γ respectively. A is the most hated, and therefore, as long as he is alive, B and C ignore each other and shoot at A . For historical reasons not developed here, A cannot stand B , and therefore he shoots only at B while the latter is still

alive. Lucky C is shot at if and only if he is in the presence of A alone or B alone. What are the survival probabilities of A , B , and C , respectively?

2.3.4 Let $\{X_n\}_{n \geq 0}$ be a homogeneous MC with state space $E = \{1, 2, 3, 4\}$ and transition matrix

$$\mathbf{P} = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \end{matrix} & \begin{pmatrix} 0.2 & 0.3 & 0.5 & 0 \\ 0 & 0.2 & 0.3 & 0.5 \\ 0.5 & 0 & 0.2 & 0.3 \\ 0.3 & 0.5 & 0 & 0.2 \end{pmatrix} \end{matrix}.$$

What is the probability that when starting from state 1, the chain hits state 3 *before* it hits state 4?

2.3.5 Compute the mean duration of the game in Example 3.3.

2.4.1 Consider the chain of Example 1.1 (Machine Replacement) and suppose that $P(U < \infty) = 1$. Show that it is irreducible and that its period is $d = \gcd\{n \geq 1; P(U_1 = n) > 0\}$.

2.4.2 Show that a necessary and sufficient condition of irreducibility of the chain of Example 2.2 (Repair Shop) is that $P(Z_1 = 0) > 0$ and $P(Z_1 \geq 2) > 0$.

2.4.3 Show that the transition graph in Figure P 2.4.3 is irreducible. Find its period and its cyclic classes.




Figure P 2.4.3

2.4.4 Let \mathbf{P} be an irreducible transition matrix on the *finite* state space E . Show that a necessary and sufficient condition for \mathbf{P} to be aperiodic is the existence of an integer m such that \mathbf{P}^m has all its entries positive.

2.4.5 Show that an irreducible transition matrix \mathbf{P} with at least one state $i \in E$ such that $p_{ii} > 0$ is aperiodic.

2.4.6 Consider the *snake chain* of Problem 2.2.4. Show that if $\{X_n\}_{n \geq 0}$ is irreducible, then so is $\{Y_n\}_{n \geq 0}$. Compare the periods of these two chains in the irreducible case.

2.4.7 Let \mathbf{P} be an irreducible transition matrix on the countable state space E . Let d be its period. Show that the restriction of \mathbf{P}^d to any cyclic class is aperiodic.

2.5.1 Give the stationary distribution of the HMC with state space $E = \{1, 2, 3\}$ and transition matrix

$$\mathbf{P} = \begin{matrix} & \begin{matrix} 1 & 2 & 3 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{pmatrix} 1 - \alpha & \alpha & 0 \\ 0 & 1 - \beta & \beta \\ \gamma & 0 & 1 - \gamma \end{pmatrix} \end{matrix},$$

where $\alpha, \beta, \gamma \in (0, 1)$.

2.5.2 Consider the *snake chain* of Problems 2.2.4 and 2.4.6. Show that if $\{X_n\}_{n \geq 0}$ has a stationary distribution π , then $\{Y_n\}_{n \geq 0}$ also has a stationary distribution. Which one?

2.5.3 Prove that the HMC of Problem 2.2.3 cannot have a stationary distribution.

2.6.1 Consider the following process $\{X_n\}_{n \geq 0}$ taking its values in \mathbb{N} . It is constructed from a sequence $\{U_n\}_{n \geq 1}$ of i.i.d random variables taking their values in $\mathbb{N}^* = \{1, 2, \dots\}$ and from the initial state X_0 , a random variable taking its values in \mathbb{N} , and independent of $\{U_n\}_{n \geq 1}$. The process $\{X_n\}_{n \geq 0}$ decreases by 1 every unit of time, except when it has reached state 0. Then, if it is the k th time it has reached 0, it jumps to the value $U_k - 1$. Prove that this is an irreducible HMC and give its transition matrix. Show that it admits a unique stationary distribution π iff $E[U_1] < \infty$, and give an expression for π in terms of the distribution of U_1 . In this case, what is the reversed chain? Do you recognize it?

2.6.2 Let $\{X_n\}_{n \geq 0}$ be an HMC on the state space E with transition matrix \mathbf{P} , and suppose that $P(X_0 = i) = \pi(i) > 0$ for all $i \in E$, where π is a stationary distribution. Define the matrix $\mathbf{Q} = \{q_{ij}\}_{i, j \in E}$ by $\pi(i)q_{ij} = \pi(j)p_{ji}$. Construct $\{X_{-n}\}_{n \geq 1}$ by

$$\begin{aligned} P(X_{-1} = i_1, X_{-2} = i_2, \dots, X_{-k} = i_k \mid X_0 = i, X_1 = j_1, \dots, X_n = j_n) \\ = P(X_{-1} = i_1, X_{-2} = i_2, \dots, X_{-k} = i_k \mid X_0 = i) = q_{ii_1} q_{i_1 i_2} \cdots q_{i_{k-1} i_k} \end{aligned}$$

for all $k \geq 1, n \geq 1, i, i_1, \dots, i_k, j_1, \dots, j_n \in E$. Show that $\{X_n\}_{n \in \mathbb{Z}}$ is an HMC with transition matrix \mathbf{P} and that $P(X_n = i) = \pi(i)$, for all $i \in E$, all $n \in \mathbb{Z}$.

2.6.3 Is the HMC of Problem 2.5.1 reversible?

2.6.4 Let \mathbf{P} be a transition matrix on the countable state space E , with the positive stationary distribution π . Let A be a subset of the state space, and define the truncation of \mathbf{P} on A to be the transition matrix \mathbf{Q} indexed by A and given by

$$\begin{aligned} q_{ij} &= p_{ij} \text{ if } i, j \in A, i \neq j, \\ q_{ii} &= p_{ii} + \sum_{k \in \bar{A}} p_{ik}. \end{aligned}$$

Show that if (\mathbf{P}, π) is reversible, then so is $(\mathbf{Q}, \frac{\pi}{\pi(A)})$.

2.7.1 Let $\{X_n\}_{n \geq 0}$ be an HMC on the state space E with transition matrix \mathbf{P} . Define the sequence $\{\tau_k\}_{k \geq 0}$ recursively by $\tau_0 = 0$ and for $k \geq 0$,

$$\tau_{k+1} = \inf \{n \geq \tau_k + 1; X_n \neq X(\tau_k)\}$$

($= +\infty$ if $\tau_k = \infty$ or if $X_n = X(\tau_k)$ for all $n \geq \tau_k + 1$). Show that for all $k \geq 0$, τ_k is a stopping time of $\{X_n\}_{n \geq 0}$. Define for all $n \geq 0$, $Y_n = X(\tau_n)$ ($= \Delta \notin E$ if $\tau_n = \infty$). Show that $\{Y_n\}_{n \geq 0}$ is an HMC, and give its state space and transition matrix.

2.7.2 Let $\{X_n\}_{n \geq 0}$ be an HMC on the state space E with transition matrix \mathbf{P} . Let $\{\tau_k\}_{k \geq 1}$ be the successive return times to a given subset $F \subset E$. Assume these times almost surely finite. Let $X_0 \equiv 0 \in F$, and define $Y_n = X(\tau_n)$. Show that $\{Y_n\}_{n \geq 0}$ is an HMC with state space F .

2.7.3 Let $\{X_n\}_{n \geq 0}$ be an irreducible HMC with state space E and transition matrix \mathbf{P} . Let H be a subset of $E \times E$. One says that a transition of type H is observed at time k if $(X_{k-1}, X_k) \in H$. Let $\tau(0), \tau(1), \dots$ be the sequence of times of transitions of type H , with $\tau(n) = +\infty$ if the total number of transitions of type H is less than or equal to n . Observe that $\tau(0) \geq 1$, since X_{-1} is not defined. Define for each $n \geq 0$, $Y_n = X(\tau_n)$ if $\tau_n < \infty$, and $Y_n = \Delta$, an arbitrary element outside E , if $\tau_n = \infty$. Suppose that $\sum_{(i,j) \in H} p_{ij} > 0$. Show that for some subset $\tilde{E} \subset E$, to be identified, $\{Y_n\}_{n \geq 0}$ is an irreducible HMC on $\tilde{E} \cup \{\Delta\}$. Compare with Problems 2.7.1 and 2.7.2.

Recurrence and Ergodicity

1 Potential Matrix Criterion

1.1 Recurrent and Transient States

Consider a Markov chain taking its values in $E = \mathbb{N}$. There is a possibility that for any initial state $i \in \mathbb{N}$ the chain will never visit i after some finite random time. This is often an undesirable feature. For example, if the chain counts the number of customers waiting in line at a service counter (we shall see Markovian models of waiting lines, or *queues*, at different places in this book), such a behavior implies that the waiting line will eventually go beyond the limits of the waiting facility. In a sense, the corresponding system is unstable.

The good notion of stability for an irreducible HMC is that of *positive recurrence*, when any given state is visited infinitely often and when, moreover, the average time between two successive visits to this state is finite. The principal problem is to find sufficient, and maybe necessary, conditions guaranteeing stability. We begin with the *potential matrix criterion* (necessary and sufficient condition), which is of mainly theoretical interest, and the *stationary distribution criterion*. Further conditions, such as *Foster's theorem*, will be given in Chapter 5.

For the time being, we introduce the relevant definitions. First recall that T_i denotes the *return time* to state i .

Definition 1.1. *Recurrence and Transience*
 State $i \in E$ is called *recurrent* if

$$P_i(T_i < \infty) = 1, \tag{1.1}$$

and otherwise it is called *transient*. A recurrent state $i \in E$ is called *positive recurrent* if

$$E_i[T_i] < \infty, \tag{1.2}$$

and otherwise it is called *null* recurrent.

Example 1.1. Success Runs

The rule of the game is the following: A coin is tossed repeatedly, and whenever the result is tails (probability $q = 1 - p$), you go one step up the ladder, but if the result is heads, you fall all the way down. If X_n is your position at time n , $\{X_n\}_{n \geq 0}$ forms a homogeneous Markov chain with state space $E = \mathbb{N}$ and is a special case of the chain with the transition graph in Figure 3.1.1.




Figure 3.1.1. Transition graph of the success-runs chain

The state space of the chain of Figure 3.1.1 is $E = \mathbb{N}$, and we impose the condition

$$p_i \in (0, 1) \text{ for all } i \in E,$$

which guarantees irreducibility. We shall compute the probability of returning to state 0 and the mean return time to state 0, and from the expressions obtained, we shall deduce the nature of state 0.

There is just one way of going from state 0 to state 0 in exactly n steps. The corresponding path is $0, 1, 2, \dots, n - 1, 0$, and therefore $P_0(T_0 = 1) = p_0$, and for $n \geq 1$,

$$P_0(T_0 = n) = (1 - p_0) \cdots (1 - p_{n-2})p_{n-1}.$$

Defining $u_0 = 1$, and for $n \geq 1$,

$$u_n = (1 - p_0) \cdots (1 - p_{n-1}),$$

we see that

$$P_0(T_0 = n) = u_{n-1} - u_n.$$

Since

$$P_0(T_0 < \infty) = \sum_{n=1}^{\infty} P_0(T_0 = n) = \lim_{m \rightarrow \infty} \sum_{n=1}^m P_0(T_0 = n) = \lim_{m \rightarrow \infty} (1 - u_m),$$

we have

$$P_0(T_0 < \infty) = 1 - \lim_{m \uparrow \infty} \prod_{i=0}^{m-1} (1 - p_i).$$

Therefore, in view of a classical result on infinite products (see Theorem 1.9 of the Appendix)

$$P_0(T_0 < \infty) = 1 \Leftrightarrow \prod_{i=0}^{\infty} (1 - p_i) = 0 \Leftrightarrow \sum_{i=0}^{\infty} p_i = \infty. \quad \diamond$$

In general, it is not easy to check whether a given state is transient or recurrent. One of the goals of the theory of Markov chains is to provide criteria of recurrence. Sometimes, one is happy with just a sufficient condition, or a necessary condition.

The problem of finding useful (easy to check) conditions of recurrence is an active area of research. However, the theory has a few conditions that qualify as useful and are applicable to many practical situations. Although the next criterion is of theoretical rather than practical interest, it can be helpful in a few situations, for instance in the study of recurrence of random walks (see Examples 1.2 and 1.3 below).

1.2 Potential Matrix

The *potential matrix* \mathbf{G} associated with the transition matrix \mathbf{P} is defined by

$$\mathbf{G} = \sum_{n \geq 0} \mathbf{P}^n.$$

Its general term

$$g_{ij} = \sum_{n=0}^{\infty} p_{ij}(n) = \sum_{n=0}^{\infty} P_i(X_n = j) = \sum_{n=0}^{\infty} E_i[\mathbf{1}_{\{X_n=j\}}] = E_i \left[\sum_{n=0}^{\infty} \mathbf{1}_{\{X_n=j\}} \right]$$

is the average number of visits to state j , given that the chain starts from state i .

Theorem 1.1. Potential Matrix Criterion

State $i \in E$ is recurrent if and only if

$$\sum_{n=0}^{\infty} p_{ii}(n) = \infty. \quad (1.3)$$

Proof. Theorem 1.1 merely rephrases Theorem 7.3 of Chapter 2. □

Example 1.2. 1-D Random Walk

The corresponding Markov chain was described in Example 2.1 of Chapter 2. The nonzero terms of its transition matrix are

$$p_{i,i+1} = p, \quad p_{i,i-1} = 1 - p,$$

where $p \in (0, 1)$. We shall study the nature (recurrent or transient) of any one of its states, say, 0. We have $p_{00}(2n + 1) = 0$ and

$$p_{00}(2n) = \frac{(2n)!}{n!n!} p^n (1 - p)^n.$$

By Stirling's equivalence formula $n! \sim (n/e)^n \sqrt{2\pi n}$, the above quantity is equivalent to

$$\frac{[4p(1 - p)]^n}{\sqrt{\pi n}}, \quad (1.4)$$

and the nature of the series $\sum_{n=0}^{\infty} p_{00}(n)$ (convergent or divergent) is that of the series with general term (1.4). If $p \neq \frac{1}{2}$, in which case $4p(1 - p) < 1$, the latter series converges, and if $p = \frac{1}{2}$, in which case $4p(1 - p) = 1$, it diverges. In summary, the states of the 1-D random walk are transient if $p \neq \frac{1}{2}$, recurrent if $p = \frac{1}{2}$.

Example 7.6 of Chapter 2 shows that for the *symmetric* ($p = \frac{1}{2}$) 1-D random walk, the states are in fact *null* recurrent. \diamond

Example 1.3. 3-D Symmetric Random Walk

The state space of this HMC is $E = \mathbb{Z}^3$. Denoting by e_1, e_2 , and e_3 the canonical basis vectors of \mathbb{R}^3 (respectively $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$), the nonnull terms of the transition matrix of the 3-D symmetric random walk are given by

$$p_{x, x \pm e_i} = \frac{1}{6}.$$

We elucidate the nature of state, say, $0 = (0, 0, 0)$. Clearly, $p_{00}(2n + 1) = 0$ for all $n \geq 0$, and (exercise)

$$p_{00}(2n) = \sum_{0 \leq i+j \leq n} \frac{(2n)!}{(i!j!(n-i-j)!)^2} \left(\frac{1}{6}\right)^{2n}.$$

This can be rewritten as

$$p_{00}(2n) = \sum_{0 \leq i+j \leq n} \frac{1}{2^{2n}} \binom{2n}{n} \left(\frac{n!}{i!j!(n-i-j)!}\right)^2 \left(\frac{1}{3}\right)^{2n}.$$

Using the *trinomial formula*

$$\sum_{0 \leq i+j \leq n} \frac{n!}{i!j!(n-i-j)!} \left(\frac{1}{3}\right)^n = 1,$$

we obtain the bound

$$p_{00}(2n) \leq K_n \frac{1}{2^{2n}} \binom{2n}{n} \left(\frac{1}{3}\right)^n,$$

where

$$K_n = \max_{0 \leq i+j \leq n} \frac{n!}{i!j!(n-i-j)!}.$$

For large values of n , K_n is bounded as follows. Let i_0 and j_0 be the values of i, j that maximize $n!/(i!j!(n+j)!)$ in the domain of interest $0 \leq i+j \leq n$. From the definition of i_0 and j_0 , the quantities

$$\frac{n!}{(i_0 - 1)!j_0!(n - i_0 - j_0 + 1)!},$$

$$\frac{n!}{(i_0 + 1)!j_0!(n - i_0 - j_0 - 1)!},$$

$$\frac{n!}{i_0!(j_0 - 1)!(n - i_0 - j_0 + 1)!},$$

$$\frac{n!}{i_0!(j_0 + 1)!(n - i_0 - j_0 - 1)!},$$

are bounded by

$$\frac{n!}{i_0!j_0!(n - i_0 - j_0)!}.$$

The corresponding inequalities reduce to

$$n - i_0 - 1 \leq 2j_0 \leq n - i_0 + 1 \text{ and } n - j_0 - 1 \leq 2i_0 \leq n - j_0 + 1,$$

and this shows that for large n , $i_0 \sim n/3$ and $j_0 \sim n/3$. Therefore, for large n ,

$$p_{00}(2n) \sim \frac{n!}{(n/3)!(n/3)!2^{2n}e^n} \binom{2n}{n}.$$

By Stirling's equivalence formula, the right-hand side of the latter equivalence is in turn equivalent to

$$\frac{3\sqrt{3}}{2(\pi n)^{3/2}},$$

the general term of a divergent series. State 0 is therefore transient. ◇

Suppose that state $i \in E$ is recurrent, and accessible from state $j \in E$. That is, starting from j , the probability of visiting i at least once is positive (accessibility of i from j), and starting from i , the average number of visits to i is infinite (recurrence of i). Therefore, starting from j the average number of visits to i is infinite:

$$E_j[N_i] = \sum_{n \geq 1} p_{ji}(n) = \infty.$$

Similarly, if i is transient, then for any state $j \in E$,

$$E_j[N_i] = \sum_{n \geq 1} p_{ji}(n) < \infty.$$

1.3 Structure of the Transition Matrix

A theoretical application of the potential matrix criterion is to the proof that recurrence is a (communication) class property.

Theorem 1.2. *Recurrence Is a Class Property*

If i and j communicate, they are either both recurrent or both transient.

Proof. By definition, i and j communicate if and only if there exist integers M and N such that $p_{ij}(M) > 0$, $p_{ji}(N) > 0$. Going from i to j in M steps, then from j to j in n steps, then from j to i in N steps, is just one way of going from i back to i in $M + n + N$ steps. Therefore, $p_{ii}(M + n + N) \geq p_{ij}(M)p_{jj}(n)p_{ji}(N)$. Similarly, $p_{jj}(N + n + M) \geq p_{ji}(N)p_{ii}(n)p_{ij}(M)$. Therefore, writing $\alpha = p_{ij}(M)p_{ji}(N)$ (a strictly positive quantity), we have $p_{ii}(M + N + n) \geq \alpha p_{jj}(n)$ and $p_{jj}(M + N + n) \geq \alpha p_{ii}(n)$. This implies that the series $\sum_{n=0}^{\infty} p_{ii}(n)$ and $\sum_{n=0}^{\infty} p_{jj}(n)$ either both converge or both diverge. Theorem 1.1 concludes the proof. \square

It will be proven later in this chapter that positive recurrence (resp., null recurrence) is also a class property, in the sense that if states i and j communicate and if one of them is positive recurrent (resp., null recurrent), then so is the other.

An irreducible Markov chain has therefore all its states of the same nature: transient, positive recurrent, or null recurrent. We shall therefore call it a transient chain, a positive recurrent chain, or a null recurrent chain, and to determine to which category it belongs, it suffices to study *one* state, selecting the state for which the computations seem easiest (such as state 0 for the chain of Example 1.1).

It follows from the above discussion that there are two types of communication classes: the *transient classes* and the *recurrent classes*. Call T the set of all transient states and R the set of all recurrent states. The set R may be composed of several disjoint communication classes R_1, R_2 , etc. Any recurrent communication class, R_1 for instance, is closed. Indeed, if the chain goes from $i \in R_1$ to some $j \in E$, it will have to come back to i , since i is recurrent, and therefore i and j must communicate, so that j must be in R_1 . The communication structure of a transition matrix is therefore as shown in Figure 3.1.2.

2 Recurrence and Invariant Measures

The notion of invariant measure plays an important technical role in the recurrence theory of Markov chains. It extends the notion of stationary distribution.

Definition 2.1. *Invariant Measure.*

A nontrivial (that is, nonnull) vector $x = \{x_i\}_{i \in E}$ is called an *invariant measure* of the stochastic matrix $\mathbf{P} = \{p_{ij}\}_{i,j \in E}$ if for all $i \in E$,

$$x_i \in [0, \infty) \quad (2.1)$$

	R_1	R_2	R_3	T	
$\mathbf{P} =$	\mathbf{P}_1	0	0	0	R_1
	0	\mathbf{P}_2	0	0	R_2
	0	0	\mathbf{P}_3	0	R_3
					T

Figure 3.1.2

and

$$x_i = \sum_{j \in E} x_j p_{ji}. \tag{2.2}$$

(In abbreviated notation, $0 \leq x < \infty$ and $x^T \mathbf{P} = x^T$.)

Theorem 2.1. *Regenerative Form of Invariant Measure*

Let \mathbf{P} be the transition matrix of an irreducible recurrent HMC $\{X_n\}_{n \geq 0}$. Let 0 be an arbitrary state and let T_0 be the return time to 0. Define for all $i \in E$

$$x_i = E_0 \left[\sum_{n \geq 1} \mathbf{1}_{\{X_n=i\}} \mathbf{1}_{\{n \leq T_0\}} \right] \tag{2.3}$$

(For $i \neq 0$, x_i is the expected number of visits to state i before returning to 0). Then, for all $i \in E$,

$$x_i \in (0, \infty), \tag{2.4}$$

and x is an invariant measure of \mathbf{P} .

Observe that for $n \in [1, T_0]$, $X_n = 0$ if and only if $n = T_0$. Therefore,

$$x_0 = 1. \tag{2.5}$$

Also, $\sum_{i \in E} \sum_{n \geq 1} \mathbf{1}_{\{X_n=i\}} \mathbf{1}_{\{n \leq T_0\}} = \sum_{n \geq 1} \{ \sum_{i \in E} \mathbf{1}_{\{X_n=i\}} \} \mathbf{1}_{\{n \leq T_0\}} = \sum_{n \geq 1} \mathbf{1}_{\{n \leq T_0\}} = T_0$, and therefore

$$\sum_{i \in E} x_i = E_0[T_0]. \tag{2.6}$$

For the proof of Theorem 2.1, we introduce the quantity

$${}_0 p_{0i}(n) \stackrel{\text{def}}{=} E_0[\mathbf{1}_{\{X_n=i\}} \mathbf{1}_{\{n \leq T_0\}}] = P_0(X_1 \neq 0, \dots, X_{n-1} \neq 0, X_n = i). \tag{2.7}$$

This is the probability, starting from state 0, of visiting i at time n before returning to 0. From the definition of x ,

$$x_i = \sum_{n \geq 1} {}_0p_{0i}(n). \quad (2.8)$$

Proof. (of Theorem 2.1) We first prove (2.2). Observe that

$${}_0p_{0i}(1) = p_{0i} \quad (2.9)$$

and, using first-step analysis, for all $n \geq 2$,

$${}_0p_{0i}(n) = \sum_{j \neq 0} {}_0p_{0j}(n-1)p_{ji} \quad (2.10)$$

(see Problem 3.2.1). Summing up all the above equalities, and taking (2.8) into account, we obtain

$$x_i = p_{0i} + \sum_{j \neq 0} x_j p_{ji},$$

that is, (2.2), since $x_0 = 1$ (see (2.5)).

Next we show that $x_i > 0$ for all $i \in E$. Indeed, iterating (2.2), we find $x^T = x^T \mathbf{P}^n$, that is, since $x_0 = 1$,

$$x_i = \sum_{j \in E} x_j p_{ji}(n) = p_{0i}(n) + \sum_{j \neq 0} x_j p_{ji}(n).$$

If x_i were null for some $i \in E$, $i \neq 0$, the latter equality would imply that $p_{0i}(n) = 0$ for all $n \geq 0$, which means that 0 and i do not communicate, in contradiction to the irreducibility assumption.

It remains to show that $x_i < \infty$ for all $i \in E$. As before, we find that

$$1 = x_0 = \sum_{j \in E} x_j p_{j0}(n)$$

for all $n \geq 1$, and therefore if $x_i = \infty$ for some i , necessarily $p_{i0}(n) = 0$ for all $n \geq 1$, and this also contradicts irreducibility. \square

Theorem 2.2. Uniqueness of Invariant Measure

The invariant measure of an irreducible recurrent stochastic matrix is unique up to a multiplicative factor.

Proof. In the proof of Theorem 2.1 we showed that for an invariant measure y of an irreducible chain, $y_i > 0$ for all $i \in E$, and therefore, one can define, for all $i, j \in E$, the matrix \mathbf{Q} by

$$q_{ji} = \frac{y_i}{y_j} p_{ij}. \quad (2.11)$$

It is a transition matrix, since $\sum_{i \in E} q_{ji} = \frac{1}{y_j} \sum_{i \in E} y_i p_{ij} = \frac{y_i}{y_j} = 1$. The general term of \mathbf{Q}^n is

$$q_{ji}(n) = \frac{y_i}{y_j} p_{ij}(n). \quad (2.12)$$

Indeed, supposing (2.12) true for n ,

$$\begin{aligned} q_{ji}(n+1) &= \sum_{k \in E} q_{jk} q_{ki}(n) = \sum_{k \in E} \frac{y_k}{y_j} p_{kj} \frac{y_i}{y_k} p_{ik}(n) \\ &= \frac{y_i}{y_j} \sum_{k \in E} p_{ik}(n) p_{kj} = \frac{y_i}{y_j} p_{ij}(n+1), \end{aligned}$$

and (2.12) follows, by induction, for all $n \geq 1$.

Clearly, \mathbf{Q} is irreducible, since \mathbf{P} is irreducible (just observe that $q_{ji}(n) > 0$ if and only if $p_{ij}(n) > 0$ in view of (2.12)). Also, $p_{ii}(n) = q_{ii}(n)$, and therefore $\sum_{n \geq 0} q_{ii}(n) = \sum_{n \geq 0} p_{ii}(n)$, and this ensures that \mathbf{Q} is recurrent by the potential matrix criterion. Call $g_{ji}(n)$ the probability, relative to the chain governed by the transition matrix \mathbf{Q} , of returning to state i for the first time at step n when starting from j . First-step analysis gives $g_{i0}(n+1) = \sum_{j \neq 0} q_{ij} g_{j0}(n)$ (see Problem 3.2.1), that is, using (2.11),

$$y_i g_{i0}(n+1) = \sum_{j \neq 0} (y_j g_{j0}(n)) p_{ji}.$$

Recall that ${}_0 p_{0i}(n+1) = \sum_{j \neq 0} {}_0 p_{0j}(n) p_{ji}$, or, equivalently,

$$y_0 {}_0 p_{0i}(n+1) = \sum_{j \neq 0} (y_0 {}_0 p_{0j}(n)) p_{ji}.$$

We therefore see that the sequences $\{y_0 {}_0 p_{0i}(n)\}$ and $\{y_i g_{i0}(n)\}$ satisfy the same recurrence equation. Their first terms ($n = 1$), respectively $y_0 {}_0 p_{0i}(1) = y_0 p_{0i}$ and $y_i g_{i0}(1) = y_i q_{i0}$, are equal in view of (2.11). Therefore, for all $n \geq 1$,

$${}_0 p_{0i}(n) = \frac{y_i}{y_0} g_{i0}(n).$$

Summing up with respect to $n \geq 1$ and using $\sum_{n \geq 1} g_{i0}(n) = 1$ (\mathbf{Q} is recurrent), we obtain the announced result $x_i = \frac{y_i}{y_0}$. \square

Equality (2.6) and the definition of positive recurrence give the following.

Theorem 2.3. *Positive vs. Null Recurrence.*

An irreducible recurrent HMC is positive recurrent if and only if its invariant measures x satisfy

$$\sum_{i \in E} x_i < \infty. \quad (2.13)$$

Remark 2.1.

An HMC may well be irreducible and possess an invariant measure, and yet not be recurrent. The simplest example is the 1-D nonsymmetric random walk, which was shown to be transient (Example 1.2) and which admits $x_i \equiv 1$ for invariant measure. \diamond

3 Positive Recurrence

3.1 Stationary Distribution Criterion

In the previous section, an irreducible Markov chain was assumed recurrent, and it was shown that it has a unique stationary distribution if it is positive recurrent. It was also observed that the existence of an invariant measure is not sufficient for recurrence. It turns out, however, that the existence of a stationary probability distribution is necessary and sufficient for an irreducible chain (not a priori assumed recurrent) to be recurrent positive.

Theorem 3.1. *Stationary Distribution Criterion*

An irreducible homogeneous Markov chain is positive recurrent if and only if there exists a stationary distribution. Moreover, the stationary distribution π is, when it exists, unique, and $\pi > 0$.

Proof. The direct part follows from Theorems 2.1 and 2.3. For the converse part, assume the existence of a stationary distribution π . Iterating $\pi^T = \pi^T \mathbf{P}$, we obtain $\pi^T = \pi^T \mathbf{P}^n$, that is, for all $i \in E$,

$$\pi(i) = \sum_{j \in E} \pi(j) p_{ji}(n).$$

If the chain were transient, then, in view of the potential matrix criterion and the discussion following it, for all states i, j ,

$$\lim_{n \uparrow \infty} p_{ji}(n) = 0,$$

and since $p_{ji}(n)$ is bounded by 1 uniformly in j and n , by the dominated convergence theorem for series (see Theorem 1.6 of the Appendix)

$$\pi(i) = \lim_{n \uparrow \infty} \sum_{j \in E} \pi(j) p_{ji}(n) = \sum_{j \in E} \pi(j) \left(\lim_{n \uparrow \infty} p_{ji}(n) \right) = 0.$$

This contradicts the assumption that π is a stationary distribution (in particular, $\sum_{i \in E} \pi(i) = 1$). The chain must therefore be recurrent, and by Theorem 2.3, it is positive recurrent.

The stationary distribution π of an irreducible positive recurrent chain is unique (use Theorem 2.2 and the fact that there is no choice for a multiplicative factor but 1). Also recall that $\pi(i) > 0$ for all $i \in E$ (see Theorem 2.1). \square

Theorem 3.2. *Mean Return Time*

Let π be the unique stationary distribution of an irreducible positive recurrent chain, and let T_i be the return time to state i . Then

$$\pi(i) E_i[T_i] = 1. \tag{3.1}$$

Proof. This equality is a direct consequence of expression (2.3) for the invariant measure. Indeed, π is obtained by normalization of x : for all $i \in E$,

$$\pi(i) = \frac{x_i}{\sum_{j \in E} x_j},$$

and in particular, for $i = 0$, using (2.5) and (2.6),

$$\pi(0) = \frac{x_0}{\sum_{j \in E} x_j} = \frac{1}{E_0[T_0]}.$$

Since state 0 does not play a special role in the analysis, (3.1) is true for all $i \in E$. \square

The situation is extremely simple when the state space is finite.

Theorem 3.3. *Finite State Space and Positive Recurrence.*

An irreducible HMC with finite state space is positive recurrent.

Proof. We first show recurrence. If the chain were transient, then, from the potential matrix criterion and the observations following it, for all $i, j \in E$,

$$\sum_{n \geq 0} p_{ij}(n) < \infty,$$

and therefore, since the state space is finite

$$\sum_{j \in E} \sum_{n \geq 0} p_{ij}(n) < \infty.$$

But the latter sum is equal to

$$\sum_{n \geq 0} \sum_{j \in E} p_{ij}(n) = \sum_{n \geq 0} 1 = \infty,$$

a contradiction. Therefore, the chain is recurrent. By Theorem 2.1 it has an invariant measure x . Since E is finite, $\sum_{i \in E} x_i < \infty$, and therefore the chain is positive recurrent, by Theorem 3.1. \square

A “talk proof” of recurrence in Theorem 3.3 is available: The states cannot be all visited only a finite number of times; otherwise, there would exist a finite random time after which no state is visited!

3.2 Examples

Example 3.1. *Random Walk Reflected at 0*

This chain has the state space $E = \mathbb{N}$ and the transition graph of Figure 3.3.1. It is assumed that p_i (and therefore $q_i = 1 - p_i$) are in the open interval $(0, 1)$ for all $i \in E$, so that the chain is irreducible.

The invariant measure equation $x^T = x^T \mathbf{P}$ takes in this case the form

$$\begin{aligned} x_0 &= x_1 q_1, \\ x_i &= x_{i-1} p_{i-1} + x_{i+1} q_{i+1}, \quad i \geq 1, \end{aligned}$$




Figure 3.3.1. Reflected random walk

with $p_0 = 1$. The general solution is, for $i \geq 1$,

$$x_i = x_0 \frac{p_0 \cdots p_{i-1}}{q_1 \cdots q_i}.$$

The positive recurrence condition $\sum_{i \in E} x_i < \infty$ is

$$1 + \sum_{i \geq 1} \frac{p_0 \cdots p_{i-1}}{q_1 \cdots q_i} < \infty, \tag{3.2}$$

and if it is satisfied, the stationary distribution π is obtained by normalization of the general solution. This gives

$$\pi(0) = \left(1 + \sum_{i \geq 1} \frac{p_0 \cdots p_{i-1}}{q_1 \cdots q_i} \right)^{-1}, \tag{3.3}$$

and for $i \geq 1$,

$$\pi(i) = \pi(0) \frac{p_0 \cdots p_{i-1}}{q_1 \cdots q_i}. \tag{3.4}$$

In the special case where $p_i = p$, $q_i = q = 1 - p$, the positive recurrence condition becomes $1 + \frac{1}{q} \sum_{j \geq 0} \left(\frac{p}{q}\right)^j < \infty$, that is to say $p < q$, or equivalently,

$$p < \frac{1}{2}. \tag{3.5} \quad \diamond$$

Example 3.2. *Success Runs and Machine Replacement*

The result of Example 1.1 will be derived once again, this time via the stationary distribution criterion. We shall set $q_i = 1 - p_i$. Equality $x^T = x^T \mathbf{P}$ takes the form

$$x_0 = p_0 x_0 + p_1 x_1 + p_2 x_2 + \cdots,$$

and for $i \geq 1$,

$$x_i = q_{i-1} x_{i-1}.$$

Therefore, leaving aside the first equality, for $i \geq 1$,

$$x_i = (q_0 q_1 \cdots q_{i-1}) x_0.$$

Discarding the possibility $x_0 \leq 0$, which would imply that x is negative or null, the first equation is satisfied if and only if

$$1 = p_0 + q_0 p_1 + q_0 q_1 p_2 + \cdots,$$

that is, since $q_0 q_1 \cdots q_{n-1} p_n = q_0 q_1 \cdots q_{n-1} - q_0 q_1 \cdots q_n$,

$$\prod_{i=0}^{\infty} q_i = 0. \quad (3.5)$$

Since $q_i = 1 - p_i$ and $p_i \in (0, 1)$, the convergence criterion for infinite products (see Theorem 1.9 of the Appendix) tells that this is in turn equivalent to

$$\sum_{i=0}^{\infty} p_i = \infty.$$

The divergence of the series $\sum_{i=0}^{\infty} p_i$ is therefore a necessary and sufficient condition of existence of an invariant measure.

Recall, however, that the existence of an invariant measure does not imply recurrence (see Remark 2.1). But existence of a stationary distribution does imply recurrence (and actually positive recurrence) by the stationary distribution criterion.

Under condition (3.5), there exists an invariant measure, and this measure has finite mass ($\sum_{i=0}^{\infty} x_i < \infty$) if and only if

$$1 + \sum_{n=1}^{\infty} \left(\prod_{i=0}^{n-1} q_i \right) < \infty. \quad (3.6)$$

The stationary distribution is then given by

$$\pi(0) = \left(1 + \sum_{n=1}^{\infty} \left(\prod_{i=0}^{n-1} q_i \right) \right)^{-1} \quad (3.7)$$

and for $i \geq 1$,

$$\pi(i) = \left(\prod_{j=0}^{i-1} q_j \right) \pi_0. \quad (3.8)$$

In Problem 3.1.2, the reader is invited to verify that the success-runs chain of Example 1.1, and the machine-replacement chain of Example 1.1, Chapter 2, are the same if one sets

$$p_i = \frac{P(U = i + 1)}{P(U > i)}.$$

Inequality (3.6) then reads $E[U] < \infty$, and (3.7) and (3.8) give

$$\pi(i) = \frac{P(U > i)}{E[U]}. \quad (3.9)$$

◇

The stationary distribution criterion can also be used to prove instability.

Example 3.3. Instability of ALOHA

A typical situation in a multiple-access satellite communications system is the following. Users—each one identified with a message—contend for access to a single-channel satellite communications link for the purpose of transmitting messages. Two or more messages in the air at the same time jam each other, and are not successfully transmitted. The users are somehow able to detect a collision of this sort and will try to retransmit later the message involved in a collision. The difficulty in such communications systems resides mainly in the absence of cooperation among users, who are all unaware of the intention to transmit of competing users.

The *slotted* ALOHA protocol imposes on the users the following rules (see Fig. 3.3.2):

(i) Transmissions and retransmissions of messages can start only at equally spaced moments; the interval between two consecutive (re-)transmission times is called a *slot*; the duration of a slot is always larger than that of any message.

(ii) All *backlogged* messages, i.e., those messages having already tried unsuccessfully—maybe more than once—to get through the link, require retransmission independently of one another with probability $\nu \in (0, 1)$ at each slot. This is the so-called *Bernoulli retransmission policy*.

(iii) The *fresh messages*—those presenting themselves for the first time—immediately attempt to get through.

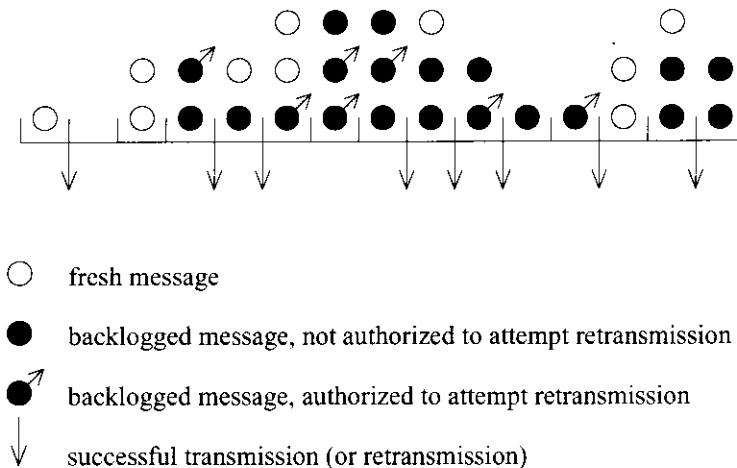


Figure 3.3.2. The ALOHA protocol

Let X_n be the number of backlogged messages at the beginning of slot n . The backlogged messages behave independently, and each one has probability ν of attempting retransmission

in slot n . In particular, if there are $X_n = k$ backlogged messages, the probability that i among them attempt to retransmit in slot n is

$$b_i(k) = \binom{k}{i} v^i (1-v)^{k-i}. \quad (3.10)$$

Let A_n be the number of fresh requests for transmission in slot n . The sequence $\{A_n\}_{n \geq 0}$ is assumed i.i.d with the distribution

$$P(A_n = j) = a_j. \quad (3.11)$$

The quantity

$$\lambda = E[A_n] = \sum_{i=1}^{\infty} i a_i \quad (3.12)$$

is the *traffic intensity*. We suppose that $a_0 + a_1 \in (0, 1)$, so that $\{X_n\}_{n \geq 0}$ is an irreducible HMC. Its transition matrix is

$$p_{ij} = \begin{cases} b_1(i)a_0 & \text{if } j = i - 1, \\ [1 - b_1(i)]a_0 + b_0(i)a_1 & \text{if } j = i, \\ [1 - b_0(i)]a_1 & \text{if } j = i + 1, \\ a_{j-i} & \text{if } j \geq i + 2. \end{cases} \quad (3.13)$$

The proof of (3.13) is by accounting. For instance, the first line corresponds to one among the i backlogged messages having succeeded to retransmit, and for this there should be no fresh arrival (probability a_0) and only one of the i backlogged messages allowed to retransmit (probability $b_1(i)$). The second line corresponds to one of the two events “no fresh arrival and zero or strictly more than two retransmission requests from the backlog” and “zero retransmission request from the backlog and one fresh arrival.”

Our objective in this example is to show that the system using the Bernoulli retransmission policy is *not stable*, in the sense that the chain $\{X_n\}_{n \geq 0}$ is *not positive recurrent*. Later on, in Example 4.2, a remedy to this situation will be proposed. To prove unstability, we must, in view of Theorem 3.1, contradict the existence of a stationary distribution π .

If such a stationary distribution existed, it should satisfy the balance equations

$$\begin{aligned} \pi(i) &= \pi(i)[[1 - b_1(i)]a_0 + b_0(i)a_1] + \pi(i - 1)[1 - b_0(i - 1)]a_1 \\ &\quad + \pi(i + 1)b_1(i + 1)a_0 + \sum_{\ell=2}^{\infty} \pi(i - \ell)a_{\ell} \end{aligned}$$

where $\pi(j) = 0$ if $j < 0$. Writing

$$P_N = \sum_{i=0}^N \pi(i)$$

and summing up the balance equations from $i = 0$ to N , we obtain

$$P_N = \pi(N)b_0(N)a_1 + \pi(N+1)b_1(N+1)a_0 + \sum_{\ell=0}^N a_\ell P_{N-\ell}.$$

This in turn gives

$$P_N(1 - a_0) = \pi(N)b_0(N)a_1 + \pi(N+1)b_1(N+1)a_0 + \sum_{\ell=1}^N a_\ell P_{N-\ell}.$$

But since P_N increases with N and $\sum_{\ell=1}^N a_\ell \leq \sum_{\ell=1}^{\infty} a_\ell = 1 - a_0$, we have

$$\sum_{\ell=1}^N a_\ell P_{N-\ell} \leq P_{N-1}(1 - a_0),$$

and therefore

$$P_N(1 - a_0) \leq \pi(N)b_0(N)a_1 + \pi(N+1)b_1(N+1)a_0 + P_{N-1}(1 - a_0),$$

from which it follows that

$$\frac{\pi(N+1)}{\pi(N)} \geq \frac{1 - a_0 - b_0(N)a_1}{b_1(N+1)a_0}.$$

Using expression (3.10), we obtain

$$\frac{\pi(N+1)}{\pi(N)} \geq \frac{(1 - a_0) - (1 - \nu)^N a_1}{(N+1)\nu(1 - \nu)^N a_0}.$$

For all values of $\nu \in (0, 1)$, the right-hand side of this inequality eventually becomes infinite, and this contradicts the equality $\sum_{N=1}^{\infty} \pi(N) = 1$ and the inequalities $\pi(N) > 0$ that π should satisfy as the stationary distribution of an irreducible Markov chain. \diamond

4 Empirical Averages

4.1 Ergodic Theorem

This subsection is devoted to the ergodic theorem for Markov chains. It gives conditions which guarantee that empirical averages of the type

$$\frac{1}{N} \sum_{k=1}^N f(X_k, \dots, X_{k+L})$$

converge to probabilistic averages.

Proposition 4.1.

Let $\{X_n\}_{n \geq 0}$ be an irreducible recurrent HMC, and let x denote the canonical invariant measure associated with state $0 \in E$,

$$x_i = E_0 \left[\sum_{n \geq 1} \mathbf{1}_{\{X_n=i\}} \mathbf{1}_{\{n \leq T_0\}} \right], \quad (4.1)$$

where T_0 is the return time to 0. Define for $n \geq 1$

$$v(n) = \sum_{k=1}^n \mathbf{1}_{\{X_k=0\}}. \quad (4.2)$$

Let $f : E \rightarrow \mathbb{R}$ be such that

$$\sum_{i \in E} |f(i)| x_i < \infty. \quad (4.3)$$

Then, for any initial distribution μ , P_μ -a.s.,

$$\lim_{N \uparrow \infty} \frac{1}{v(N)} \sum_{k=1}^N f(X_k) = \sum_{i \in E} f(i) x_i. \quad (4.4)$$

◇

Before the proof, we shall harvest the most interesting consequences.

Theorem 4.1. Ergodic Theorem

Let $\{X_n\}_{n \geq 0}$ be an irreducible positive recurrent Markov chain with the stationary distribution π , and let $f : E \rightarrow \mathbb{R}$ be such that

$$\sum_{i \in E} |f(i)| \pi(i) < \infty. \quad (4.5)$$

Then for any initial distribution μ , P_μ -a.s.,

$$\lim_{n \uparrow \infty} \frac{1}{N} \sum_{k=1}^N f(X_k) = \sum_{i \in E} f(i) \pi(i). \quad (4.6)$$

◇

Proof. Apply Proposition 4.1 to $f \equiv 1$. Condition (4.3) is satisfied, since in the positive recurrent case, $\sum_{i \in E} x_i < \infty$. Therefore, P_μ -a.s.,

$$\lim_{N \uparrow \infty} \frac{N}{v(N)} = \sum_{j \in E} x_j.$$

Now, f satisfying (4.5) also satisfies (4.3), since x and π are proportional, and therefore, P_μ -a.s.,

$$\lim_{N \uparrow \infty} \frac{1}{v(N)} \sum_{k=1}^N f(X_k) = \sum_{i \in E} f(i)x_i.$$

Combination of the above equalities gives, P_μ -a.s.,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N f(X_k) = \lim_{N \rightarrow \infty} \frac{v(N)}{N} \frac{1}{v(N)} \sum_{k=1}^N f(X_k) = \frac{\sum_{i \in E} f(i)x_i}{\sum_{j \in E} x_j},$$

from which (4.6) follows, since π is obtained by normalization of x . \square

Corollary 4.1.

Let $\{X_n\}_{n \geq 1}$ be an irreducible positive recurrent Markov chain with the stationary distribution π , and let $g : E^{L+1} \rightarrow \mathbb{R}$ be such that

$$\sum_{i_0, \dots, i_L} |g(i_0, \dots, i_L)| \pi(i_0) p_{i_0 i_1} \cdots p_{i_{L-1} i_L} < \infty. \quad (4.7)$$

Then for all initial distributions μ , P_μ -a.s.

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N g(X_k, X_{k+1}, \dots, X_{k+L}) = \sum_{i_0, i_1, \dots, i_L} g(i_0, i_1, \dots, i_L) \pi(i_0) p_{i_0 i_1} \cdots p_{i_{L-1} i_L}. \quad (4.8)$$

Proof. Apply Theorem 4.1 to the snake chain $\{(X_n, X_{n+1}, \dots, X_{n+L})\}_{n \geq 0}$ (see Problems 2.2.4, 2.4.6, and 2.5.2), which is irreducible recurrent and admits the stationary distribution

$$\pi(i_0) p_{i_0 i_1} \cdots p_{i_{L-1} i_L}. \quad \square$$

Note that

$$\sum_{i_0, i_1, \dots, i_L} g(i_0, i_1, \dots, i_L) \pi(i_0) p_{i_0 i_1} \cdots p_{i_{L-1} i_L} = E_\pi [g(X_0, \dots, X_L)]$$

Proof. (of Proposition 4.1.) Let $T_0 = \tau_1, \tau_2, \tau_3, \dots$ be the successive return times to state 0, and define

$$U_p = \sum_{n=\tau_p+1}^{\tau_{p+1}} f(X_n).$$

In view of the regenerative cycle theorem (Theorem 7.4 of Chapter 2), $\{U_p\}_{p \geq 1}$ is an i.i.d sequence. Moreover, assuming $f \geq 0$ and using the strong Markov property,

$$\begin{aligned} E[U_1] &= E_0 \left[\sum_{n=1}^{T_0} f(X_n) \right] \\ &= E_0 \left[\sum_{n=1}^{T_0} \sum_{i \in E} f(i) 1_{\{X_n=i\}} \right] = \sum_{i \in E} f(i) E_0 \left[\sum_{n=1}^{T_0} 1_{\{X_n=i\}} \right] \\ &= \sum_{i \in E} f(i) x_i. \end{aligned}$$

By hypothesis, this quantity is finite, and therefore the strong law of large numbers applies, to give

$$\lim_{n \uparrow \infty} \frac{1}{n} \sum_{p=1}^n U_p = \sum_{i \in E} f(i)x_i,$$

that is,

$$\lim_{n \uparrow \infty} \frac{1}{n} \sum_{k=\tau_0+1}^{\tau_{n-1}} f(X_k) = \sum_{i \in E} f(i)x_i. \quad (4.9)$$

Observing that

$$\tau_{v(n)} \leq n < \tau_{v(n)+1},$$

we have

$$\frac{\sum_{k=1}^{\tau_{v(n)}} f(X_k)}{v(n)} \leq \frac{\sum_{k=1}^n f(X_k)}{v(n)} \leq \frac{\sum_{k=1}^{\tau_{v(n)+1}} f(X_k)}{v(n)}.$$

Since the chain is recurrent, $\lim_{n \uparrow \infty} v(n) = \infty$, and therefore, from (4.9), the extreme terms of the above chain of inequality tend to $\sum_{i \in E} f(i)x_i$ as n goes to ∞ , and this implies (4.4). The case of a function f of arbitrary sign is obtained by considering (4.4) written separately for $f^+ = \max(0, f)$ and $f^- = \max(0, -f)$, and then taking the difference of the two equalities obtained this way. The difference is not an undetermined form $\infty - \infty$ due to hypothesis (4.3). \square

The version of the ergodic theorem for Markov chains featured in Theorem 4.1 is a kind of strong law of large numbers, and it can be used in simulations to compute, when π is unknown, quantities of the type $E_\pi[f(X_0)]$.

4.2 Examples

Example 4.1. Fixed-Age Retirement

We adopt the machine replacement interpretation of the success-runs Markov chain (Example 1.1). Assume positive recurrence. A visit of the chain to state 0 corresponds to a breakdown of a machine, and therefore, in view of the ergodic theorem,

$$\pi(0) = \lim_{N \uparrow \infty} \frac{1}{N} \sum_{k=1}^N 1_{\{X_k=0\}}$$

is the empirical frequency of breakdowns. Recall that

$$\pi(0) = E_0[T_0]^{-1},$$

where T_0 is the return time to 0. Here,

$$E_0[T_0] = E[U],$$

and therefore

$$\lim_{N \uparrow \infty} \frac{1}{N} \sum_{k=1}^N 1_{\{X_k=0\}} = \frac{1}{E[U]}. \quad (4.10)$$

Suppose that the cost of a breakdown is so important that it is better to replace a working machine during its lifetime (breakdown implies costly repairs, whereas replacement only implies moderate maintenance costs). The *fixed-age retirement policy* fixes an integer $T \geq 1$ and requires that a machine having reached age T be immediately replaced. We are interested in computing the empirical frequency of breakdowns (not replacements).

The success-runs chain corresponding to this situation is the same as before, except that the times U_n are replaced by $V_n = U_n \wedge T$. Also, a replacement (not breakdown) occurs at time n if and only if $X_n = 0$ and $X_{n-1} = T - 1$. But $X_{n-1} = T - 1$ implies $X_n = 0$, and therefore a replacement occurs at time n if and only if

$$X_{n-1} = T - 1.$$

The empirical frequency of replacements is, therefore, in view of the ergodic theorem,

$$\lim_{N \uparrow \infty} \frac{1}{N} \sum_{k=1}^N 1_{\{X_k=T-1\}} = \pi(T - 1).$$

Formula (3.9) gives

$$\pi(T - 1) = \frac{P(V \geq T)}{E[V]},$$

and therefore, since $V = U \wedge T$,

$$\pi(T - 1) = \frac{P(U \geq T)}{E[U \wedge T]}.$$

The empirical frequency of visits to state 0 is, by (4.10),

$$\frac{1}{E[U \wedge T]}.$$

The empirical frequency of breakdowns is therefore

$$\frac{1}{E[U \wedge T]} - \frac{P(U \geq T)}{E[U \wedge T]} = \frac{P(U < T)}{E[U \wedge T]}. \quad \diamond$$

Example 4.2. Cash Management

The level of cash in a bank at the beginning of day n is X_n , the number of idling kilodollars in the safe. For each day a kilodollar spends idling in the safe, the banker incurs a loss of r dollars. In the absence of control, the cash level fluctuates as a symmetric random walk.

The management policy uses two integers s and S , where $0 < s < S$. The number S is an upper barrier: If at the beginning of day n , the cash level is $X_n = S - 1$ and if during




Figure 3.4.1. The (s, S) policy of cash management

day n one more kilodollar enters the bank, then during the night from day n to day $n + 1$, $S - s$ kilodollars are transferred to a more profitable occupation (treasury bills, etc.), and the level of cash at the beginning of day $n + 1$ is therefore $X_{n+1} = s$.

Also, if $X_n = 1$ and one more kilodollar is removed from the safe during day n , the safe is replenished at level s during the night, so that $X_{n+1} = s$.

Therefore, the process $\{X_n\}_{n \geq 0}$ takes its values in $\{1, \dots, S - 1\}$ if we assume that $X_0 \in [1, S - 1]$. A typical trajectory of the cash level is as shown in Figure 3.4.1.a.

The process $\{X_n\}_{n \geq 0}$ is an HMC with the transition graph of Figure 3.4.1.b, where all transition probabilities are equal to $\frac{1}{2}$. In addition to the storing cost of $\$r$ per day per kilodollar, the banker incurs a transaction cost of $\$\alpha$ whenever one kilodollar is moved into or out of the safe. The long-run cost per day of the (s, S) policy is therefore

$$\lim_{N \uparrow \infty} \frac{1}{N} \sum_{n=1}^N \left[\left(\sum_{k=1}^{S-1} rk 1_{\{X_n=k\}} \right) + \alpha(S-s) 1_{\{X_n=S-1, X_{n+1}=s\}} + \alpha s 1_{\{X_n=1, X_{n+1}=s\}} \right].$$

The chain $\{X_n\}_{n \geq 0}$ is obviously irreducible, and since the state space is finite, it is positive recurrent. Let π be its stationary distribution.

By the ergodic theorem,

$$\lim_{N \uparrow \infty} \frac{1}{N} \sum_{n=1}^N \left(\sum_{k=1}^{S-1} rk 1_{\{X_n=k\}} \right) = r \sum_{k=1}^{S-1} k \pi(k)$$

and

$$\begin{aligned} \lim_{N \uparrow \infty} \frac{1}{N} \sum_{n=1}^N (\alpha(S-s) \mathbf{1}_{\{X_n=S-1, X_{n+1}=s\}} + \alpha s \mathbf{1}_{\{X_n=1, X_{n+1}=s\}}) \\ = \alpha(S-s)\pi(S-1)\frac{1}{2} + \alpha s \pi(1)\frac{1}{2}, \end{aligned}$$

so that the long run cost per day is

$$C = r \sum_{k=1}^{S-1} k\pi(k) + \frac{1}{2}\alpha(S-s)\pi(S-1) + \frac{1}{2}\alpha s \pi(1).$$

The stationary distribution π is determined by the global balance equations

$$\begin{aligned} \pi(1) &= \frac{1}{2}\pi(2), \\ \pi(2) &= \frac{1}{2}\pi(1) + \frac{1}{2}\pi(3), \\ &\vdots \\ \pi(s-1) &= \frac{1}{2}\pi(s-2) + \frac{1}{2}\pi(s), \\ \pi(s) &= \frac{1}{2}\pi(s-1) + \frac{1}{2}\pi(s+1) + \frac{1}{2}\pi(S-1) + \frac{1}{2}\pi(1), \\ \pi(s+1) &= \frac{1}{2}\pi(s) + \frac{1}{2}\pi(s+2), \\ &\vdots \\ \pi(S-2) &= \frac{1}{2}\pi(S-3) + \frac{1}{2}\pi(S-1), \\ \pi(S-1) &= \frac{1}{2}\pi(S-2), \end{aligned}$$

together with the normalizing condition $\sum_{i=1}^{S-1} \pi(i) = 1$. Since the rank of the transition matrix \mathbf{P} is $|E| - 1$, one of the $S - 1$ balance equations can be removed, and it can be chosen arbitrarily. We choose the equation relative to s .

From equations 1 to $s - 1$, we obtain for $i \in [1, s]$,

$$\pi(i) = i\pi(1).$$

Equations $s + 1$ to S give for $j \in [1, S - s]$,

$$\pi(S - j) = j\pi(S - 1).$$

Taking $i = s$ and $j = S - s$, we have $\pi(s) = s\pi(1) = (S - s)\pi(S - 1)$, so that

$$\pi(S - 1) = \frac{s}{S - s}\pi(1).$$

From the normalization equation we have

$$\pi(1)(1 + \cdots + s) + \pi(S-1)(1 + \cdots + (S-s-1)) = 1,$$

that is,

$$\pi(1) \frac{s(s+1)}{2} + \pi(S-1) \frac{(S-s-1)(S-s)}{2} = 1,$$

and taking into account the expression of $\pi(S-1)$ in terms of $\pi(1)$, we have

$$\pi(1) \left[\frac{s(s+1)}{2} + \frac{(S-s-1)s}{2} \right] = 1,$$

that is,

$$\pi(1) = \frac{2}{sS}.$$

Therefore, finally,

$$\pi(i) = \begin{cases} \frac{2i}{sS} & \text{if } i \in [1, s], \\ \frac{2(S-i)}{(S-s)S} & \text{if } i \in [s+1, S]. \end{cases}$$

The cost C can therefore be computed in terms of s and S and this can be exploited for optimization purposes; see, for instance, (Taylor and Karlin, 1984), pp. 110-113. \diamond

4.3 Renewal Reward Theorem

Theorem 4.2. *Renewal Reward Theorem*

Let $\{S_n\}_{n \geq 1}$ be an i.i.d sequence of positive random variables such that $E[S_1] < \infty$, and let R_0 be a finite nonnegative random variable independent of this sequence. Define for all $n \geq 0$, $R_{n+1} = R_n + S_{n+1}$ and for $t \geq 0$, $N(t) = \sum_{n=1}^{\infty} 1_{\{R_n \leq t\}}$.

Now let $\{Y_n\}_{n \geq 1}$ be an i.i.d sequence of random variables such that $E[|Y_1|] < \infty$. Then

$$\lim_{t \uparrow \infty} \frac{N(t)}{t} = \frac{1}{E[S_1]} \quad (4.11)$$

and

$$\lim_{t \uparrow \infty} \frac{\sum_{n=1}^{N(t)} Y_n}{t} = \frac{E[Y_1]}{E[S_1]}. \quad (4.12)$$

Proof. Since $R_{N(t)} \leq t < R_{N(t)+1}$, we have

$$\frac{N(t)}{R_{N(t)+1}} < \frac{N(t)}{t} \leq \frac{N(t)}{R_{N(t)}}.$$

But the rightmost term is the inverse of

$$\frac{R_{N(t)}}{N(t)} = \frac{R_0 + \sum_{n=1}^{N(t)} S_n}{N(t)}.$$

By the strong law of large numbers and the fact that $\lim_{t \uparrow \infty} N(t) = \infty$ (the S_n 's are finite), this quantity tends to $E[S_1]$, and similarly, $\frac{R_{N(t)+1}}{N(t)} = \frac{R_{N(t)+1}}{N(t)+1} \frac{N(t)+1}{N(t)}$ tends to $E[S_1]$ as $t \rightarrow \infty$.

The proof of (4.12) follows from the SLLN and (4.11), since

$$\frac{\sum_{n=1}^{N(t)} Y_n}{t} = \frac{\sum_{n=1}^{N(t)} Y_n}{N(t)} \cdot \frac{N(t)}{t}. \quad \square$$

Note that the above result covers the discrete-time case (the S_n 's and t are integers) as well as the continuous-time case (the S_n 's and t are real).

The next result extends formula (2.3).

Example 4.3. *Regenerative Form of the Stationary Distribution*

Let $\{X_n\}_{n \geq 0}$ be a positive recurrent HMC with state space E and stationary distribution π . Let S be a stopping time of this chain and let i be a state, such that, P_i -a.s., $S \in (0, \infty)$ and $X_S = i$. Then for all $j \in E$,

$$E_i \left[\sum_{k=0}^{S-1} 1_{\{X_n=j\}} \right] = E_i [S] \pi(j). \quad (4.13)$$

To prove this, observe that S , being a stopping time of the chain, can be written as

$$S = \phi(X_0, X_1, \dots)$$

for some functional ϕ , namely

$$S = \sum_{m=0}^{\infty} m \psi_m(X_0, \dots, X_m),$$

where

$$1_{\{S=m\}} = \psi_m(X_0, \dots, X_m).$$

Define the sequence $\{S_n\}_{n \geq 1}$ by $S_1 = S$, and for $k \geq 0$,

$$S_{k+1} = \phi(X_{R_k}, X_{R_k+1}, \dots),$$

where $R_0 = 0$ and, for $k \geq 1$, $R_k = S_1 + \dots + S_k$. As a matter of fact, since S is a stopping time, S_{k+1} depends only on the cycle $C_k = (X_{R_k}, \dots, X_{R_k+S_{k+1}-1})$ (recall that $X_{R_k+S_{k+1}} = i$). But by the strong Markov property, the cycles $\{C_k\}_{k \geq 0}$ are i.i.d under P_i , and in particular, $\{S_n\}_{n \geq 1}$ is i.i.d.

We first suppose that $E_i [S] < \infty$. The renewal reward theorem with

$$Y_n = \sum_{k=R_{n-1}}^{R_n-1} 1_{\{X_k=j\}}$$

(observe that $E_i [|Y_1|] \leq E_i [S_1] < \infty$) gives

$$\lim_{t \uparrow \infty} \frac{\sum_{n=1}^{N(t)} Y_n}{t} = \frac{E_i \left[\sum_{k=0}^{S-1} 1_{\{X_k=j\}} \right]}{E_i [S]}.$$

But

$$\frac{\sum_{n=1}^{N(t)} Y_n}{t} = \frac{\sum_{k=R_0}^{R_{N(t)}-1} 1_{\{X_k=j\}}}{t} = \frac{\sum_{k=R_0}^{R_{N(t)}-1} 1_{\{X_k=j\}}}{R_{N(t)}} \cdot \frac{R_{N(t)}}{t}.$$

Now, $\lim_{t \uparrow \infty} \frac{R_{N(t)}}{t} = 1$ (exercise), and therefore, by the ergodic theorem,

$$\lim_{t \uparrow \infty} \frac{\sum_{n=1}^{N(t)} Y_n}{t} = \pi(j).$$

This proves the theorem when $E_i [S] < \infty$. If $E_i [S] = \infty$, consider $S^{(n)} = S \wedge \tau_n(i)$, where $\tau_n(i)$ is the n -return time to i , write (4.13) for $S^{(n)}$, and let $n \rightarrow \infty$, to obtain $E_i \left[\sum_{k=0}^{S-1} 1_{\{X_k=j\}} \right] = \infty$. \diamond

Example 4.4.

(Aldous and Fill, 1998) Let i and j be two distinct states and let S be the first time of return to i after the first visit to j . Then $E_i [S] = E_i [T_j] + E_j [T_i]$ (use the strong Markov property at T_j). Also,

$$E_i \left[\sum_{n=0}^{S-1} 1_{\{X_n=j\}} \right] = E_i \left[\sum_{n=T_j}^{S-1} 1_{\{X_n=j\}} \right] = E_j \left[\sum_{n=0}^{T_i-1} 1_{\{X_n=j\}} \right],$$

where the last equality is the strong Markov property. Therefore, by (4.13),

$$E_j \left[\sum_{n=0}^{T_i-1} 1_{\{X_n=j\}} \right] = \pi(j) (E_i [T_j] + E_j [T_i]). \quad (4.14)$$

Using words, the left-hand side of this equality is

$$E_j [\text{number of visits to } j \text{ before } i].$$

Now, the probability that j is not visited between two successive visits of i is $P_i (T_j > T_i)$. Therefore, the number of visits to i (including time 0) before T_j has a geometric distribution with parameter $p = P_i (T_j > T_i)$, and the average number of such visits is

$$\frac{1}{P_i (T_j < T_i)}.$$

Therefore, by (4.14), after exchanging the roles of i and j ,

$$P_i (T_j < T_i) = \frac{1}{\pi(i) (E_i [T_j] + E_j [T_i])}. \quad (4.15)$$

Theorem 6.4 of Chapter 8 gives an algebraic method of computation of $E_i [T_j]$ when the state space is finite. \diamond

Problems

3.1.1 For an HMC with state space $E = \{0, 1\}$ and transition matrix

$$\mathbf{P} = \begin{matrix} & 0 & 1 \\ \begin{matrix} 0 \\ 1 \end{matrix} & \begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix} \end{matrix},$$

compute $P_0(T_0 = n)$ and $E_0[T_0]$.

3.1.2 Verify that the success-runs chain of Example 1.1 and the machine-replacement chain of Example 1.1, Chapter 2, are the same if one sets

$$p_i = \frac{P(U = i + 1)}{P(U > i)}.$$

3.1.3 Let $\{X_n\}_{n \geq 0}$ be a HMC with state space E and n -step transition matrix $\mathbf{P}^n = \{p_{ij}(n)\}_{i, j \in E}$. Prove directly that if $i \in E$ is a transient state, then $\lim_{n \uparrow \infty} p_{ji}(n) = 0$ for all $j \in E$.

3.1.4 Consider the HMC of Example 2.2, Chapter 2. Recall that $X_n \in E \equiv \mathbb{N}$, and for all $n \geq 0$,

$$X_{n+1} = (X_n - 1)^+ + Z_{n+1},$$

where $a^+ = \sup(a, 0)$ and $\{Z_n\}_{n \geq 1}$ is an i.i.d sequence of random variables with values in \mathbb{N} . Prove directly, using the strong law of large numbers, that if $E[Z_1] > 1$, state 0 is transient, and that if $E[Z_1] < 1$, state 0 is recurrent.

3.1.5 Show that the transition matrix \mathbf{P} and the potential matrix \mathbf{G} are related by

$$\mathbf{G}(I - \mathbf{P}) = I,$$

where I is the identity matrix. Conclude from this that if the state space E is finite, the potential matrix \mathbf{G} cannot have all its entries finite. Conclude that an irreducible HMC with finite state space is recurrent.

3.1.6 The symmetric random walk on \mathbb{Z}^2 is an HMC $\{X_n\}_{n \geq 0}$ with state space $E = \mathbb{Z}^2$ and with a transition matrix having for nonzero entries

$$p_{x, x \pm e_i} = \frac{1}{4}$$

($x \in \mathbb{Z}^2, i = 1, 2$), where $e_1 = (1, 0)$ and $e_2 = (0, 1)$. Compute $P(X_{2n} = (0, 0) \mid X_0 = (0, 0))$ and show that the above HMC is irreducible and recurrent. Show that it is in fact null-recurrent.

3.2.1 Let $\{X_n\}_{n \geq 0}$ be an irreducible recurrent HMC with state space E and transition matrix \mathbf{P} . Let $0 \in E$ be some state, and define ${}_0p_{0i}(n) = E_0[\mathbf{1}_{\{X_n = i, n \leq T_0\}}]$ and $f_{ij}(n) = P_i(T_j = n)$. Show that for all $n \geq 0, i \in E$,

$${}_0p_{0i}(n + 1) = \sum_{j \neq 0} {}_0p_{0j}(n) p_{ji}$$

and

$$f_{i0}(n + 1) = \sum_{j \neq 0} p_{ij} f_{j0}(n).$$

3.2.2 A countable number of particles move independently in the countable space E , each according to a Markov chain with the transition matrix \mathbf{P} . Let $A_n(i)$ be the number of particles in state $i \in E$ at time $n \geq 0$, and suppose that the random variables $A_0(i)$, $i \in E$, are independent Poisson random variables with respective means $\mu(i)$, $i \in E$, where $\mu = \{\mu(i)\}_{i \in E}$ is an invariant measure of \mathbf{P} . Show that for all $n \geq 1$, the random variables $A_n(i)$, $i \in E$, are independent Poisson random variables with respective means $\mu(i)$, $i \in E$.

3.2.3 A stochastic matrix \mathbf{P} on the state space E is called *doubly stochastic* if for all states i , $\sum_{j \in E} p_{ji} = 1$. Suppose in addition that \mathbf{P} is irreducible, and that E is *infinite*. Find the invariant measure of \mathbf{P} . Show that \mathbf{P} cannot be positive recurrent.

3.2.4 Consider the transition matrix

$$\mathbf{P} = \begin{pmatrix} 1 - \alpha & \alpha & 0 \\ 0 & 1 - \beta & \beta \\ \gamma & 0 & 1 - \gamma \end{pmatrix},$$

where $\alpha, \beta, \gamma \in (0, 1)$. Show that it is irreducible and compute directly its stationary probability with the help of Formula (2.3).

3.3.1 Let τ be the first return time to initial state of an irreducible positive recurrent HMC $\{X_n\}_{n \geq 0}$, that is,

$$\tau = \inf \{n \geq 1; X_n = X_0\},$$

with $\tau = +\infty$ if $X_n \neq X_0$ for all $n \geq 1$. Compute the expectation of τ when the initial distribution is the stationary distribution π . Conclude that it is finite if and only if E is finite. When E is infinite, is this in contradiction to positive recurrence?

3.3.2 Let $\{X_n\}_{n \geq 0}$ be an HMC with state space E and transition matrix \mathbf{P} , and let $\{Y_n\}_{n \geq 0}$ be the sequence of *new values* of $\{X_n\}_{n \geq 0}$. For instance,

$$\begin{array}{rcccccccccccc} n & = & 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 \\ X_n & = & 1 & 1 & 1 & 2 & 2 & 1 & 3 & 3 & 3 & 2 & 1 \\ & & Y_0 = 1 & & & Y_1 = 2 & & Y_2 = 1 & Y_3 = 3 & & & Y_4 = 2 & Y_5 = 1 \end{array}$$

Show that $\{Y_n\}_{n \geq 0}$ is a HMC and give its transition matrix \mathbf{Q} in terms of \mathbf{P} .

Show that $\{X_n\}_{n \geq 0}$ is irreducible and recurrent if and only if $\{Y_n\}_{n \geq 0}$ is irreducible and recurrent.

In the irreducible recurrent case, express the invariant measure of $\{Y_n\}_{n \geq 0}$ in terms of $(p_{ii}, i \in E)$ and the invariant measure of $\{X_n\}_{n \geq 0}$.

In the irreducible recurrent case, give a necessary and sufficient condition for $\{Y_n\}_{n \geq 0}$ to be positive recurrent when $\{X_n\}_{n \geq 0}$ is positive recurrent in terms of the stationary distribution π of $\{X_n\}_{n \geq 0}$ and $(p_{ii}, i \in E)$.

In the irreducible recurrent case, give a necessary and sufficient condition for $\{Y_n\}_{n \geq 0}$ to be positive recurrent when $\{X_n\}_{n \geq 0}$ is null recurrent in terms of an invariant measure μ of $\{X_n\}_{n \geq 0}$ and p_{ii} .

3.3.3 Stones S_1, \dots, S_M are placed in line. At each time n a stone is selected at random, and this stone and the one ahead of it in the line exchange positions. If the selected stone is at the head of the line, nothing is changed. For instance, with $M = 5$, Let the current configuration be $S_2 S_3 S_1 S_5 S_4$. If S_5 is selected, the new situation is $S_2 S_3 S_5 S_1 S_4$, whereas If S_1 is selected, the configuration is not altered. At each step, stone S_i is selected with probability $\alpha_i > 0$. Call X_n the situation at time n , for instance $X_n = (S_{i_1}, \dots, S_{i_M})$, meaning that stone S_{i_j} is in the j th position.

Show that $\{X_n\}_{n \geq 0}$ is an irreducible positive recurrent HMC and that its stationary distribution is given by the formula

$$\pi(S_{i_1}, \dots, S_{i_M}) = C \alpha_{i_1}^M \alpha_{i_2}^{M-1} \dots \alpha_{i_M},$$

for some normalizing constant C .

3.3.4 Let A, B, C , and D be four points on the unit circle forming a square, and let $\{Y_n\}_{n \geq 0}$ and $\{Z_n\}_{n \geq 0}$ be two stochastic processes with state space $\{A, B, C, D\}$ denoted by $\{0, 1, 2, 3\}$, and with the dynamics

$$\begin{aligned} Y_{n+1} &= Y_n + B_{n+1} \pmod{4}, \\ Z_{n+1} &= Z_n + C_{n+1} \pmod{4}, \end{aligned}$$

where $\{B_n\}_{n \geq 1}$ and $\{C_n\}_{n \geq 1}$ are independent i.i.d sequences with $P(B_n = \pm 1) = P(C_n = \pm 1) = \frac{1}{3} = P(B_n = 0) = P(C_n = 0)$. Suppose, moreover, that $Y_0, Z_0, \{B_n\}_{n \geq 1}$, and $\{C_n\}_{n \geq 1}$ are mutually independent. What is the average time separating two successive coincidences of the processes $\{Y_n\}_{n \geq 1}$ and $\{Z_n\}_{n \geq 1}$?

3.3.5 A knight moves randomly on a chessboard, making each admissible move with equal probability, and starting from a corner. What is the average time he takes to return to the corner he started from?

3.4.1 Consider a production line where each manufactured item may be defective with probability $p \in (0, 1)$. The following inspection plan is proposed with a view to detecting defective items without checking every single one.

It has 2 phases: In phase A, the probability of inspecting an article is $r \in (0, 1)$. In phase B, all the articles are inspected. One switches from phase A to phase B as soon as a defective item is detected. One switches from phase B to phase A as soon as a sequence of N successive acceptable items has been found.

Let $\{X_n\}_{n \geq 0}$ be the process taking the values E_0, \dots, E_N , where if $j \in [0, N - 1]$, E_j means that the inspection plan is in phase B with j successive good items observed, and E_N means that the plan is in phase A.

Prove that $\{X_n\}_{n \geq 0}$ is an irreducible HMC, give its transition graph, and show that it is ergodic. Find its stationary distribution.

Find the proportion of items inspected and give the *efficiency* of the inspection plan, which is by definition equal to the ratio of the long-run proportion of *detected* defective items over the proportion of defective items.

3.4.2 A sequence of A's and B's is formed as follows. The first item is chosen at random, $P(A) = P(B) = \frac{1}{2}$, as is the second item, independently of the first one. When the first

$n \geq 2$ items have been selected, the $(n + 1)$ st is chosen, conditionally with respect to the pair at position $n - 1$ and n , as follows:

$$P(A | AA) = \frac{1}{2}, P(A | AB) = \frac{1}{2}, P(A | BA) = \frac{1}{4}, P(A | BB) = \frac{1}{4}.$$

What is the proportion of A 's and B 's in a long chain?

3.4.3 Let $\{X_n\}_{n \geq 0}$ be an irreducible positive recurrent HMC with stationary distribution π . Let A be a subset of the state space E and let $\{\tau(k)\}_{k \geq 1}$ be the sequence of return times to A . Show that

$$\lim_{k \uparrow \infty} \frac{\tau(k)}{k} = \frac{1}{\sum_{i \in A} \pi(i)}.$$

(This extends Theorem 3.2.)

3.4.4 Let \mathbf{P} be a transition matrix on $E = \{0, 1, \dots, N - 1, N\}$ with three communicating classes, $\{0\}$, $\{N\}$, and $\{1, \dots, N - 1\}$, respectively recurrent, recurrent, and transient (in particular, 0 and N are closed states). Suppose, moreover, that 0 and N are not isolated states, in the sense that there exists at least one $i \in E$ with $p_{i0} > 0$ and one $j \in E$ with $p_{jN} > 0$. In this situation 0 and N are two absorbing states.

We wish to compute u_k , the mean time to absorption in $\{0, N\}$ when starting from $k \in [1, N - 1]$.

To do this, consider the transition matrix $\hat{\mathbf{P}}$ on E obtained by modifying the first and last rows of \mathbf{P} in such a way that $\hat{p}_{0k} = 1 = \hat{p}_{Nk}$. The matrix $\hat{\mathbf{P}}$ is easily seen to be irreducible, and since E is finite, it has a unique stationary distribution, denoted by $\hat{\pi}$.

Prove that

$$u_k = \frac{1}{\hat{\pi}(0) + \hat{\pi}(N) - 1}.$$

(Use the result of Problem 3.4.3.)

3.4.5 Consider the urn of Ehrenfest (Examples 2.6 and 5.2 of Chapter 2) with a total number of particles $N = 2M$. Let T_0 and T_M be, respectively, the first time when the compartment A becomes empty and the first time when the two compartments have an equal number of particles. Compute the ratio

$$\frac{P_0(T_M < T_0)}{P_M(T_0 < T_M)}.$$

3.4.6 Consider the array of a 's and b 's in part A of the figure below where the letter in a given position has been chosen at random in the set $\{a, b\}$, independently of the other letters, with $P(a) = P(b) = \frac{1}{2}$. Suppose we are interested in counting the occurrences of the special pattern of part B of the figure below without counting overlapping patterns. For instance, with the data shown in part A of the figure below, 2 such patterns are counted; the third one is not counted because it overlaps the second one (part C).

Find an automaton that successively reads the 2-letter columns from left to right, and that has a privileged state $*$ with the following property: the automaton enters (or stays in) state $*$ if and only if it has just discovered a special pattern that is nonoverlapping with a previously discovered one. What is the long-run proportion of nonoverlapping patterns?

$b \ a \ b \ b \ a \ b \ a \ b \ b \ b$
 $b \ a \ a \ a \ b \ b \ a \ a \ a \ a$
A

$b \ . \ b$
 $. \ a \ a$
B

$\boxed{b \ a \ b} \ b \ a \ \boxed{b \ a \ b} \ b \ b$
 $\boxed{b \ a \ a} \ a \ b \ \boxed{b \ a \ a} \ a \ a$
C

YES YES NO

Long Run Behavior

1 Coupling

1.1 Convergence in Variation

Consider an HMC that is irreducible and positive recurrent. In particular, if its initial distribution is the stationary distribution, it keeps the same distribution at all times. The chain is then said to be in the *stationary regime*, or in *equilibrium*, or in *steady state*.

A question arises naturally: What is the long-run behavior of the chain when the initial distribution μ is *arbitrary*? For instance, will it *converge to equilibrium*, and in which sense?

When the HMC is reducible, another type of problem is of interest. Suppose, for instance, that the set of transient states is not empty and that each remaining state is absorbing. One may want to compute the probability of reaching a given absorbing state when the initial state is transient, or the probability of remaining forever in the transient set. In this special case, where all recurrent states are absorbing, the probability of leaving the transient set is exactly the probability of converging. We are dealing here with almost-sure convergence.

For an ergodic HMC, the type of convergence of interest is not almost-sure convergence but convergence in variation of the distribution at time n to the stationary distribution. This type of convergence is relative to a metric structure that we proceed to define.

Definition 1.1. *Distance in Variation*

Let E be a countable space and let α and β be probability distributions on E . The *distance in variation* $d_V(\alpha, \beta)$ between α and β is defined by

$$d_V(\alpha, \beta) = \frac{1}{2}|\alpha - \beta| = \frac{1}{2} \sum_{i \in E} |\alpha(i) - \beta(i)|. \quad (1.1)$$

The distance in variation between two random variables X and Y with values in E and respective distributions $\mathcal{L}(X)$ and $\mathcal{L}(Y)$ is $d_V(\mathcal{L}(X), \mathcal{L}(Y))$, and it is denoted with a slight abuse of notation by $d_V(X, Y)$.

That d_V is indeed a distance is clear.

Lemma 1.1.

Let X and Y be two random variables with values in the same countable space E . Then

$$\sup_{A \subset E} |P(X \in A) - P(Y \in A)| = \sup_{A \subset E} \{P(X \in A) - P(Y \in A)\} = d_V(X, Y). \quad (1.2)$$

Proof. For the first equality observe that for each A there is a B such that $|P(X \in A) - P(Y \in A)| = P(X \in B) - P(Y \in B)$ (take $B = A$ or \bar{A}). For the second equality, write

$$P(X \in A) - P(Y \in A) = \sum_{i \in E} 1_A(i) \{P(X = i) - P(Y = i)\}$$

and observe that the right-hand side is maximal for

$$A = \{i \in E; P(X = i) > P(Y = i)\}$$

Also, for any $A \subset E$,

$$\sum_{i \in E} 1_A(i) \{P(X = i) - P(Y = i)\} + \sum_{i \in E} 1_{\bar{A}}(i) \{P(X = i) - P(Y = i)\} = 0$$

because $\sum_{i \in E} \{P(X = i) - P(Y = i)\} = 0$. For the specific set A above, $P(X = i) - P(Y = i)$ equals $|P(X = i) - P(Y = i)|$ on A , and equals $-|P(X = i) - P(Y = i)|$ on \bar{A} . Therefore, for this particular set A ,

$$\begin{aligned} \sum_{i \in E} 1_A(i) \{P(X = i) - P(Y = i)\} &= \sum_{i \in E} 1_A(i) |P(X = i) - P(Y = i)| \\ &= \sum_{i \in E} 1_{\bar{A}}(i) |P(X = i) - P(Y = i)| \\ &= \frac{1}{2} \sum_{i \in E} |P(X = i) - P(Y = i)|. \quad \square \end{aligned}$$

For two probability distributions α and β on the countable set E , let $\mathcal{D}(\alpha, \beta)$ be the collection of random vectors (X, Y) taking their values in $E \times E$, and with marginal distributions α and β , that is,

$$\alpha = \mathcal{L}(X), \quad \beta = \mathcal{L}(Y). \quad (1.3)$$

Theorem 1.1. Maximal Coincidence

For any $(X, Y) \in \mathcal{D}(\alpha, \beta)$,

$$P(X = Y) \leq 1 - d_V(\alpha, \beta), \quad (1.4)$$

and equality in (1.4) is attained by some pair $(X, Y) \in \mathcal{D}(\alpha, \beta)$, which is then said to *realize maximal coincidence*.

Proof. For arbitrary $A \subset E$,

$$\begin{aligned} P(X \neq Y) &\geq P(X \in A, Y \in \bar{A}) = P(X \in A) - P(X \in A, Y \in A) \\ &\geq P(X \in A) - P(Y \in A), \end{aligned}$$

and therefore

$$P(X \neq Y) \geq \sup_{A \subset E} \{P(X \in A) - P(Y \in A)\} = d_V(\alpha, \beta).$$

To finish the proof, it suffices to construct $(X, Y) \in \mathcal{D}(\alpha, \beta)$ realizing equality. We shall need the following observations (Problem 4.1.1):

$$\frac{1}{2}|\alpha - \beta| = \sum_{i \in E} (\alpha(i) - \beta(i))^+ = \sum_{i \in E} (\beta(i) - \alpha(i))^+ = 1 - \sum_{i \in E} \min(\alpha(i), \beta(i)). \quad (1.5)$$

Let U, Z, V , and W be independent random variables; U takes its values in $\{0, 1\}$, and Z, V, W take their values in E . The distributions of these random variables is given by

$$\begin{aligned} P(U = 1) &= 1 - d_V(\alpha, \beta), \\ P(Z = i) &= \min(\alpha(i), \beta(i)) / (1 - d_V(\alpha, \beta)), \\ P(V = i) &= (\alpha(i) - \beta(i))^+ / d_V(\alpha, \beta), \\ P(W = i) &= (\beta(i) - \alpha(i))^+ / d_V(\alpha, \beta). \end{aligned}$$

Defining

$$X = UZ + (1 - U)V, \quad Y = UZ + (1 - U)W,$$

we have

$$\begin{aligned} P(X = i) &= P(U = 1, Z = i) + P(U = 0, V = i) \\ &= P(U = 1)P(Z = i) + P(U = 0)P(V = i) \\ &= \min(\alpha(i), \beta(i)) + (\alpha(i) - \beta(i))^+ = \alpha(i), \end{aligned}$$

and similarly, $P(Y = i) = \beta(i)$. Therefore, $(X, Y) \in \mathcal{D}(\alpha, \beta)$. Also, $P(X = Y) = P(U = 1) = 1 - d_V(\alpha, \beta)$. □

Example 1.1.

One seeks a pair of $\{0, 1\}$ -valued random variables with prescribed marginals

$$P(X = 1) = a, \quad P(Y = 1) = b,$$

where $a, b \in (0, 1)$, and such that $P(X = Y)$ is maximal. In the notation of the above theory,

$$\alpha = (1 - a, a), \quad \beta = (1 - b, b),$$

and therefore

$$d_V(\alpha, \beta) = |a - b|.$$

Suppose for definiteness that $a \geq b$. The random U, Z, V, W of the construction of Theorem 1.1 have the following distributions.

$$\begin{aligned} P(U = 1) &= 1 - a + b, \\ P(Z = 1) &= \frac{b}{1 - a + b}, \quad P(Z = 0) = \frac{1 - a}{1 - a + b}, \\ V &= 1, \\ W &= 0. \end{aligned}$$

Here $X = UZ + 1 - U, Y = UZ$. ◇

Definition 1.2. *Convergence in Variation*

Let $\{\alpha_n\}_{n \geq 0}$ and β be probability distributions on a countable space E . If $\lim_{n \uparrow \infty} d_V(\alpha_n, \beta) = 0$, the sequence $\{\alpha_n\}_{n \geq 0}$ is said to *converge in variation* to the probability distribution β .

Let $\{X_n\}_{n \geq 0}$ be an E -valued stochastic process. If for some probability distribution π on E , the distribution $\mathcal{L}(X_n)$ of the random variable X_n converges in variation to π , i.e., if

$$\lim_{n \uparrow \infty} \sum_{i \in E} |P(X_n = i) - \pi(i)| = 0, \quad (1.6)$$

then $\{X_n\}_{n \geq 0}$ is said to *converge in variation* to π .

There is some abuse of terminology in the above definition (it is the state random variable, not the process, that converges in variation). However, in this book, such abuse turns out to be harmless and very convenient.

If the process $\{X_n\}_{n \geq 0}$ converges in variation to π , then

$$\lim_{n \uparrow \infty} E[f(X_n)] = \pi(f) \quad (1.7)$$

for all bounded functions $f : E \rightarrow R$, where

$$\pi(f) = \sum_{i \in E} \pi(i) f(i). \quad (1.8)$$

Indeed, if M is an upper bound of $|f|$, then

$$|E[f(X_n)] - \pi(f)| = \left| \sum_{i \in E} f(i) (P(X_n = i) - \pi(i)) \right| \leq M \sum_{i \in E} |P(X_n = i) - \pi(i)|.$$

1.2 The Coupling Method

Coupling is an old idea of Doeblin (1938), revived in Markov-chain theory by the influential work of Griffeath (1975) and Pitman (1974), and brought to fame by Lindvall (1977) who gave a purely probabilistic proof of the renewal theorem (section 3 of the present chapter is devoted to discrete-time renewal theory). The coupling method has a wide range of

applications and the reader is directed to the book (Lindvall, 1992) for additional information and historical comments.

Observe that Definition 1.2 concerns only the marginal distributions of the process, not the process itself. Therefore, if there exists another process $\{X'_n\}_{n \geq 0}$ with $\mathcal{L}(X_n) = \mathcal{L}(X'_n)$ for all $n \geq 0$, and if there exists a third process $\{X''_n\}_{n \geq 0}$ such that $\mathcal{L}(X''_n) = \pi$ for all $n \geq 0$, then (1.6) follows from

$$\lim_{n \uparrow \infty} d_V(X'_n, X''_n) = 0. \tag{1.9}$$

This trivial observation is useful because of the resulting freedom in the choice of $\{X'_n\}$ and $\{X''_n\}$. In particular, one can use dependent versions, and the most interesting case occurs when there exists a finite random time τ such that $X'_n = X''_n$ for all $n \geq \tau$. It follows then, as will be proven in Theorem 1.2, that

$$d_V(X'_n, X''_n) \leq P(\tau > n). \tag{1.10}$$

Finiteness of τ is equivalent to $\lim_{n \uparrow \infty} P(\tau > n) = 0$, and therefore (1.9) is a consequence of (1.10).

The above program can be carried out in a variety of situations, and most notably for ergodic Markov chains. For the time being, the preliminary definitions and results relative to the *coupling method* for proving convergence in variation of ergodic chains will be collected.

Definition 1.3. Coupling

Two stochastic processes $\{X'_n\}_{n \geq 0}$ and $\{X''_n\}_{n \geq 0}$ taking their values in the same countable state space E are said to *couple* if there exists an almost surely *finite* random time τ such that

$$n \geq \tau \Rightarrow X'_n = X''_n. \tag{1.11}$$

The random variable τ is called a *coupling time* of the two processes.

Theorem 1.2. Coupling Inequality

Inequality (1.10) holds for any coupling time τ of $\{X'_n\}_{n \geq 0}$ and $\{X''_n\}_{n \geq 0}$.

Proof. For all $A \subset E$,

$$\begin{aligned} P(X'_n \in A) - P(X''_n \in A) &= P(X'_n \in A, \tau \leq n) + P(X'_n \in A, \tau > n) \\ &\quad - P(X''_n \in A, \tau \leq n) - P(X''_n \in A, \tau > n) \\ &= P(X'_n \in A, \tau > n) - P(X''_n \in A, \tau > n) \\ &\leq P(X'_n \in A, \tau > n) \\ &\leq P(\tau > n). \end{aligned}$$

Inequality (1.10) then follows from Lemma 1.1. □

The framework of the coupling method is now in place. It remains to construct $\{X'_n\}_{n \geq 0}$ and $\{X''_n\}_{n \geq 0}$ that couple and that mimic $\{X_n\}_{n \geq 0}$ and π in the sense that $\mathcal{L}(X'_n) = \mathcal{L}(X_n)$ and $\mathcal{L}(X''_n) = \pi$ for all $n \geq 0$.

2 Convergence to Steady State

2.1 Positive Recurrent Case

The main result concerns ergodic (i.e., irreducible positive recurrent, and *aperiodic*) HMCs.

Theorem 2.1. *Convergence to Steady State*

Let \mathbf{P} be an ergodic transition matrix on the countable state space E . For all probability distributions μ and ν on E ,

$$\lim_{n \uparrow \infty} d_V(\mu^T \mathbf{P}^n, \nu^T \mathbf{P}^n) = 0. \quad (2.1)$$

In particular, if ν is the stationary distribution π ,

$$\lim_{n \uparrow \infty} |\mu^T \mathbf{P}^n - \pi^T| = 0,$$

and with $\mu = \delta_j$, the probability distribution putting all its mass on j ,

$$\lim_{n \uparrow \infty} \sum_{i \in E} |p_{ji}(n) - \pi(i)| = 0.$$

From the discussion preceding Definition 1.3, it suffices to construct two coupling chains with initial distributions μ and ν , respectively.

Theorem 2.2. *Independent Coupling*

Let $\{X_n^{(1)}\}_{n \geq 0}$ and $\{X_n^{(2)}\}_{n \geq 0}$ be two *independent* ergodic HMCs with the same transition matrix \mathbf{P} and initial distributions μ and ν , respectively. Let $\tau = \inf \{n \geq 0; X_n^{(1)} = X_n^{(2)}\}$, with $\tau = \infty$ if the chains never intersect. Then τ is, in fact, almost surely finite. Moreover, the process $\{X'_n\}_{n \geq 0}$ defined by

$$X'_n = \begin{cases} X_n^{(1)} & \text{if } n \leq \tau, \\ X_n^{(2)} & \text{if } n \geq \tau \end{cases} \quad (2.2)$$

is an HMC with transition matrix \mathbf{P} . (See Fig. 4.2.1.)

Proof. Consider the product HMC $\{Z_n\}_{n \geq 0}$ defined by $Z_n = (X_n^{(1)}, X_n^{(2)})$. It takes values in $E \times E$, and the probability of transition from (i, k) to (j, ℓ) in n steps is $p_{ij}(n)p_{k\ell}(n)$. This chain is irreducible. Indeed, since \mathbf{P} is irreducible and *aperiodic*, by Theorem 4.3 of Chapter 2, there exists m such that for all pairs (i, j) and (k, ℓ) , $n \geq m$ implies $p_{ij}(n)p_{k\ell}(n) > 0$. This implies that the period of the product chain is 1, again by Theorem 4.3 of Chapter 2.

Clearly, $\{\pi(i)\pi(j)\}_{(i,j) \in E^2}$ is a stationary distribution for the product chain, where π is the stationary distribution of \mathbf{P} . Therefore, by the stationary distribution criterion, the product chain is positive recurrent. In particular, it reaches the diagonal of E^2 in finite time, and consequently, $P(\tau < \infty) = 1$.

It remains to show that $\{X'_n\}_{n \geq 0}$ given by (2.2) is an HMC with transition matrix \mathbf{P} . This is a consequence of the strong Markov property applied to the product chain. The details are left for the reader (Problem 4.2.1). \square




Figure 4.2.1. Independent coupling

Theorem 2.1 concerns ergodic chains, and aperiodicity is needed there to guarantee that the product chain is irreducible (see Problem 4.2.2). For periodic chains, the situation is different, but the result follows directly from the ergodic case.

Theorem 2.3. Periodic Case

Let \mathbf{P} be an irreducible positive recurrent transition matrix on the countable space E , with period d . Let π be its stationary distribution. If μ is a probability distribution such that $\mu(C_0) = 1$ for some cyclic class C_0 , then

$$\lim_{n \uparrow \infty} \sum_{i \in C_0} |(\mu^T \mathbf{P}^{nd})_i - d\pi(i)| = 0. \tag{2.3}$$

Proof. Consider the restriction of \mathbf{P}^d to C_0 , which is irreducible and aperiodic (see Problem 2.4.7). It is positive recurrent, since it has an invariant measure with finite mass, namely π restricted to C_0 . It remains to show that $d\pi$ restricted to C_0 is a probability distribution, that is, $\pi(C_0) = 1/d$. By the ergodic theorem,

$$\lim_{N \uparrow \infty} \frac{1}{N} \sum_{n=1}^N 1_{\{X_n \in C_0\}} = \pi(C_0),$$

and since $X_n \in C_0$ once every d steps, the left-hand side equals $1/d$. □

2.2 Null Recurrent Case

The last two theorems concern the positive recurrent case. The null recurrent case is a little more difficult.

Theorem 2.4. Orey's Theorem

Let \mathbf{P} be an irreducible null recurrent transition matrix on E . Then for all $i, j \in E$,

$$\lim_{n \uparrow \infty} p_{ij}(n) = 0. \tag{2.4}$$

Proof. The periodic case follows from the aperiodic case by considering the restriction of \mathbf{P}^d to C_0 , an arbitrary cyclic class, and observing that this restriction is also null recurrent. Therefore, \mathbf{P} will be assumed aperiodic.

In this case the product HMC $\{Z_n\}_{n \geq 0} = \{X_n^{(1)}, X_n^{(2)}\}_{n \geq 0}$ defined in the proof of Theorem 2.2 is irreducible and aperiodic. However, it cannot be argued that it is recurrent, although its components are recurrent. One must therefore separate the two possible cases.

If the product chain is transient, its n -step transition probability from (i, i) to (j, j) is $[p_{ij}(n)]^2$, and in view of the discussion of Section 1.2 of Chapter 3, it tends to 0 as $n \rightarrow \infty$, and the result is proven.

If the product chain is recurrent, the coupling argument used in the proof of Theorem 2.1 applies and yields

$$\lim_{n \uparrow \infty} |\mu^T \mathbf{P}^n - \nu^T \mathbf{P}^n| = 0 \quad (2.5)$$

for arbitrary initial distributions μ and ν . Suppose now that for some $i, j \in E$, (2.4) is not true. One can then find a sequence $\{n_k\}_{k \geq 0}$ of integers strictly increasing to ∞ such that

$$\lim_{k \uparrow \infty} p_{ij}(n_k) = \alpha > 0.$$

For fixed $i \in E$ chosen as above, the sequence $(\{p_{is}(n_k), s \in E\})_{k \geq 0}$ of vectors of $[0, 1]^E$ is compact in the topology of pointwise convergence. Therefore (see, however, Theorem 1.10 of the Appendix for an elementary proof), there exists a subsequence $\{m_\ell\}_{\ell \geq 0}$ of integers strictly increasing to ∞ and a vector $\{x_s, s \in E\} \in [0, 1]^E$ such that for all $s \in E$,

$$\lim_{\ell \uparrow \infty} p_{is}(m_\ell) = x_s.$$

Now, $x_j = \alpha > 0$, and therefore $\{x_s, s \in E\}$ is nontrivial. Since $\sum_{s \in E} p_{is}(m_\ell) = 1$, it follows from Fatou's lemma (see Theorem 1.8 of the Appendix), that

$$\sum_{s \in E} x_s \leq 1.$$

In (2.5), take $\mu = \delta_i$ (all the mass is on i) and $\nu^T = \delta_i^T \mathbf{P}$ to obtain, for all $s \in E$,

$$\lim_{\ell \uparrow \infty} |p_{is}(m_\ell) - p_{is}(m_\ell + 1)| = 0. \quad (2.6)$$

But $p_{is}(m_\ell + 1) = \sum_{k \in E} p_{ik}(m_\ell) p_{ks}$. Therefore, by dominated convergence in (2.6), we obtain for all $s \in E$,

$$x_s = \sum_{k \in E} x_k p_{ks}.$$

In other words, $\{x_s, s \in E\}$ is an invariant measure of \mathbf{P} with finite mass, which implies that \mathbf{P} is positive recurrent, a contradiction. Therefore, (2.4) cannot be contradicted. \square

2.3 Thermodynamic Irreversibility

Zermelo's Refutation

We shall now take a small pause, and go back to the Ehrenfest diffusion model. Its fame is due to the insight it gives to the once controversial issue of thermodynamic irreversibility. It is the right time to try to understand this, because we have the notion of recurrence and the theorem of convergence to steady state of ergodic chains.

According to the macroscopic theory of thermodynamics, systems progress in an orderly and irreversible manner towards equilibrium. Consider, for instance, a system of N particles in a box divided in two similar compartments A and B by a fictive membrane. If at the origin of time, all particles are in A , they will rather quickly reorganize themselves, and they will settle to equilibrium, a macroscopic state in which the contents of A and B are thermodynamically equivalent.

Boltzmann claimed that there was an arrow of time in the direction of increasing entropy, and indeed, in the diffusion experiment, equality between the thermodynamic quantities in both compartments corresponds to maximal entropy.

Zermelo, who obviously was not sleeping in the back of the classroom, argued that in view of the time reversibility of the laws of physics, the Boltzmann theory should at least be discussed. Zermelo held a strong position in this controversy. Indeed, there is a famous result of mechanics, Poincaré's recurrence theorem, which implies that in the situation where at time 0 all molecules are in A , then whatever the time T , there will be a subsequent time $t > T$ at which all the molecules will again gather in A . This phenomenon predicted by irrefutable mathematics is, of course, never observed in daily life, where it would imply that the chunk of sugar that one patiently dissolves in one's cup of coffee could escape ingestion by reforming itself at the bottom of the cup.

Convergence vs. Recurrence

Boltzmann's theory was really hurt by this striking and seemingly inescapable argument. Things had to be clarified. Fortunately, Tatiana and Paul Ehrenfest came up with their Markov chain model, and in a sense saved the edifice that Boltzmann had constructed.

Without going into the details, let us just say that the Ehrenfest model is an approximation of the real diffusion phenomenon that is *philosophically correct* from the point of view of statistical mechanics. Also, at first sight, it is subject to Zermelo's attack, presenting both features that the latter found incompatible: an irreversible tendency towards equilibrium, and recurrence. Here the role of Poincaré's recurrence theorem is played by the Markov chain recurrence theorem, stating that an irreducible chain with a stationary distribution visits any fixed state, say 0, infinitely often. As for the irreversible tendency towards equilibrium, one has the theorem of convergence to steady state, according to which the distribution at time n converges to the stationary distribution whatever the initial distribution as n tends to infinity (*stricto sensu* this statement is not true, due to the periodicity of the chain. However, such periodicity is an artifact created by the discretization of time, and it would disappear

in the continuous-time model, or in a slight modification of the discrete-time model). Thus, according to Markov-chain theory, convergence to statistical equilibrium and recurrence are not antagonistic, and we are here at the epicenter of Zermelo's refutation.

One can show that recurrence is *not observable* for states far from $L = \frac{N}{2}$, assuming that N is even. For instance, the average time to reach 0 from state L is

$$\frac{1}{2L} 2^{2L} (1 + O(L)) \quad (2.7)$$

whereas the average time to reach state L from state 0 is less than

$$L + L \log L + O(1). \quad (2.8)$$

(See Chapter III, Section 5 of (Bhattacharya and Waymire, 1990) for a derivation of the above estimates.) With $L = 10^6$ and one unit of mathematical time equal to 10^{-5} second, the return time to equilibrium when compartment A is initially empty is on the order of a second, whereas it would take on the order of

$$\frac{1}{2 \cdot 10^{11}} \times 2^{2^{10^6}} \text{ seconds}$$

to go from L to empty, which is an astronomical time. These numbers teach us not to spend too much time stirring the coffee, or hurry to swallow it for fear of recrystallization of the chunk of sugar. From a mathematical point of view, being in the steady state at a given time does not prevent the chain from being in a rare state, only it is there rarely. The rarity of the state is equivalent to long recurrence times, so long that when there are more than a few particles in the boxes, it would take an astronomical time to witness the effects of Poincaré's recurrence theorem. Note that Boltzmann rightly argued that the recurrence times in Poincaré's theorem are extremely long, but his heuristic arguments failed to convince.

Here is another manifestation of thermodynamic irreversibility in the Ehrenfest model.

Newton's Law of Cooling

Let $g(z, n) = E[z^{X_n}]$ be the generating function of X_n . Using the basic rules of conditional expectation (see Section 7 of Chapter 1), we have

$$\begin{aligned} E[z^{X_{n+1}} | X_n] &= E[z^{X_{n+1}} 1_{\{X_{n+1}=X_n-1\}} | X_n] + E[z^{X_{n+1}} 1_{\{X_{n+1}=X_n+1\}} | X_n] \\ &= E[z^{X_n-1} 1_{\{X_{n+1}=X_n-1\}} | X_n] + E[z^{X_n+1} 1_{\{X_{n+1}=X_n+1\}} | X_n] \\ &= z^{X_n-1} P(X_{n+1} = X_n - 1 | X_n) + z^{X_n+1} P(X_{n+1} = X_n + 1 | X_n), \end{aligned}$$

and therefore, taking into account the dynamics of the Ehrenfest model,

$$\begin{aligned} E[z^{X_{n+1}} | X_n] &= z^{X_n-1} \frac{X_n}{N} + z^{X_n+1} \left(1 - \frac{X_n}{N}\right) \\ &= z \cdot z^{X_n} + \frac{1}{N} (1 - z^2) X_n z^{X_n-1}. \end{aligned}$$

Taking expectations gives

$$g(z, n+1) = zg(z, n) + \frac{1}{N}(1-z^2)E[X_n z^{X_n-1}],$$

that is,

$$g(z, n+1) = zg(z, n) + \frac{1}{N}(1-z^2)g'(z, n).$$

Differentiation of the above identity yields

$$g'(z, n+1) = g(z, n) + zg'(z, n) - \frac{2z}{N}g'(z, n) + \frac{1}{N}(1-z^2)g''(z, n).$$

Letting $z = 1$, we obtain (see Section 5 of Chapter 1)

$$E[X_{n+1}] = 1 + E[X_n] - \frac{2}{N}E[X_n].$$

Supposing N even ($N = 2L$) and rearranging terms, we have

$$E[X_{n+1} - L] = \left(1 - \frac{1}{L}\right)E[X_n - L],$$

and therefore

$$E[X_n - L] = E[X_0 - L] \left(1 - \frac{1}{L}\right)^n.$$

Supposing $P(X_0 = i) = 1$, we then have

$$E \left[\frac{X_n - L}{L} \right] = \left(\frac{i - L}{L} \right) \left(1 - \frac{1}{L} \right)^n.$$

In the kinetic theory of heat, $\frac{X_n - L}{L}$ is interpreted as the temperature difference between state X_n and the "equilibrium" state L . To account for a large number of particles, we let L tend to infinity and make i depend on L in such a way that

$$\lim_{L \rightarrow \infty} \frac{i(L) - L}{L} = \theta(0).$$

The number $\theta(0)$ is interpreted as the initial deviation from the equilibrium temperature in compartment A . Also, the discrete time unit 1 is now interpreted as Δ units of real time. If X_n is the state at real time t , we must then have

$$t = n\Delta.$$

Interpreting $\frac{1}{2L\Delta}$ as the average proportion of the total energy that passes through the membrane per unit of real time, we require that as L tends to infinity, this quantity remains constant, that is,

$$\frac{1}{L\Delta} = \gamma.$$

In particular, Δ tends to zero as L increases to infinity. Therefore, observing that $n = L\gamma t$,

$$\lim_{L \rightarrow \infty} \left(1 - \frac{1}{L}\right)^n = \lim_{L \rightarrow \infty} \left[\left(1 - \frac{1}{L}\right)^L \right]^{\gamma t} = e^{-\gamma t},$$

and finally, with $\theta(t) = \lim_{L \rightarrow \infty} E\left[\frac{X_n - L}{L}\right]$,

$$\theta(t) = \theta(0)e^{-\gamma t}.$$

We conclude this section with a classic reference on thermodynamic irreversibility as probabilists understand it, the article of M. Kac (1947), “Random Walk and the Theory of Brownian Motion”, *American Mathematical Monthly*, 54, 369-391.

2.4 Convergence Rates via Coupling

Knowing that an ergodic Markov chain converges to equilibrium, the next question is, How fast? The first result below is not explicit, but it can be used *in principle* for chains with an infinite number of states.

We recall at this point the meaning of the *o* (small *o*) symbol. It represents a function defined in a neighborhood of zero and such that $\lim_{t \rightarrow 0} \frac{o(t)}{t} = 0$.

Theorem 2.5. Rate of Convergence

Suppose that the coupling time τ in Theorem 1.2 satisfies

$$E[\psi(\tau)] < \infty \tag{2.9}$$

for some nondecreasing function $\psi : \mathbb{N} \rightarrow \mathbb{R}_+$ such that $\lim_{n \uparrow \infty} \psi(n) = \infty$. Then for any initial distributions μ and ν

$$|\mu^T \mathbf{P}^n - \nu^T \mathbf{P}^n| = o\left(\frac{1}{\psi(n)}\right). \tag{2.10}$$

Proof. Since ψ is nondecreasing, $\psi(\tau)1_{\{\tau > n\}} \geq \psi(n)1_{\{\tau > n\}}$, and therefore

$$\psi(n)P(\tau > n) \leq E[\psi(\tau)1_{\{\tau > n\}}].$$

Now,

$$\lim_{n \uparrow \infty} E[\psi(\tau)1_{\{\tau > n\}}] = 0$$

by dominated convergence (see Theorem 3.2 of the Appendix), since $\lim_{n \uparrow \infty} \psi(\tau)1_{\{\tau > n\}} = 0$, by the finiteness of τ , and $\psi(\tau)1_{\{\tau > n\}}$ is bounded by the integrable random variable $\psi(\tau)$. □

Time τ in Theorem 1.2 is the entrance time of the product chain in the diagonal set of $E \times E$. In principle, the distribution of τ can be explicitly computed. However the actual computations are usually difficult when the state space is infinite. The reader will find in (Lindvall, 1992) examples of application of Theorem 2.5 to infinite state space HMCs. When the state space is finite, convergence is exponential (geometric).

Theorem 2.6. *Exponential Convergence Rate of Finite HMCs*

Let \mathbf{P} be an ergodic transition matrix on the finite state space E . Then for any initial distributions μ and ν , one can construct two HMCs $\{X_n\}_{n \geq 0}$ and $\{Y_n\}_{n \geq 0}$ on E with the same transition matrix \mathbf{P} , and the respective initial distributions μ and ν , in such a way that they couple at a finite time τ such that $E[e^{\alpha\tau}] < \infty$ for some positive α .

Proof. See Problem 4.1.4. □

In Chapter 6, we shall say more on convergence rates in the finite state space case. The reader will find in (Lindvall, 1992) examples of application of Theorem 2.5 for infinite state space HMCs.

3 Discrete-Time Renewal Theory

3.1 Renewal Equation

In the analytic approach to Markov chains, the proof of convergence to steady state of an ergodic HMC is a consequence of a result on power series called the *renewal theorem* by the probabilists. This result forms the matter of the present section. However, the renewal theorem will not be used as the essential step towards the convergence theorem, but on the contrary, it will be obtained as a corollary of the latter.

We start with the basic definitions. Let $\{S_n\}_{n \geq 1}$ be an i.i.d sequence of random variables with values in $\bar{\mathbb{N}} = \{1, 2, \dots, +\infty\}$ and with the probability distribution

$$P(S_1 = k) = f_k. \quad (3.1)$$

Define for $n \geq 0$,

$$R_{n+1} = R_n + S_{n+1}, \quad (3.2)$$

where R_0 is an arbitrary random variable with values in $\bar{\mathbb{N}}$ (in particular, $R_0 < \infty$).

Definition 3.1. *Renewal Sequence*

The sequence $\{R_n\}_{n \geq 0}$ is called a *delayed* (by R_0) *renewal sequence* with the *renewal distribution* $\{f_k\}_{k \geq 1}$. If $R_0 \equiv 0$, one speaks of an *undelayed* renewal sequence, or, more simply, of a renewal sequence. If $P(S_1 = \infty) = 0$, the renewal sequence (delayed or not) is called a *proper renewal sequence*, and $\{f_k\}_{k \geq 1}$ is called a *proper renewal distribution*. Otherwise, one speaks of a *defective renewal sequence* and of a *defective renewal distribution*.

The quantity

$$\alpha = P(S_1 = \infty)$$

is the *defect* of the renewal distribution. The random time R_k is the *kth renewal time*, and the sequence $\{S_n\}_{n \geq 1}$ is the *interrenewal sequence*.

Definition 3.2. *Renewal Equation*

With the renewal distribution $\{f_k\}_{k \geq 1}$ is associated the *renewal equation*

$$u_n = v_n + \sum_{k=1}^n f_k u_{n-k} \quad (3.3)$$

(for $n = 0$, this reduces to $u_0 = v_0$). The sequence $\{u_n\}_{n \geq 0}$ is the *unknown sequence*, and $\{v_n\}_{n \geq 0}$ is the *data*, a sequence of real numbers such that

$$\sum_{k=0}^{\infty} |v_k| < \infty. \quad (3.4)$$

Since u_n can be computed recursively as a function of $u_0, \dots, u_{n-1}, v_0, \dots, v_n$, a solution of the renewal equation always exists and is unique.

Example 3.1. *Lifetime of a Defective Renewal Sequence*

Define the lifetime L of a defective renewal sequence by

$$L = \inf \{R_k; k \geq 0, S_{k+1} = \infty\}.$$

It is the last renewal time at finite distance. We shall see that $u_n = P(L > n)$ satisfies a renewal equation. For this, write

$$1_{\{L > n\}} = 1_{\{L > n\}} 1_{\{S_1 > n\}} + 1_{\{L > n\}} 1_{\{S_1 \leq n\}}.$$

Observe that $\{L > n, S_1 > n\} = \{n < S_1 < \infty\}$. Also, denoting by \hat{L} the lifetime associated with the renewal sequence $\{R_{n+1} - R_1\}_{n \geq 0}$, we have the set identity $\{L > n, S_1 \leq n\} = \{\hat{L} > n - S_1, S_1 \leq n\}$. Therefore,

$$P(L > n) = P(n < S_1 < \infty) + P(\hat{L} > n - S_1, S_1 \leq n).$$

Now, L and \hat{L} have the same distribution, and \hat{L} is independent of S_1 . Therefore,

$$P(\hat{L} > n - S_1, S_1 \leq n) = \sum_{k=1}^n P(\hat{L} > n - k)P(S_1 = k) = \sum_{k=1}^n P(L > n - k)P(S_1 = k).$$

This shows that u_n satisfies the renewal equation with data $v_n = P(n < S_1 < \infty)$. \diamond

Definition 3.3. *Fundamental Solution*

Define the *Dirac sequence* $\{\delta_n\}_{n \geq 0}$ by $\delta_0 = 1, \delta_n = 0$ for $n \geq 1$. When the data is the Dirac sequence, the renewal equation is called the *basic renewal equation*, and its solution the *fundamental solution*.

The fundamental solution will be denoted by $\{h_n\}_{n \geq 0}$, and therefore $h_0 = 1$, and for $n \geq 1$,

$$h_n = \sum_{k=1}^n f_k h_{n-k}. \quad (3.5)$$

The fundamental solution has a very simple interpretation. Indeed, h_n is the probability that n is a renewal time (we then say, for short, “ n is renewal”). It suffices to show that $u_n = P(n \text{ is renewal})$ is the unique solution of the basic renewal equation. Clearly, $u_0 = 1$. Also,

$$\begin{aligned}
 P(n \text{ is renewal}) &= \sum_{k=0}^{n-1} P(n \text{ is renewal, last renewal strictly before } n \text{ is } k) \\
 &= \sum_{i=0}^{\infty} \sum_{k=0}^{n-1} P(S_{i+1} = n - k, k = R_i) \\
 &= \sum_{i=0}^{\infty} \sum_{k=0}^{n-1} P(S_{i+1} = n - k) P(k = R_i) \\
 &= \sum_{k=0}^{n-1} P(S_1 = n - k) \left(\sum_{i=0}^{\infty} P(k = R_i) \right) \\
 &= \sum_{k=0}^{n-1} P(S_1 = n - k) P(k \text{ is renewal}) \\
 &= \sum_{k=0}^{n-1} u_k f_{n-k} = \sum_{k=1}^n f_k u_{n-k}.
 \end{aligned}$$

Therefore,

$$h_k = P(k \text{ is a renewal time}). \quad (3.6)$$

In particular, if v_n is the number of renewal times R_k in the interval $[0, n]$, then

$$v_n = \sum_{k=0}^n h_k. \quad (3.7)$$

We now introduce a definition and a convenient notation. The *convolution* of two real sequences $\{x_n\}_{n \geq 0}$ and $\{y_n\}_{n \geq 0}$ is the real sequence $\{z_n\}_{n \geq 0}$ defined by

$$z_n = \sum_{k=0}^n x_k y_{n-k}.$$

This is written for short as $z = x * y$.

Theorem 3.1. Solution of the Renewal Equation

The renewal equation (3.3) has a unique solution

$$u = h * v. \quad (3.8)$$

Proof. Existence and uniqueness have already been observed. To check that the announced solution is correct, write the renewal equation as $u = v + f * u$ (with $f_0 = 0$) and the fundamental equation as $h = \delta + f * h$. Inserting (3.8) into the renewal equation gives $h * v = v + f * (h * v)$ which is indeed true, since the right-hand side is $v + (f * h) * v = v + (h - \delta) * v = v + h * v - \delta * v$, that is, $h * v$, because $\delta * v = v$. \square

Example 3.2. *Geometric Interrenewal Times*

When the the distribution of the typical interrenewal time is geometric, i.e., for $k \geq 1$,

$$P(S_1 = k) = p(1 - p)^{k-1},$$

the fundamental solution is given by $h_0 = 1$, and

$$h_n = p,$$

for $n \geq 1$, as can be readily checked. The solution of the general renewal equation is then

$$u_n = v_n + p(v_0 + \cdots + v_{n-1}).$$

One observes in this particular case that since $\lim_{n \uparrow \infty} v_n = 0$ in view of assumption (3.4),

$$\lim_{n \uparrow \infty} u_n = p \sum_{k=0}^{\infty} v_k = \frac{\sum_{k \geq 0} v_k}{\sum_{k \geq 1} k f_k}.$$

This result will be generalized by the renewal theorem. ◇

3.2 Renewal Theorem

Definition 3.4. Lattice Distributions

The renewal distribution $\{f_k\}_{k \geq 1}$ is called *lattice* (resp., *nonlattice*) if $d \stackrel{\text{def}}{=} \gcd\{k ; k \geq 1, f_k > 0\} > 1$ (resp., $= 1$); the integer d is called the *span* of the renewal distribution.

The main result can now be stated.

Theorem 3.2. The Renewal Theorem

Let $\{f_k\}_{k \geq 1}$ be a *nonlattice* and *proper* renewal distribution. For the unique solution of the renewal equation with data satisfying assumption (3.4),

$$\lim_{n \uparrow \infty} u_n = \frac{\sum_{k \geq 0} v_k}{\sum_{k \geq 1} k f_k}, \quad (3.9)$$

where the ratio on the right-hand side is 0 if $\sum_{k \geq 1} k f_k = \infty$.

Proof. In two parts.

A. Assume the result true for the fundamental solution, that is,

$$\lim_{n \uparrow \infty} h_n = \frac{1}{\sum_{k \geq 1} k f_k} \stackrel{\text{def}}{=} h_{\infty}. \quad (3.10)$$

From expression (3.8) of the solution in terms of the fundamental solution, we obtain

$$\sum_{k=0}^n (h_{n-k} - h_{\infty}) v_k = u_n - h_{\infty} \sum_{k=0}^n v_k.$$

The result follows if we can prove that the left-hand side of the above equality converges to 0 as $n \rightarrow \infty$. Indeed, with $g(n, k) = (h_{n-k} - h_\infty)v_k 1_{\{k \leq n\}}$, we have for fixed k , $\lim_{n \uparrow \infty} g(n, k) = 0$, and $|g(n, k)| \leq |v_k|$, where $\sum_{k \geq 0} |v_k| < \infty$. Therefore, by dominated convergence for series (see Theorem 1.6 of the Appendix), $\lim_{n \uparrow \infty} \sum_{k \geq 0} g(n, k) = \sum_{k \geq 0} \lim_{n \uparrow \infty} g(n, k) = 0$.

B. It remains to prove (3.10). For this, introduce a Markov chain with state space $E = \mathbb{N}$ if the support of $\{f_k\}_{k \geq 1}$ is unbounded, and state space $E = \{0, \dots, M - 1\}$ if $M < \infty$ is the largest value of S_1 . We suppose for definiteness that $E = \mathbb{N}$. The nonzero entries of the transition matrix are

$$p_{i, i-1} = 1, \quad i \geq 1,$$

$$p_{0i} = f_{i+1}, \quad i \geq 0.$$

The corresponding transition graph is shown in Figure 4.3.1. Note that this is the transition graph of the forward recurrence time HMC $\{X_n\}_{n \geq 0}$ defined by

$$X_n = \inf \{R_k ; R_k \geq n\} - n.$$




Figure 4.3.1. Transition graph of the forward recurrence chain

This chain is clearly irreducible. The distribution of the return time to state 0 is

$$P_0(T_0 = n) = f_n.$$

Event $\{T_0 = n\}$ implies event $\{X_n = 0\}$, and therefore $P_0(X_n = 0) \geq P_0(T_0 = n)$. Consequently, the set $A = \{n \geq 1; p_{00}(n) > 0\}$ contains the set $B = \{n \geq 1; f_n > 0\}$, and therefore the g.c.d. of A is smaller than or equal to the g.c.d. of B . Therefore, the g.c.d. of A equals 1, that is, the chain is aperiodic.

Since the renewal distribution is assumed proper, we have $P_0(T_0 < \infty) = \sum_{n \geq 1} f_n = 1$, and therefore the chain is recurrent. If $E_0[T_0] < \infty$, it is ergodic, and then

$$\lim_{n \uparrow \infty} p_{00}(n) = \pi_0 = \frac{1}{E_0\{T_0\}}.$$

If the chain is not ergodic but only null recurrent, then $\lim_{n \uparrow \infty} p_{00}(n) = 0$ by Orey's theorem. In both cases, since $E_0[T_0] = \sum_{k \geq 1} k f_k$,

$$\lim_{n \uparrow \infty} p_{00}(n) = \frac{1}{\sum_{k \geq 1} k f_k}.$$

The proof of (3.9) is complete because $p_{00}(n) = P(n \text{ is renewal}) = h_n$. \square

Theorem 3.3. Lattice Renewal Theorem

Under the same conditions as in Theorem 3.2, except that the span d of the renewal distribution is now strictly greater than 1, the solution of the renewal equation (3.3) with data satisfying (3.4) satisfies, for all $r \in [0, d - 1]$,

$$\lim_{N \uparrow \infty} u_{r+Nd} = d \frac{\sum_{k \geq 0} v_{r+kd}}{\sum_{k \geq 1} k f_k}. \quad (3.11)$$

Proof. Observe that when $\{f_k\}_{k \geq 1}$ is proper and lattice with span d , the distribution $\{f_{Nd}\}_{N \geq 1}$ is proper and nonlattice. On the other hand, the renewal equation (3.3) splits into d renewal equations. The r th one ($r \in [0, d - 1]$) is

$$u_{r+Nd} = v_{r+Nd} + \sum_{\ell=1}^N f_{\ell d} u_{r+Nd-\ell d},$$

where N is the time variable. The renewal theorem can be applied to each one, and we obtain (3.11) after observing that

$$\sum_{N=1}^{\infty} N f_{Nd} = \frac{1}{d} \sum_{N=1}^{\infty} N d f_{Nd} = \frac{1}{d} \sum_{k=1}^{\infty} k f_k. \quad \square$$

3.3 Defective Renewal Sequences

Theorem 3.4. Defective Renewal Theorem

Suppose that the renewal distribution is defective, and that the data sequence of the renewal equation is nonnegative and satisfies (instead of (3.4))

$$\lim_{n \uparrow \infty} v_n = v_{\infty} < \infty. \quad (3.12)$$

The solution of the renewal equation then satisfies

$$\lim_{n \uparrow \infty} u_n = \frac{v_{\infty}}{\alpha}, \quad (3.13)$$

where $\alpha = P(S_1 = \infty)$ is the defect of the renewal distribution.

Proof. The forward recurrence HMC in the proof of Theorem 3.2 now has $E = \mathbb{N} \cup \{+\infty\}$ for state space and ∞ is a closed state. All states besides $+\infty$ are transient. In particular, the average number of visits to 0 is finite:

$$v_\infty = \sum_{k=0}^{\infty} h_k < \infty.$$

From the expression of the solution

$$u_n = \sum_{k=0}^n h_k v_{n-k},$$

we therefore obtain by the dominated convergence for series (see Theorem 1.6 of the Appendix),

$$\lim_{n \uparrow \infty} u_n = \left(\sum_{k=0}^{\infty} h_k \right) v_\infty = v_\infty v_\infty.$$

Now, the probability of n visits to 0 is $(1 - \alpha)^{n-1} \alpha$, and therefore the average number of visits to 0 is $v_\infty = \frac{1}{\alpha}$. \square

Theorem 3.5. Exponential Decay of Defective Renewal Sequences

Suppose that the renewal distribution is nonlattice and defective, that there exists $\gamma > 1$ such that

$$\sum_{k=0}^{\infty} \gamma^k f_k = 1, \quad (3.14)$$

and that the data sequence satisfies

$$\sum_{k=0}^{\infty} \gamma^k |v_k| < \infty. \quad (3.15)$$

The solution of the renewal equation then satisfies

$$\lim_{n \uparrow \infty} \gamma^n u_n = \frac{\sum_{k=0}^{\infty} \gamma^k v_k}{\sum_{k=0}^{\infty} k \gamma^k f_k}. \quad (3.16)$$

Proof. Observe that if we define $\tilde{f}_n = \gamma^n f_n$, $\tilde{v}_n = \gamma^n v_n$, $\tilde{u}_n = \gamma^n u_n$, then

$$\tilde{u}_n = \tilde{v}_n + \sum_{k=1}^n \tilde{f}_k \tilde{u}_{n-k}.$$

This renewal equation is nonlattice and proper, and therefore the announced result follows from the renewal theorem. \square

Remark 3.1.

A consequence of (3.14) is the exponential decay of the renewal distribution. This shows in particular that (3.14) is an assumption that is *not always satisfied*. \diamond

Example 3.3. *Convergence Rate in the Defective Renewal Theorem*

The situation is that of Theorem 3.4, where in addition the renewal distribution is assumed nonlattice, and moreover, (3.14) is true for some $\gamma > 1$. We seek to understand how u_n tends to u_∞ . For this we define $\hat{u}_n = u_n - u_\infty$. Rewriting the renewal equation for u_n as

$$u_n - u_\infty = v_n - u_\infty + \sum_{k=1}^n f_k(u_{n-k} - u_\infty) + u_\infty \sum_{k=1}^n f_k,$$

we see that \hat{u}_n satisfies the renewal equation with data

$$\hat{v}_n = v_n - u_\infty P(S_1 > n).$$

We can therefore apply Theorem 3.5 to obtain, after rearrangement,

$$\lim_{n \uparrow \infty} \gamma^n (u_n - \frac{v_\infty}{\alpha}) = \frac{1}{\gamma} \left\{ \frac{\sum_{k=0}^\infty \gamma^k v_k}{\sum_{k=0}^\infty \gamma^k P(S_1 > k)} - \frac{v_\infty}{P(S_1 = \infty)} \right\}. \quad \diamond$$

An *excessive* renewal equation is one for which $\sum_{k=1}^\infty f_n > 1$. Theorem 3.5 then has an obvious counterpart. Note however that in the excessive case (3.14) *always* has a solution γ , and of course it is in $(0, 1)$.

Example 3.4. *Mothers and Daughters*

At each time $n \in \mathbb{Z}$, an average number u_n of daughters is born. Each of them gives birth independently of the other women. The average number of daughters of any given woman in the k th year of her life, $k \geq 1$, is f_k . At time 0 the population has $\alpha(i)$ women of age i . Expressing that u_n is the sum of v_n , the average number of daughters born at time n from mothers born at or before time 0, and of r_n , the average number of daughters born at time n from mothers born strictly after time 0 and up to time n , we obtain the renewal equation with data sequence

$$v_n = \sum_{i=0}^\infty \alpha(i) f_{n+i}.$$

In this context, the renewal equation is known as the *Lotka–Volterra equation*. Denote by

$$\rho = \sum_{k=1}^\infty f_n$$

the average number of daughters of any given woman, and assume that this number is positive and finite. Assume also that it is different from 1. Assume that γ defined by (3.14) exists and that the renewal distribution is nonlattice. Denoting by C the right-hand side of (3.16),

$$\lim_{n \uparrow \infty} \gamma^n u_n = C.$$

Note that $\gamma < 1$ if $\rho > 1$, and $\gamma > 1$ if $\rho < 1$. The first case corresponds to exponential explosion, whereas the second case is that of exponential extinction. \diamond

4 Regenerative Processes

4.1 Renewal Equation of a Regenerative Process

In the introductory lines of the previous section, we mentioned that the renewal theorem can be proven independently of the basic results of convergence to steady state, and that it could even be used to prove such convergence results. Therefore, it seems that the probabilistic approach to convergence gives a negligible status to the renewal theorem, which enjoys a central position in the analytic approach. However, the renewal theorem remains indispensable to prove convergence to equilibrium of stochastic processes of a more general nature than homogeneous Markov chains, namely regenerative processes. The common feature that such processes share with the homogeneous recurrent Markov chains is the existence of regenerative cycles.

Definition 4.1. *Regenerative Process*

Let $\{Z_n\}_{n \geq 0}$ be a stochastic process with values in an arbitrary state space E and let $\{R_n\}_{n \geq 0}$ be a delayed renewal sequence. The process $\{Z_n\}_{n \geq 0}$ is said to be regenerative with respect to the renewal sequence $\{R_n\}_{n \geq 0}$ if for all $k \geq 0$, $\{Z_{n+R_k}\}_{n \geq 0}$ is independent of R_0, S_1, \dots, S_k and has the same distribution as $\{Z_{n+R_0}\}_{n \geq 0}$.

Note that the definition does not require that $\{Z_{n+R_k}\}_{n \geq 0}$ be independent of $\{Z_n, n \in [0, R_k - 1]\}$, although in many examples this is satisfied. The freedom resulting from the relaxed conditions of Definition 4.1 can be very useful.

Example 4.1. *Recurrent Markov Chains*

Let $\{X_n\}_{n \geq 0}$ be an irreducible recurrent HMC, with arbitrary initial distribution. Let $\{R_n\}_{n \geq 0}$ be the successive *hitting* times of state 0. The regenerative cycle theorem (Theorem 7.4 of Chapter 2) tells us that $\{X_n\}_{n \geq 0}$ is regenerative with respect to $\{R_n\}_{n \geq 0}$. \diamond




Figure 4.4.1. A sample path of the reliability process

Example 4.2. *Reliability*

Let $\{U_n\}_{n \geq 1}$ and $\{V_n\}_{n \geq 1}$ be two independent i.i.d sequences of positive integer-valued random variables. Define the sequence $\{S_n\}_{n \geq 1}$ by $S_n = U_n + V_n$, and let $\{R_n\}_{n \geq 0}$ be the associated nondelayed renewal sequence ($R_0 \equiv 0$). Define a $\{0, 1\}$ -valued process $\{Z_n\}_{n \geq 0}$ as in Figure 4.4.1. Clearly, $\{Z_n\}_{n \geq 0}$ is a regenerative process with respect to $\{R_n\}_{n \geq 0}$. \diamond

Regenerative processes generate renewal equations and are the main motivation for the study of such equations. For instance, if $f : E \rightarrow \mathbb{R}$ is a nonnegative function, and if $\{Z_n\}_{n \geq 0}$ is an E -valued process regenerative with respect to the nondelayed renewal sequence $\{R_n\}_{n \geq 0}$, then the sequence $\{u_n\}_{n \geq 0}$, where $u_n = E[f(Z_n)]$, satisfies a renewal equation. Indeed,

$$E[f(Z_n)] = E[f(Z_n)\mathbf{1}_{\{n < S_1\}}] + E[f(Z_n)\mathbf{1}_{\{n \geq S_1\}}],$$

and, setting $\tilde{Z}_n = Z_{n+S_1}$, we have

$$\begin{aligned} E[f(Z_n)\mathbf{1}_{\{n \geq S_1\}}] &= E[f(\tilde{Z}_{n-S_1})\mathbf{1}_{\{n \geq S_1\}}] \\ &= \sum_{k=1}^{\infty} E[f(\tilde{Z}_{n-S_1})\mathbf{1}_{\{n \geq S_1\}}\mathbf{1}_{\{S_1=k\}}] \\ &= \sum_{k=1}^n E[f(\tilde{Z}_{n-k})\mathbf{1}_{\{S_1=k\}}] \\ &= \sum_{k=1}^n E[f(\tilde{Z}_{n-k})]P(S_1 = k) \\ &= \sum_{k=1}^n E[f(Z_{n-k})]P(S_1 = k), \end{aligned}$$

where the independence of S_1 and $\{\tilde{Z}_n\}_{n \geq 0}$, as well as the assumption of equidistribution of $\{\tilde{Z}_n\}_{n \geq 0}$ and $\{Z_n\}_{n \geq 0}$, have been taken into account. Therefore,

$$E[f(Z_n)] = E[f(Z_n)\mathbf{1}_{\{n < S_1\}}] + \sum_{k=1}^n E[f(Z_{n-k})]P(S_1 = k),$$

which is precisely the renewal equation with data

$$v_n = E[f(Z_n)\mathbf{1}_{\{n < S_1\}}].$$

4.2 Regenerative Theorem

Observe that

$$\sum_{n=0}^{\infty} |v_n| = \sum_{n=0}^{\infty} |E[f(Z_n)\mathbf{1}_{\{n < S_1\}}]| \leq E\left[\sum_{n=0}^{\infty} |f(Z_n)\mathbf{1}_{\{n < S_1\}}|\right] = E\left[\sum_{n=0}^{S_1-1} |f(Z_n)|\right].$$

Therefore, by the renewal theorem, we have the following:

Theorem 4.1. The Regenerative Theorem

Let $\{Z_n\}_{n \geq 0}$ be a nondelayed ($R_0 = 0$) regenerative process and let $f : E \rightarrow \mathbb{R}$ be such that

$$E\left[\sum_{n=0}^{S_1-1} |f(Z_n)|\right] < \infty. \quad (4.1)$$

If the distribution of S_1 is proper and nonlattice, then

$$\lim_{n \uparrow \infty} E[f(Z_n)] = \frac{E \left[\sum_{n=0}^{S_1-1} f(Z_n) \right]}{E[S_1]} \tag{4.2}$$

The latter formula is called *Smith's formula*.

Example 4.3. Reliability Formula

This is a continuation of Example 4.2. We assume that $S_1 = U_1 + V_1$ is proper and nonlattice. Applying the regenerative theorem with $f(z) = 1_{\{0\}}(z)$, and assuming $E[U_1] < \infty$, we find that

$$\lim_{n \uparrow \infty} P(Z_n = 0) = \frac{E[U_1]}{E[U_1] + E[V_1]}$$

since $E[f(Z_n)] = E[1_{\{0\}}(Z_n)] = P(Z_n = 0)$, and $\sum_{n=0}^{S_1-1} 1_{\{0\}}(Z_n) = U_1$. ◇

Example 4.4. The Bus Paradox

Consider the renewal sequence of Definition 3.1, with $R_0 = 0$. Define for each $n \geq 0$ the backward recurrence time B_n and the forward recurrence time F_n by

$$B_n = n - L_n, \quad F_n = N_n - n,$$


where

$$L_n = \sup\{R_k; k \geq 0, R_k \leq n\}$$

and

$$N_n = \inf\{R_k; k > 0, R_k > n\}.$$

In particular, if $n = R_m$ for some m , then $B_n = 0$ and $F_n = R_{m+1} - R_m = S_{m+1}$. Observe that $F_n \geq 1$ for all $n \geq 0$. Also, if $n \in [R_m, R_{m+1})$, then $B_n + F_n = S_{m+1}$.



Here, index k is random

Figure 4.4.2. Backward and forward recurrence times

The regenerative theorem with $Z_n = (B_n, F_n)$ and $f(Z_n) = 1_{\{B_n=i\}} 1_{\{F_n=j\}}$ gives, provided that the distribution of S_1 is proper and nonlattice,

$$\lim_{n \uparrow \infty} P(B_n = i, F_n = j) = \frac{P(S_1 = i + j)}{E[S_1]} \tag{4.3}$$

Indeed the sum $\sum_{n=0}^{S_1-1} 1_{\{B_n=i, F_n=j\}}$ has at most one non-zero term, in which case it is equal to 1. For this term, say corresponding to the index $n = n_0$, $B_{n_0} + F_{n_0} = S_1 = i + j$. Therefore the sum equals to $1_{\{S_1=i+j\}}$.

Summing up (4.3) from $j = 1$ to ∞ , and recalling that $F_n \geq 1$, one obtains

$$\lim_{n \uparrow \infty} P(B_n = i) = \frac{P(S_1 > i)}{E[S_1]}. \quad (4.4)$$

Similarly, for the forward recurrence time,

$$\lim_{n \uparrow \infty} P(F_n = j) = \frac{P(S_1 \geq j)}{E[S_1]}.$$

The roles of B_n and F_n are not symmetric. To restore symmetry, one must consider B_n and $F'_n = F_n - 1$ (recall that $F_n \geq 1$). Then

$$\lim_{n \uparrow \infty} P(F'_n = j) = \frac{P(S_1 > j)}{E[S_1]}.$$

Since $B_n + F_n = S_m$ for some (random) m determined by the condition $n \in [R_m, R_{m+1})$, one might expect that $P(B_n + F_n = k) = P(S_1 = k)$. But this is in general false, and constitutes the apparent *paradox of recurrence times*, also called the *bus paradox* (see Problem 4.4.2). It is true that $P(B_n + F_n = k) = P(S_m = k)$, but m is random, and therefore there is no reason why S_m should have the same distribution as S_1 . As a matter of fact,

$$\begin{aligned} \lim_{n \uparrow \infty} P(B_n + F_n = k) &= \lim_{n \uparrow \infty} \sum_{\substack{i,j \\ i+j=k}} P(B_n = i, F_n = j) \\ &= \sum_{\substack{i,j \\ i+j=k}} \frac{P(S_1 = i + j)}{E[S_1]} = \frac{k P(S_1 = k)}{E[S_1]}. \end{aligned} \quad \diamond$$

Theorem 4.2. *Delayed Regenerative Theorem*

Let $\{Z_n\}_{n \geq 0}$ be a possibly delayed regenerative process (recall, however, that $R_0 < \infty$). Let $f : E \rightarrow \mathbb{R}$ be such that

$$\lim_{n \uparrow \infty} E[f(Z_n) 1_{\{n < R_0\}}] = 0 \quad (4.5)$$

and

$$E \left[\sum_{k=R_0}^{R_1-1} |f(Z_k)| \right] < \infty. \quad (4.6)$$

Then, if the renewal distribution is proper and nonlattice,

$$\lim_{n \uparrow \infty} E[f(Z_n)] = \frac{E \left[\sum_{k=R_0}^{R_1-1} f(Z_k) \right]}{E[S_1]}. \quad (4.7)$$

Proof. It suffices to show that the limit of $E[f(Z_n)1_{n \geq R_0}]$ equals the right-hand side of (4.7). Introduce $\{\tilde{Z}_n\}_{n \geq 0} = \{Z_{n+R_0}\}_{n \geq 0}$, and observe that this process is an undelayed regenerative process with respect to $\{R_n - R_0\}_{n \geq 1}$ that is proper and nonlattice. We have

$$E[f(Z_n)1_{\{n \geq R_0\}}] = E[f(\tilde{Z}_{n-R_0})1_{\{n \geq R_0\}}] = \sum_{k=0}^n E[f(\tilde{Z}_{n-k})P(R_0 = k)].$$

By the nondelayed version of the regenerative theorem, we have that

$$\lim_{n \uparrow \infty} E[f(\tilde{Z}_n)] = \frac{E\left[\sum_{k=R_0}^{R_1-1} f(Z_k)\right]}{E[R_1 - R_0]},$$

and therefore, by dominated convergence for series (see Theorem 1.6 of the Appendix),

$$\lim_{n \uparrow \infty} \sum_{k=0}^n E[f(\tilde{Z}_{n-k})P(R_0 = k)] = \frac{E\left[\sum_{k=R_0}^{R_1-1} f(Z_k)\right]}{E[R_1 - R_0]}. \quad \square$$

A useful case where (4.5) is satisfied is that of f bounded (use dominated convergence).

5 Life Before Absorption

5.1 Infinite Sojourns

The results obtained so far concern irreducible recurrent HMCs. They must be completed by the study of chains with several communication classes.

As we learned in Chapter 3, the state space E can be decomposed as

$$E = T + \sum_j R_j,$$

where R_1, R_2, \dots are the disjoint recurrent classes and T is the collection of transient states. (Note that the number of recurrent classes as well as the number of transient states may be infinite.) From the discussion in subsection 1.3 of Chapter 3, the transition matrix can be block-partitioned as

$$\mathbf{P} = \begin{matrix} & \begin{matrix} R_1 & R_2 & \cdots & T \end{matrix} \\ \begin{matrix} R_1 \\ R_2 \\ \vdots \\ T \end{matrix} & \begin{pmatrix} \mathbf{P}_1 & 0 & 0 & 0 \\ 0 & \mathbf{P}_2 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ B(1) & B(2) & \cdots & \mathbf{Q} \end{pmatrix} \end{matrix}, \quad (5.1)$$

or in condensed notation,

$$\mathbf{P} = \begin{pmatrix} D & 0 \\ B & \mathbf{Q} \end{pmatrix}. \quad (5.2)$$

The form (5.1) accounts for the fact that one cannot go from a state in a given recurrent class to any state not belonging to this recurrent class. In other words, a recurrent class is closed.

A typical question is, What is the probability of being absorbed by a given recurrent class when starting from a given transient state? This kind of problem was already addressed in Chapter 2 to introduce first-step analysis (see, for instance, the tennis Example 3.2 and the cat-eats-rat-eats-cheese Example 3.3 of Chapter 2). It led to systems of linear equations with boundary conditions, for which the solution was unique, due to the finiteness of the state space. With an infinite state space, the uniqueness issue cannot be overlooked, and the absorption problem will be reconsidered with this in mind, and also with the intention of finding general matrix-algebraic expressions for the solutions. Another phenomenon not manifesting itself in the finite case is the possibility, when the set of transient states is infinite, of never being absorbed by the recurrent set. We shall consider this problem first, and then proceed to derive the distribution of the time to absorption by the recurrent set, and the probability of being absorbed by a given recurrent class.

Let A be a subset of the state space E (typically the set of transient states, but not necessarily). We aim at computing for any initial state $i \in A$ the probability of remaining forever in A ,

$$v(i) = P_i(X_r \in A; r \geq 0). \quad (5.3)$$

Defining

$$v_n(i) = P_i(X_1 \in A, \dots, X_n \in A),$$

we have, by monotone sequential continuity,

$$\lim_{n \uparrow \infty} \downarrow v_n(i) = v(i).$$

But for $j \in A$, $P_i(X_1 \in A, \dots, X_{n-1} \in A, X_n = j) = \sum_{i_1 \in A} \dots \sum_{i_{n-1} \in A} P_{ii_1} \dots P_{i_{n-1}j}$ is the general term $q_{ij}(n)$ of the n th iterate of \mathbf{Q} , the restriction of \mathbf{P} to the set A . Therefore,

$$v_n(i) = \sum_{j \in A} q_{ij}(n).$$

The latter equality reads

$$v_n = \mathbf{Q}^n \mathbf{1}_A, \quad (5.4)$$

where $\mathbf{1}_A$ is the column vector indexed by A with all entries equal to 1. From this equality we obtain

$$v_{n+1} = \mathbf{Q}v_n, \quad (5.5)$$

and by dominated convergence $v = \mathbf{Q}v$. Moreover, $\mathbf{0}_A \leq v \leq \mathbf{1}_A$, where $\mathbf{0}_A$ is the column vector indexed by A with all entries equal to 0. The above result can be refined.

Theorem 5.1. *Infinite Sojourn*

The vector v is the maximal solution of

$$v = \mathbf{Q}v, \quad (5.6)$$

$$\mathbf{0}_A \leq v \leq \mathbf{1}_A. \quad (5.7)$$

Moreover, either $v = \mathbf{0}_A$ or $\sup_{i \in A} v(i) = 1$.

Proof. Only maximality and the last statement remain to be proven. To prove maximality consider a vector u indexed by A such that $u = Qu$ and $0_A \leq u \leq 1_A$. Iteration of $u = Qu$ yields $u = Q^n u$, and $u \leq 1_A$ implies that $Q^n u \leq Q^n 1_A = v_n$. Therefore, $u \leq v_n$, which gives $u \leq v$ by passage to the limit.

To prove the last statement of the theorem, let $c = \sup_{i \in A} v(i)$. From $v \leq c1_A$, we obtain $v \leq cv_n$ as above, and therefore, at the limit, $v \leq cv$. This implies either $v = 0_A$ or $c = 1$. \square

Equation $v = Qv$ in (5.7) reads

$$v(i) = \sum_{j \in A} p_{ij} v(j) \quad (i \in A). \tag{5.8}$$

First-step analysis gives (5.8) as a *necessary* condition. However, it does not help to determine which solution of (5.8) to choose, in case there are several.

Example 5.1. *Random Walk Reflected at 0*

The transition matrix of the random walk on \mathbb{N} with a reflecting barrier at 0,

$$P = \begin{pmatrix} 0 & 1 & & & \\ q & 0 & p & & \\ & q & 0 & p & \\ & & q & 0 & p \\ & & & \ddots & \ddots \end{pmatrix},$$

where $p \in (0, 1)$, is clearly irreducible. Intuitively, if $p > q$, there is a drift to the right, and one expects the chain to be transient. This will be proven formally by showing that the probability $v(i)$ of never visiting state 0 when starting from state $i \geq 1$ is strictly positive. In order to apply Theorem 5.1 with $A = \mathbb{N} - \{0\}$, we must find the general solution of $u = Qu$. This equation reads

$$\begin{aligned} u(1) &= pu(2), \\ u(2) &= qu(1) + pu(3), \\ u(3) &= qu(2) + pu(4), \\ &\dots \end{aligned}$$

and its general solution is

$$u(i) = u(1) \sum_{j=0}^{i-1} \left(\frac{q}{p}\right)^j.$$

The largest value of $u(1)$ respecting the constraint $u(i) \in [0, 1]$ is $u(1) = 1 - \left(\frac{q}{p}\right)$. The solution $v(i)$ is obtained by assigning this value, and therefore

$$v(i) = 1 - \left(\frac{q}{p}\right)^i. \tag{5.9}$$

\diamond

Remark 5.1. *Finite Transient Set*

If the set of transient states T is *finite*, the probability of infinite sojourn in T is null. This intuitively obvious fact can also be deduced from application of Formula (5.4) with $A = T$. Indeed, with the notation of (5.2),

$$\mathbf{P}^n = \begin{pmatrix} D^n & 0 \\ B_n & \mathbf{Q}^n \end{pmatrix}, \quad (5.10)$$

and therefore the general term of \mathbf{Q}^n is $p_{ij}(n)$. But we know that for any transient state j , $\lim_{n \uparrow \infty} p_{ij}(n) = 0$ for all $i \in E$, and therefore if T is finite, $v(i) = \lim_{n \uparrow \infty} \sum_{j \in T} p_{ij}(n) = 0$. \diamond

Example 5.2. *Repair Shop*

This HMC was introduced in Example 2.2 of Chapter 2. Its state space is $E = \mathbb{N}$, and its transition matrix is

$$\mathbf{P} = \begin{pmatrix} a_0 & a_1 & a_2 & a_3 & \cdots \\ a_0 & a_1 & a_2 & a_3 & \cdots \\ & a_0 & a_1 & a_2 & \cdots \\ & & a_0 & a_1 & \cdots \\ & & & & \cdots \end{pmatrix},$$

where $\{a_i ; i \geq 0\}$ is a probability distribution on \mathbb{N} of mean

$$\rho = \sum_{k=0}^{\infty} k a_k.$$

It is assumed that $a_0 > 0$ and $a_0 + a_1 < 1$, so that the chain is irreducible (see Problem 2.4.2 of Chapter 2).

One can use Theorem 5.1 to prove that if $\rho \leq 1$, the chain is recurrent. For this consider the matrix

$$\mathbf{Q} = \begin{pmatrix} a_1 & a_2 & a_3 & \cdots \\ a_0 & a_1 & a_2 & \cdots \\ a_0 & a_1 & \cdots & \cdots \end{pmatrix},$$

which is the restriction of \mathbf{P} to $A_i = \{i + 1, i + 2, \dots\}$ for any $i \geq 0$. Thus, the maximal solution of $v = \mathbf{Q}v$, $\mathbf{0}_A \leq v \leq \mathbf{1}_A$ (where $A = A_i$), has in view of Theorem 5.1 the following two interpretations. Firstly, for $i \geq 1$,

$$v(i) = P_i(X_n \geq 1, n \geq 0), \quad (5.11)$$

that is, $v(i)$ is the probability of never visiting 0 when starting from $i \geq 1$. Secondly, $v(1)$ is the probability of never visiting $\{0, 1, \dots, i\}$ when starting from $i + 1$. Since when starting from $i + 1$, the chain visits $\{0, 1, \dots, i\}$ if and only if it visits i , $v(1)$ is the probability of never visiting i when starting from $i + 1$.

We can write the probability of visiting 0 when starting from $i + 1$ as

$$1 - v(i + 1) = (1 - v(1))(1 - v(i)),$$

since in order to go from $i + 1$ to 0 one must first reach i , and then go to 0. Therefore, with

$$v(1) = 1 - \beta,$$

for all $i \geq 1$,

$$v(i) = 1 - \beta^i. \quad (5.12)$$

This solution depends on a parameter $\beta \in [0, 1]$. To determine β , let us write the first equality of $v = \mathbf{Q}v$:

$$v(1) = a_1 v(1) + a_2 v(2) + \dots,$$

that is,

$$(1 - \beta) = a_1(1 - \beta) + a_2(1 - \beta^2) + \dots.$$

Taking $\sum_{i \geq 0} a_i = 1$ into account, this reduces to

$$\beta = g(\beta) \quad (5.13)$$

where $g(z)$ is the generating function of the distribution $(a_k, k \geq 0)$. Also, all other equations of $v = \mathbf{Q}v$ reduce to (5.13).

Recall Theorem 5.1 of Chapter 1, which tells that under the irreducibility assumptions $a_0 > 0, a_0 + a_1 < 1$, (5.13) has only one solution in $[0, 1]$, namely $\beta = 1$ if $\rho \leq 1$, whereas if $\rho > 1$, it has two solutions in $[0, 1]$, this probability is $\beta = 1$ and $\beta = \beta_0 \in (0, 1)$.

Theorem 5.1 says that one must take the smallest β satisfying (5.13). Therefore, if $\rho > 1$, the probability of visiting state 0 when starting from state $i \geq 1$ is $1 - v(i) = \beta_0^i < 1$, and therefore the chain is transient. If $\rho \leq 1$, the latter probability is $1 - v(i) = 1$, and therefore the chain is recurrent.

We know from Example 5.5, Chapter 2, that for $\rho = 1$, there is no stationary distribution. Therefore,

If $\rho = 1$, the chain is null recurrent.

If $\rho > 1$, the chain is transient.

It remains to elucidate the case $\rho < 1$, and we shall see in Chapter 5 that we then have positive recurrence. \diamond

5.2 Time to Absorption

We now turn to the determination of the distribution of τ , the time of exit from the transient set T . Theorem 5.1 tells that $v = \{v(i)\}_{i \in T}$, where $v(i) = P_i(\tau = \infty)$, is the largest solution of $v = \mathbf{Q}v$ subject to the constraints $\mathbf{0}_T \leq v \leq \mathbf{1}_T$, where \mathbf{Q} is the restriction of \mathbf{P} to the transient set T .

The probability distribution of τ when the initial state is $i \in T$ is readily computed starting from the identity

$$P_i(\tau = n) = P_i(\tau \geq n) - P_i(\tau \geq n + 1)$$

and the observation $\{\tau \geq n\} = \{X_{n-1} \in T\}$ for $n \geq 1$, from which we obtain, for $n \geq 1$,

$$P_i(\tau = n) = P_i(X_{n-1} \in T) - P(X_n \in T) = \sum_{j \in T} (p_{ij}(n-1) - p_{ij}(n)).$$

Now, $p_{ij}(n)$ is for $i, j \in T$ the general term of \mathbf{Q}^n , and therefore we have the following result.

Theorem 5.2. *Distribution of the Time to Absorption*

$$P_i(\tau = n) = \{(\mathbf{Q}^{n-1} - \mathbf{Q}^n)\mathbf{1}_T\}_i. \quad (5.14)$$

In particular, if $P_i(\tau = \infty) = 0$,

$$P_i(\tau > n) = \{\mathbf{Q}^n \mathbf{1}_T\}_i. \quad (5.15)$$

In particular, when T is finite, for any distribution ν such that $\nu(T) = 1$, we have

$$P_\nu(\tau > n) = \nu^T \mathbf{Q}^n \mathbf{1}_T.$$

(The reader is begged to excuse the author for the conflict of notation in the above formula, where T denotes transposition in ν^T and a set in $\mathbf{1}_T$!)

6 Absorption

6.1 Fundamental Matrix

The results of the previous section concern the probability of remaining forever in the transient set, or, taking the dual point of view, the probability of never being absorbed by the recurrent set. It remains to compute the probability of being absorbed by a given recurrent class R_j when starting from an initial state $i \in T$. For this, we shall need the notion of *fundamental matrix*, which is linked to that of potential matrix, and therefore to the number of visits to a particular state j when starting from a state i .

As we shall see, only the case $i \in T, j \in T$ requires analysis, and in this case we are looking at the number of visits to $j \in T$ before absorption by $R = R_1 \cup R_2 \cup \dots$ when starting from $i \in T$.

When starting from any state i , the mean number of visits to an arbitrary state j is

$$E_i \left[\sum_{n=0}^{\infty} \mathbf{1}_{\{X_n=j\}} \right] = \sum_{n=0}^{\infty} p_{ij}(n), \quad (6.1)$$

and this is the (i, j) element of the potential matrix

$$\mathbf{G} = \sum_{n=0}^{\infty} \mathbf{P}^n. \quad (6.2)$$

Only the pairs (i, j) where i and j are transient require treatment, because in all other situations the quantity (6.1) is either null or infinite. For instance, if i and j are in the same recurrent class, the number of visits to j starting from i is infinite, whereas it is zero if i and j belong to different recurrent classes, or if i is a recurrent state and j is a transient state. If i is transient and j is recurrent, the corresponding term is infinite if j is accessible from i , and it is zero if j is not accessible from j .

To evaluate (6.2) for $i \in T, j \in T$, one uses the representation (5.2) of the transition matrix, to obtain

$$\mathbf{G} = \begin{pmatrix} E & 0 \\ F & S \end{pmatrix}, \quad (6.3)$$

where

$$S = \sum_{n=0}^{\infty} \mathbf{Q}^n. \quad (6.4)$$

Definition 6.1. *Absorbing Chain and Fundamental Matrix*

An HMC with at least one transient state and one recurrent state is called *absorbing*, and the matrix S is its *fundamental matrix*.

Remark 6.1.

Note that there are *two* different types of fundamental matrices, one for ergodic chains (see Chapter 8) and one for absorbing chains. \diamond

Denoting by I the identity matrix indexed by T , and observing that

$$(I - \mathbf{Q}) \left(\sum_{n=0}^N \mathbf{Q}^n \right) = I - \mathbf{Q}^{N+1},$$

we see, by letting N go to ∞ , that S is a solution of

$$S(I - \mathbf{Q}) = I. \quad (6.5)$$

If T is finite, S is therefore the inverse of $I - \mathbf{Q}$. When T is infinite, there may very well be several solutions of (6.5). However, note the following theorem:

Theorem 6.1. *Maximal Characterization of the Fundamental Matrix*

The matrix S given by (6.4) is the smallest solution of $S(I - \mathbf{Q}) = I$ subject to the constraint $S \geq 0$.

Proof. Let S be a solution. Writing $S(I - \mathbf{Q}) = I$ as $S = I + S\mathbf{Q}$ and iterating the latter equality gives

$$S = I + \mathbf{Q} + \cdots + \mathbf{Q}^n + S\mathbf{Q}^{n+1}. \quad (6.6)$$

But $S \geq 0$ implies $S\mathbf{Q}^{n+1} \geq 0$, and therefore

$$S \geq I + \mathbf{Q} + \cdots + \mathbf{Q}^n.$$

Passing to the limit, we obtain $S \geq \sum_{n=0}^{\infty} \mathbf{Q}^n$. \square

Example 6.1.

The transition matrix

$$\mathbf{P} = \begin{pmatrix} 0.4 & 0.3 & 0.3 & & & & & \\ 0 & 0.6 & 0.4 & & & & & \\ 0.5 & 0.5 & 0 & & & & & \\ & & & 0 & 1 & & & \\ & & & 0.8 & 0.2 & & & \\ 0 & 0 & 0 & & & 0.4 & 0.6 & 0 \\ 0.4 & 0.4 & 0 & & & 0 & 0 & 0.2 \\ 0.1 & 0 & 0.3 & & & 0.6 & 0 & 0 \end{pmatrix}$$

with state space $E = \{1, 2, 3, 4, 5, 6, 7, 8\}$ has two recurrent classes $R_1 = \{1, 2, 3\}$ and $R_2 = \{4, 5\}$ and one transient class $T = \{6, 7, 8\}$. Here

$$\mathbf{Q} = \begin{pmatrix} 0.4 & 0.6 & 0 \\ 0 & 0 & 0.2 \\ 0.6 & 0 & 0 \end{pmatrix}$$

and

$$S = (I - \mathbf{Q})^{-1} = \frac{1}{66} \begin{pmatrix} 125 & 75 & 15 \\ 15 & 75 & 15 \\ 75 & 45 & 75 \end{pmatrix}.$$

Also, no state of T leads to R_2 , and therefore $g_{ij} = 0$ for $i \in T, j \in R_2$. All states of T lead to R_1 , and therefore $g_{ij} = +\infty$ for $i \in T, j \in R_1$. In summary:

$$\mathbf{G} = \begin{pmatrix} +\infty & +\infty & +\infty & 0 & 0 & 0 & 0 & 0 \\ +\infty & +\infty & +\infty & 0 & 0 & 0 & 0 & 0 \\ +\infty & +\infty & +\infty & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & +\infty & +\infty & 0 & 0 & 0 \\ 0 & 0 & 0 & +\infty & +\infty & 0 & 0 & 0 \\ +\infty & +\infty & +\infty & 0 & 0 & & & \\ +\infty & +\infty & +\infty & 0 & 0 & & & \\ +\infty & +\infty & +\infty & 0 & 0 & & & \end{pmatrix},$$

where the missing part is S . ◇

6.2 Absorption Matrix

We seek to compute the probability of absorption by a given recurrent class when starting from a given transient state. As we shall see later, it suffices for the theory to treat the case where the recurrent classes are singletons. We therefore suppose that the transition matrix has the form

$$\mathbf{P} = \begin{pmatrix} I & 0 \\ B & Q \end{pmatrix}. \quad (6.7)$$

Let f_{ij} be the probability of absorption by recurrent class $R_j = \{j\}$ when starting from the transient state i . From (6.7),

$$\mathbf{P}^n = \begin{pmatrix} I & 0 \\ L_n & \mathbf{Q}^n \end{pmatrix},$$

where $L_n = (I + \mathbf{Q} + \cdots + \mathbf{Q}^n)B$. Therefore, $\lim_{n \uparrow \infty} L_n = SB$. For $i \in T$, the (i, j) term of L_n is

$$L_n(i, j) = P(X_n = j | X_0 = i).$$

Now, if T_j is the first time of visit to $R_j = \{j\}$ after time 0, then

$$L_n(i, j) = P_i(T_j \leq n),$$

since $R_j = \{j\}$ is a closed state. Letting n go to ∞ gives the following:

Theorem 6.2. Probability of Absorption

Consider a HMC with transition matrix \mathbf{P} as in (6.7). The probability of absorption by recurrent class $R_j = \{j\}$ starting from transient state i is

$$P_i(T_{R_j} < \infty) = (SB)_{ij}. \quad (6.8)$$

The general case, where the recurrence classes are not necessarily singletons, can be reduced to the singleton case as follows. Let \mathbf{P}^* be the matrix obtained from the transition matrix \mathbf{P} given by (6.1), by grouping for each j the states of recurrent class R_j into a single state \hat{j} :

$$\mathbf{P}^* = \begin{matrix} & \hat{1} & \hat{2} & \cdots & T \\ \hat{1} & \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ b_{\hat{1}} & b_{\hat{2}} & \cdots & \mathbf{Q} \end{pmatrix} & & & \\ \hat{2} & & & & \\ \vdots & & & & \\ T & & & & \end{matrix}, \quad (6.9)$$

where $b_j = B(j)\mathbf{1}_T$ is obtained by summation of the columns of $B(j)$, the matrix consisting of the columns $i \in R_j$ of B .

The probability f_{iR_j} of absorption by class R_j when starting from $i \in T$ equals $\hat{f}_{i\hat{j}}$, the probability of ever visiting \hat{j} when starting from i , computed for the chain with transition matrix \mathbf{P}^* .

Example 6.2.

Consider the chain with state space $E = \{1, 2, 3, 4, 5, 6, 7\}$ and transition matrix

$$\mathbf{P} = \begin{pmatrix} 0.5 & 0.5 & & & & & \\ 0.8 & 0.2 & & & & & \\ & & 0 & 0.4 & 0.6 & & \\ & & 1 & 0 & 0 & & \\ & & 1 & 0 & 0 & & \\ 0.1 & 0 & 0.2 & 0.1 & 0.2 & 0.3 & 0.1 \\ 0.1 & 0.1 & 0.1 & 0 & 0.1 & 0.2 & 0.4 \end{pmatrix}.$$

It has two recurrent classes $R_1 = \{1, 2\}$, $R_2 = \{3, 4, 5\}$ and one transient class $T = \{6, 7\}$. In the notations introduced in the current subsection,

$$B_1 = \begin{pmatrix} 0.1 & 0 \\ 0.1 & 0.1 \end{pmatrix}, B_2 = \begin{pmatrix} 0.2 & 0.1 & 0.2 \\ 0.1 & 0 & 0.1 \end{pmatrix}, \mathbf{Q} = \begin{pmatrix} 0.3 & 0.1 \\ 0.2 & 0.4 \end{pmatrix},$$

and therefore

$$b_1 = \begin{pmatrix} 0.1 \\ 0.2 \end{pmatrix}, b_2 = \begin{pmatrix} 0.5 \\ 0.2 \end{pmatrix}, \hat{B} = \begin{pmatrix} 0.1 & 0.5 \\ 0.2 & 0.2 \end{pmatrix}.$$

Also,

$$\mathbf{P}^* = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0.1 & 0.5 & 0.3 & 0.1 \\ 0.2 & 0.2 & 0.2 & 0.4 \end{pmatrix}.$$

Computations yield

$$S = (I - \mathbf{Q})^{-1} = \begin{pmatrix} 1.5 & 0.25 \\ 0.5 & 1.75 \end{pmatrix}, S\hat{B} = \begin{pmatrix} 0.2 & 0.8 \\ 0.4 & 0.6 \end{pmatrix}.$$

In particular the probability of absorption from transient state 6 to class $\{3, 4, 5\}$ is 0.8. \diamond

Example 6.3. Sibmating

In diploid organisms a given hereditary character is carried by a pair of *genes*. We consider the situation in which each gene can take two forms, called *alleles*, denoted by a and A . Such was the case in the historical experiments performed around the year 1865 by the monk Gregor Mendel, who studied the hereditary transmission of the nature of the skin in a species of green peas. The two alleles corresponding to the genes or character “nature of the skin” are a for “wrinkled” and A for “smooth.” The genes are grouped by pairs, and there are two alleles, and thus three *genotypes* are possible for the character under study: aa , Aa (same as aA), and AA . With each genotype is associated a *phenotype*, which is the external appearance corresponding to the genotype. Genotypes aa , AA have different phenotypes (otherwise, no character could be isolated), and the phenotype of Aa lies somewhere between the phenotypes of aa and AA . Sometimes an allele is *dominant*, say A , and the phenotype of Aa is then the same of the phenotype of AA .

During the reproduction process, each parent contributes to the genetic heritage of the descendant by providing *one* allele of their pair. This is done by intermediaries of the reproduction process, the cells called *gametes* (in the human species, the *spermatozoid* and the *ovula*), which carry only one gene of the pair of genes characteristic of each parent. The gene carried by the gamete is chosen at random from among the pair of genes of the parent. The actual process occurring in the reproduction of diploid cells is called *meiosis* (see Figure 4.6.1).

A given cell possesses two chromosomes. A chromosome can be viewed as a string of genes, each gene being at a specific location in the string.




Figure 4.6.1. Meiosis

Let us start from an idealistically infinite population where the genotypes are found in the following proportions:

$$\begin{matrix} AA & Aa & aa \\ x & 2z & y \end{matrix}$$

Here x , y , and z are numbers between 0 and 1, and

$$x + 2z + y = 1.$$

The two parents are chosen independently (*random mating*), and each of their respective gametes chooses an allele at random independently of each other in the pair carried by the corresponding parent. We leave as an exercise for the reader to prove that the genotypic distribution of all the generations starting with the third one are the same, depending only on the proportions of alleles of type A in the initial population: This is the famous law of Hardy and Weinberg, which holds for random mating.

In *sibmating* (sister–brother mating) two individuals are mated and two individuals from their offspring are chosen at random to be mated, and this incestuous process goes on through the generations.

We shall denote by X_n the genetic types of the mating pair at the n th generation. Clearly, $\{X_n\}_{n \geq 0}$ is an HMC with six states representing the different pairs of genotypes $AA \times AA$, $aa \times aa$, $AA \times Aa$, $Aa \times Aa$, $Aa \times aa$, $AA \times aa$, denoted respectively 1, 2, 3, 4, 5, 6. The following table gives the probabilities of occurrence of the 3 possible genotypes in the descent of a mating pair:

	AA	Aa	aa	
$AA \times AA$	1	0	0	}
$aa \times aa$	0	0	1	
$AA \times Aa$	1/2	1/2	0	
$Aa \times Aa$	1/4	1/2	1/4	
$Aa \times aa$	0	1/2	1/2	
$AA \times aa$	0	1	0	
	descendant's genotype			

The transition matrix of $\{X_n\}_{n \geq 0}$ is then easily deduced:

$$\mathbf{P} = \begin{pmatrix} 1 & & & & & & \\ & 1 & & & & & \\ 1/4 & & 1/2 & 1/4 & & & \\ 1/16 & 1/16 & 1/4 & 1/4 & 1/4 & 1/8 & \\ & 1/4 & & 1/4 & 1/2 & & \\ & & & & & 1 & \end{pmatrix}.$$

The set $R = \{1, 2\}$ is absorbing, and the restriction of the transition matrix to the transient set $T = \{3, 4, 5, 6\}$ is

$$\mathbf{Q} = \begin{pmatrix} 1/2 & 1/4 & 0 & 0 \\ 1/4 & 1/4 & 1/4 & 1/8 \\ 0 & 1/4 & 1/2 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}.$$

The fundamental matrix is

$$\mathbf{S} = (\mathbf{I} - \mathbf{Q})^{-1} = \frac{1}{6} \begin{pmatrix} 16 & 8 & 4 & 1 \\ 8 & 16 & 8 & 2 \\ 4 & 8 & 16 & 1 \\ 8 & 16 & 8 & 8 \end{pmatrix},$$

and the absorption probability matrix is

$$\mathbf{SB} = \mathbf{S} \begin{pmatrix} 1/4 & 0 \\ 1/16 & 1/16 \\ 0 & 1/4 \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 3/4 & 1/4 \\ 1/2 & 1/2 \\ 1/4 & 3/4 \\ 1/2 & 1/2 \end{pmatrix}.$$

For instance, the $(3, 2)$ entry, $\frac{3}{4}$, is the probability that when starting from a couple of ancestors of type $Aa \times aa$, the race will end up in the genotype $aa \times aa$. \diamond

We have seen that the absorption probabilities by the recurrent classes corresponding can be obtained as a byproduct of the *access matrix*

$$\mathbf{F} = \{f_{ij}\},$$

where

$$f_{ij} = P_i(T_j < \infty),$$

because $f_{iR_j} = f_{ik}$ for any $k \in R_j$. For the absorption probabilities, we need only the entries of the access matrix corresponding to the indices $i \in T, j \in R$. Note, however, that the other entries of the access matrix can be obtained from the potential matrix. Indeed,

(α) If i and j are in the same recurrent class, $f_{ij} = 1$.

(β) If i and j are in two different recurrent classes, $f_{ij} = 0$.

(γ) If i is recurrent and j is transient, $f_{ij} = 0$.

(δ) If i and j are transient,

$$f_{ij} = \frac{g_{ij}}{g_{jj}} \text{ if } j \neq i \text{ and } f_{jj} = \frac{g_{jj} - 1}{g_{jj}}. \quad (6.10)$$

To prove (6.10), use Theorem 7.2 of Chapter 2.

Problems

4.1.1

(a) Prove (1.5).

(b) Show that

$$d_V(\alpha, \beta) = \sup \left(\sum_{i=1}^r \alpha(i)y(i) - \sum_{i=1}^r \beta(i)y(i); \sup |y(i)| = 1 \right).$$

(c) Show that if $\lambda_i \in [0, 1]$ and $\sum_{i=1}^K \lambda_i = 1$, then

$$d_V \left(\sum_{i=1}^K \lambda_i \alpha_i, \sum_{i=1}^K \lambda_i \beta_i \right) \leq \sum_{i=1}^K \lambda_i d_V(\alpha_i, \beta_i),$$

where α_i and β_i are probability distributions on E for all $i \in [0, K]$.

4.1.2 Let f_1 and f_2 be two probability density functions on \mathbb{R}^k , and call $\mathcal{D}(f_1, f_2)$ the collection of couples of \mathbb{R}^k -valued random variables (X, Y) such that X and Y have the p.d.f. f_1 and f_2 respectively.

Show that for any $(X, Y) \in \mathcal{D}(f_1, f_2)$,

$$P(X = Y) \leq 1 - \frac{1}{2} \int_{\mathbb{R}^k} |f_1(x) - f_2(x)| dx$$

and that equality can be attained.

Exhibit (X, Y) achieving equality in the following cases.

(i) $k = 1$, X uniform on $[0, \theta_1]$, Y uniform on $[0, \theta_2]$

(ii) $k = 2$, X uniform on $[0, \theta_1]^2$, Y uniform on $[0, \theta_2]^2$

(iii) $k = 1$, X exponential with parameter θ_1 , Y exponential with parameter θ_2

4.1.3 Consider the following situation. A random variable X has a distribution depending on the value of another random variable $\Theta \in \{1, 2\}$. For $i = 1, 2$, conditionally to $\Theta = i$ the random variable X has the p.d.f. $f_i(x)$, and $P(\Theta = i) = \frac{1}{2}$. The game is the following. Find a partition $A_1 + A_2 = \mathbb{R}$ such that if we define $\Theta' = i$ if $X \in A_i$, then the probability of error, that is, the probability that the estimate Θ' is different from the actual value Θ , is minimized. What is then the probability of error?

4.1.4 Let \mathbf{P} be a transition matrix on the finite state space E . Define

$$\delta(\mathbf{P}) = \frac{1}{2} \sup_{i, j \in E} \sum_{k \in E} |p_{ik} - p_{jk}|.$$

Let $\{X_n\}_{n \geq 0}$ and $\{Y_n\}_{n \geq 0}$ be two E -valued stochastic processes with the following properties:

The initial states are maximally coupled with marginal distributions μ and ν , respectively. Also, $\{(X_n, Y_n)\}_{n \geq 0}$ is an HMC with transition probabilities $q_{i,j;k,l}$ where for fixed (i, j) , the probability distribution $q_{i,j;\cdot,\cdot}$ realizes maximal coincidence of $p_{i,\cdot}$ and $p_{j,\cdot}$. (this type of coupling is called *Markov coupling*.)

Show that $\{X_n\}_{n \geq 0}$ and $\{Y_n\}_{n \geq 0}$ are two HMCs with the transition matrix \mathbf{P} .

Show that if $\delta(\mathbf{P}) < 1$, then $\{X_n\}_{n \geq 0}$ and $\{Y_n\}_{n \geq 0}$ couple in finite time τ and that $E[e^{\alpha\tau}] < \infty$ for some positive α .

Show the same when only $\delta(\mathbf{P})^m < \infty$ is assumed, for some integer $m > 1$.

Deduce from this that if \mathbf{P} is ergodic, $\{X_n\}_{n \geq 0}$ and $\{Y_n\}_{n \geq 0}$ couple in finite time τ and that $E[e^{\alpha\tau}] < \infty$ for some positive α .

4.2.1 Prove in detail the last assertion of Theorem 2.2.

4.2.2 Let $\{X_n^1\}_{n \geq 0}$ and $\{X_n^2\}_{n \geq 0}$ be two independent irreducible HMCs with the same transition matrix \mathbf{P} . Give a counterexample showing that if \mathbf{P} is not aperiodic, the product chain $\{(X_n^1, X_n^2)\}_{n \geq 0}$ may be reducible.

4.2.3 Let $\{X_n\}_{n \geq 0}$ and $\{Y_n\}_{n \geq 0}$ be two independent HMCs on the state space $E = \{1, 2\}$, with the same transition matrix

$$\mathbf{P} = \begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix},$$

where $\alpha, \beta \in (0, 1)$. Let τ be the first time n when $X_n = Y_n$. Compute the probability distribution of τ when $X_0 = 1, Y_0 = 2$.

4.2.4 Let \mathbf{P} be an ergodic transition matrix on the countable state space E . Show that a necessary and sufficient condition for the corresponding HMC to be reversible is that for all states i, i_1, \dots, i_{k-1}

$$p_{ii_1} p_{i_1 i_2} \cdots p_{i_{k-1} i} = p_{ii_{k-1}} \cdots p_{i_2 i_1} p_{i_1 i}.$$

4.2.5 A man commutes every weekend between his town house and his country house. Each Saturday, he leaves town for the countryside, and he comes back on Sunday. He owns N

umbrellas which he takes between his two houses. He takes one, if available, if it is raining when he leaves town (resp., the countryside). The probability of rain on a given day is p .

Find the steady state probability of having i umbrellas in the country house on a given Saturday night.

4.3.1 At a crosswalk, cars pass on a single lane at times $R_0 = 0, R_1, R_2, \dots$, where $\{R_n\}_{n \geq 0}$ is a proper renewal sequence. A pedestrian arriving at time 0 crosses the lane as soon as he sees a time interval $x > 0$ between two consecutive cars. How long must he wait, on the average?

4.3.2 Let L be the lifetime of a defective renewal sequence. Show that $\lim_{n \uparrow \infty} P(L > n) = 0$, and give the rate of convergence. Treat in detail the case where the interrenewal sequence is geometric.

4.3.3 Let $v((a, b])$ be the average number of renewal epochs in the integer interval $(a, b]$ of a proper nonlattice renewal sequence. What is the limit as $n \uparrow \infty$ of $v((n + a, n + b])$?

4.3.4 Suppose that the typical interrenewal time S_1 of a renewal sequence is proper but that $P(S_1 = 0) = f_0 > 0$. Otherwise, suppose that $\gcd\{n \geq 1; f_n > 0\} = 1$. Show that the solution of the *extended* renewal equation

$$u_0 = v_0, \quad u_n = v_n + \sum_{k=0}^n f_k u_{n-k}, \quad n \geq 1$$

(notice the additional term in the sum, corresponding to $k = 0$) satisfies, under the summability condition (3.4),

$$\lim_{n \uparrow \infty} u_n = \frac{\sum_{k \geq 0} v_k}{\sum_{k=1}^{\infty} k f_k}.$$

4.3.5 In the population model of Example 3.4 what is, in the critical case $\rho = 1$, the average number of daughters born at a given large time, when $\alpha(i) = 0$ for all $i > 0$, and $\alpha(0) = 1$? Suppose now that $f_k = e^{-\beta \frac{\theta^{k-1}}{(k-1)!}}$ for $k \geq 1$. Discuss the asymptotic behavior of u_n , the average number of daughters born at time n , in terms of the positive parameters β and θ (use the same initial conditions as in the first part of the problem).

4.4.1 A given machine can be in either one of three states: G (*good*), M (*in maintenance*), or R (*in repair*). Its successive periods where it is in state G (resp., M, R) form an independent and identically distributed sequence $\{S_n\}_{n \geq 0}$ (resp., $\{U_n\}_{n \geq 0}$, $\{V_n\}_{n \geq 0}$) with finite mean. All these sequences are assumed mutually independent. The maintenance policy uses a number $T > 0$. If the machine has age T and has not failed, it goes to state M. If it fails before it has reached age T , it enters state R. From states M and R, the next state is G. Find the steady state probability that the machine is operational. (Note that “good” does not mean “operational.” The machine can be “good” but, due to the operations policy, in maintenance, and therefore not operational. However, after a period of maintenance or of repair, we consider that the machine starts anew, and enters a G period.)

4.4.2 In Example 4.4, when the typical interrenewal time is geometric, compute $\lim_{n \rightarrow \infty} P(F_n + B_n = k)$, $\lim_{n \uparrow \infty} P(F_n = k)$, and $\lim_{n \uparrow \infty} P(F_n + B_n = k)$.

4.4.3 In Example 4.4, under what circumstances do we have $\lim_{n \uparrow \infty} P(F_n + B_n = k) = P(S_1 = k)$ for all $k \geq 1$?

4.4.4 In Example 4.3, we abandon the nonlattice hypothesis. Show that

$$\lim_{t \uparrow \infty} \int_0^t \mathbf{1}_{\{Z(s)=0\}} ds = \frac{E[U_1]}{E[U_1] + E[V_1]},$$

where $Z(s) = 0$ means that the machine is in repair at time s .

4.5.1 Let $\{X_n\}_{n \geq 0}$ be an irreducible HMC with finite state space $E = \{1, 2, \dots, r\}$ and transition matrix \mathbf{P} . Let T_j be the return time to $j \in E$. Show that

$$P_i(T_j > n) = \{\mathbf{Q}_j^n \mathbf{1}_{E \setminus \{j\}}\}_i,$$

where \mathbf{Q}_j is obtained from \mathbf{P} by deleting the j th column and the j th row, and $\mathbf{1}_{E \setminus \{j\}}$ is a column vector of dimension $r - 1$ with all its entries equal to 1.

4.5.2 Find the distribution of the first meeting time of two independent HMC with state space $E = \{1, 2\}$ and transition matrix

$$\mathbf{P} = \begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix},$$

where $\alpha, \beta \in (0, 1)$, when their initial states are different.

4.5.3 Let $\{X_n\}_{n \geq 0}$ be an HMC on the state space $E = \{1, 2, 3, 4, 5\}$ with transition matrix

$$\mathbf{P} = \begin{pmatrix} 1 & & & & \\ & 0.4 & 0.6 & & \\ & 0.5 & 0.5 & & \\ 0.4 & 0.2 & 0.2 & 0.1 & 0.1 \\ 0.5 & 0 & 0.1 & 0.2 & 0.2 \end{pmatrix}.$$

Compute the probability of absorption in the recurrent class $\{2, 3\}$ when starting from the transient state 5.

4.5.4 Give a direct proof of the result in Remark 5.1 stating that infinite sojourn in a *finite* transient set is almost surely impossible.

4.5.5 Consider an HMC with a nonempty recurrent set R and a transient set T with at least two distinct elements i and j . Give a general method to compute the probability of never visiting j and of being absorbed by a given recurrent class R_k when starting from i . Apply this to the HMC of Example 6.2, with $i = 6$, $j = 7$, and $k = 2$.

4.6.1 You have 1 dollar and you owe the mafia 5 dollars. Being just a probabilist, you decide to make a series of bets with a fair coin. You have the choice between two strategies. Bet 1

dollar each time (the timid strategy), or bet as much as you can but no more than necessary to reach a fortune of 5 dollars (the bold strategy). Find your probability of staying alive in each strategy.

4.6.2 Consider an HMC with a nonempty set of transient states T and with at least one positive recurrent class R that is aperiodic. Let π be the unique stationary distribution of the restriction of the chain to R . Prove that for any $i \in T, j \in R, \lim_{n \uparrow \infty} p_{ij}(n) = f_{iR}\pi(j)$, where f_{iR} is the probability of absorption in R when starting from i .

Lyapunov Functions and Martingales

1 Lyapunov Functions

1.1 Foster's Theorem

The present chapter contains a potpourri of topics around potential theory and martingale theory. More exactly, it is a brief introduction to these topics, with the limited purpose of showing the power of martingale theory and the rich interplay between probability and analysis. An important aspect of this chapter concerns the various results complementing the study of recurrence of Chapter 3. In this respect, the single most important result is Foster's theorem below.

The stationary distribution criterion of positive recurrence of an irreducible chain requires solving the balance equation, a too often hopeless enterprise except in a few textbook situations. The following sufficient condition is more tractable, and indeed quite powerful.

Theorem 1.1. *Foster's Theorem*

Let the transition matrix \mathbf{P} on the countable state space E be irreducible and suppose that there exists a function $h : E \rightarrow \mathbb{R}$ such that $\inf_i h(i) > -\infty$ and

$$\sum_{k \in E} p_{ik} h(k) < \infty \text{ for all } i \in F, \tag{1.1}$$

$$\sum_{k \in E} p_{ik} h(k) \leq h(i) - \epsilon \text{ for all } i \notin F, \tag{1.2}$$

for some finite set F and some $\epsilon > 0$. Then the corresponding HMC is positive recurrent.

One could call the function h *uniformly strict* superharmonic outside F , but this is not standard terminology.

The proof below does not use explicitly martingale theory, but Foster's theorem is very much related to results in the Section 3 that are proved via the martingale convergence theorem (Theorem 3.1).

Proof. We use the notation X_0^n for (X_0, \dots, X_n) . Since $\inf_i h(i) > -\infty$, one may assume without loss of generality that $h \geq 0$, by adding a constant if necessary. Call τ the return time to F , and define $Y_n = h(X_n)1_{\{n < \tau\}}$. Equality (1.2) is just $E[h(X_{n+1}) \mid X_n = i] \leq h(i) - \epsilon$ for all $i \notin F$. For $i \notin F$,

$$\begin{aligned} E_i[Y_{n+1} \mid X_0^n] &= E_i[Y_{n+1}1_{\{n < \tau\}} \mid X_0^n] + E_i[Y_{n+1}1_{\{n \geq \tau\}} \mid X_0^n] \\ &= E_i[Y_{n+1}1_{\{n < \tau\}} \mid X_0^n] \leq E_i[h(X_{n+1})1_{\{n < \tau\}} \mid X_0^n] \\ &= 1_{\{n < \tau\}} E_i[h(X_{n+1}) \mid X_0^n] = 1_{\{n < \tau\}} E_i[h(X_{n+1}) \mid X_n] \\ &\leq 1_{\{n < \tau\}} h(X_n) - \epsilon 1_{\{n < \tau\}}, \end{aligned}$$

where the third equality comes from the fact that $1_{\{n < \tau\}}$ is a function of X_0^n and uses rule (7.8) of Chapter 1, the fourth equality is the Markov property, and the last inequality is true because P_i -a.s., $X_n \notin F$ on $n < \tau$. Therefore, P_i -a.s.,

$$E_i[Y_{n+1} \mid X_0^n] \leq Y_n - \epsilon 1_{\{n < \tau\}},$$

and taking expectations,

$$0 \leq E_i[Y_{n+1}] \leq E_i[Y_n] - \epsilon P_i(\tau > n).$$

Iterating the above equality, and observing that Y_n is nonnegative, we obtain

$$0 \leq E_i[Y_0] - \epsilon \sum_{k=0}^n P_i(\tau > k).$$

But $Y_0 = h(i)$, P_i -a.s., and $\sum_{k=0}^{\infty} P_i(\tau > k) = E_i[\tau]$. Therefore, for all $i \notin F$,

$$E_i[\tau] \leq \epsilon^{-1} h(i).$$

For $j \in F$, first-step analysis yields

$$E_j[\tau] = 1 + \sum_{i \notin F} p_{ji} E_i[\tau].$$

Thus $E_j[\tau] \leq 1 + \epsilon^{-1} \sum_{i \notin F} p_{ji} h(i)$, and this quantity is finite in view of assumption (1.1). Therefore, the return time to F starting anywhere in F has finite expectation. Since F is a finite set, this implies positive recurrence in view of the following lemma. \square

Lemma 1.1. Let $\{X_n\}_{n \geq 0}$ be an irreducible HMC, let F be a finite subset of the state space E , and let $\tau(F)$ be the return time to F . If $E_j[\tau(F)] < \infty$ for all $j \in F$, the chain is positive recurrent.

Proof. Select $i \in F$, and let T_i be the return time of $\{X_n\}$ to i . Let $\tau_1 = \tau(F)$, τ_2, τ_3, \dots be the successive return times to F . It follows from the strong Markov property that $\{Y_n\}_{n \geq 0}$ defined by $Y_0 = X_0 = i$ and $Y_n = X_{\tau_n}$ for $n \geq 1$ is an HMC with state space F (Problem 2.7.2). Since $\{X_n\}$ is irreducible, so is $\{Y_n\}$. Since F is finite, $\{Y_n\}$ is positive recurrent, and in particular, $E_i[\tilde{T}_i] < \infty$, where \tilde{T}_i is the return time to i of $\{Y_n\}$. Defining $S_0 = \tau_1$ and $S_k = \tau_{k+1} - \tau_k$ for $k \geq 1$, we have

$$T_i = \sum_{k=0}^{\infty} S_k 1_{\{k < \tilde{T}_i\}},$$

and therefore

$$E_i[T_i] = \sum_{k=0}^{\infty} E_i[S_k 1_{\{k < \tilde{T}_i\}}].$$

Now,

$$E_i[S_k 1_{\{k < \tilde{T}_i\}}] = \sum_{\ell \in F} E_i[S_k 1_{\{k < \tilde{T}_i\}} 1_{\{X_{\tau_k} = \ell\}}],$$

and by the strong Markov property applied to $\{X_n\}$ and the stopping time τ_k , and the fact that the event $\{k < \tilde{T}_i\}$ belongs to the past of $\{X_n\}$ at time τ_k ,

$$\begin{aligned} E_i[S_k 1_{\{k < \tilde{T}_i\}} 1_{\{X_{\tau_k} = \ell\}}] &= E_i[S_k \mid k < \tilde{T}_i, X_{\tau_k} = \ell] P_i(k < \tilde{T}_i, X_{\tau_k} = \ell) \\ &= E_i[S_k \mid X_{\tau_k} = \ell] P_i(k < \tilde{T}_i, X_{\tau_k} = \ell). \end{aligned}$$

Observing that $E_i[S_k \mid X_{\tau_k} = \ell] = E_\ell[\tau(F)]$, we see that the latter expression is bounded by $(\max_{\ell \in F} E_\ell[\tau(F)]) P_i(k < \tilde{T}_i, X_{\tau_k} = \ell)$, and therefore

$$E_i[T_i] \leq \left(\max_{\ell \in F} E_\ell(\tau(F)) \right) \sum_{k=0}^{\infty} P_i(\tilde{T}_i > k) = \left(\max_{\ell \in F} E_\ell(\tau(F)) \right) E_i[\tilde{T}_i] < \infty. \quad \square$$

The function h in Foster's theorem is called a *Lyapunov function* because it plays a role similar to the Lyapunov functions in the stability theory of ordinary differential equations. It has a tendency to decrease along the trajectories of the process, at least outside a finite set of states, the *refuge*. Since it is nonnegative, it cannot decrease forever, and therefore it eventually enters the refuge.

Corollary 1.1. *Pakes's Lemma*

Let $\{X_n\}_{n \geq 0}$ be an irreducible HMC on $E = \mathbb{N}$ such that for all $n \geq 0$ and all $i \in E$,

$$E[X_{n+1} \mid X_n = i] < \infty \tag{1.3}$$

and

$$\limsup_{i \uparrow \infty} E[X_{n+1} - X_n \mid X_n = i] < 0. \tag{1.4}$$

Such an HMC is positive recurrent.

Proof. Let -2ϵ be the left-hand side of (1.4). In particular, $\epsilon > 0$. By (1.4), for i sufficiently large, say $i > i_0$, $E[X_{n+1} - X_n \mid X_n = i] < -\epsilon$. We are therefore in the conditions of Foster's theorem with $h(i) = i$ and $F = \{i; i \leq i_0\}$. \square

Example 1.1. *A Random Walk on \mathbb{N}*

Let $\{Z_n\}_{n \geq 1}$ be an i.i.d sequence of integrable random variables with values in \mathbb{Z} such that

$$E[Z_1] < 0,$$

and define $\{X_n\}_{n \geq 0}$, an HMC with state space $E = \mathbb{N}$, by

$$X_{n+1} = (X_n + Z_{n+1})^+,$$

where X_0 is independent of $\{Z_n\}_{n \geq 1}$. Assume irreducibility (the reader is invited to find the necessary and sufficient condition for this). Here

$$\begin{aligned} E[X_{n+1} - i \mid X_n = i] &= E[(i + Z_{n+1})^+ - i] = E[-i 1_{\{Z_{n+1} \leq -i\}} + Z_{n+1} 1_{\{Z_{n+1} > -i\}}] \\ &\leq E[Z_1 1_{\{Z_1 > -i\}}]. \end{aligned}$$

By dominated convergence (see Theorem 3.2 of the appendix), the limit of $E[Z_1 1_{\{Z_1 > -i\}}]$ as i tends to ∞ is $E[Z_1] < 0$ and therefore, by Pakes's lemma, the HMC is positive recurrent. \diamond

Example 1.2. *Repair Shop*

Arguments very similar to those in the previous example show that in the repair shop HMC (assumed irreducible; see Examples 2.2 and 5.5 of Chapter 2), condition $E[Z_1] < 1$ implies positive recurrence. Actually, in view of the results of Example 5.5, Chapter 2, condition $E[Z_1] < 1$ is necessary and sufficient for positive recurrence. \diamond

Next example features a quadratic Lyapunov function. The computations involved are a bit tedious, although straightforward once the correct form of the Lyapunov function is given.

Example 1.3. *A 2-D Random Walk with Reflection*

The reader is referred to the monograph (Fayolle, Malyshev, and Menshikov, 1995) for a general theory of reflected random walks and for bibliographical details.

The state space is $E = \mathbb{N}^2$ and therefore $X_n = (X_n^1, X_n^2)$. The process $\{X_n\}_{n \geq 0}$ evolves according to

$$X_{n+1}^\ell = X_n^\ell + Z_{n+1}^\ell, \quad (1.5)$$

($\ell \in \{1, 2\}$), where Z_{n+1}^ℓ is a *bounded* random variable such that $\{X_n\}_{n \geq 0}$ remains in the first quadrant, and with a conditional distribution given the past of the chain up to time n depending only on the last state X_n . In particular, with $Z_n = (Z_n^1, Z_n^2)$,

$$E[Z_{n+1} \mid X_n = i] = d(i), \quad (1.6)$$

where

$$d(i) = (\alpha(i), \beta(i)). \quad (1.7)$$

The vector $d(i)$ is the *drift* at position $i = (i_1, i_2)$. It will be assumed that

$$(\alpha(i), \beta(i)) = \begin{cases} (\alpha, \beta) & \text{if } i_1 > 0, i_2 > 0 \\ (\alpha_1, \beta_1) & \text{if } i_1 > 0, i_2 = 0 \\ (\alpha_2, \beta_2) & \text{if } i_1 = 0, i_2 > 0 \\ (\alpha_0, \beta_0) & \text{if } i_1 = i_2 = 0 \end{cases} \quad (1.8)$$

Of course, since the chain must remain in the nonnegative quadrant, we have the conditions $\alpha_0 \geq 0, \beta_0 \geq 0$, and $\beta_1 \geq 0, \alpha_2 \geq 0$. Actually we shall assume slightly more:

$$\beta_1 > 0, \alpha_2 > 0, \quad (1.9)$$

just to prevent the chain to be absorbed by a boundary.

It will be shown that for positive recurrence, it suffices to have

$$\alpha < 0, \beta < 0, \alpha\beta_1 - \alpha_1\beta < 0, \beta\alpha_2 - \alpha_2\beta < 0. \quad (1.10)$$

For this, Foster's theorem will be applied with a quadratic Lyapunov function

$$h(i) = h(i_1, i_2) = \frac{1}{2}(ui_1^2 + vi_2^2) + wi_1i_2, \quad (1.11)$$

for suitable, u, v, w .

First u, v, w will be selected in such a way that h is a positive definite quadratic form. This is guaranteed by

$$u > 0, v > 0, w^2 > uv. \quad (1.12)$$

Next a further condition on u, v, w must be found in order to guarantee that

$$E[h(X_{n+1}) - h(X_n) \mid X_n = i] \leq -\epsilon.$$

Elementary computations yield

$$\begin{aligned} h(X_{n+1}) - h(X_n) &= h(X_{n+1} - X_n) + [u(X_{n+1}^1 - X_n^1) + w(X_{n+1}^2 - X_n^2)]X_n^1 \\ &\quad + [v(X_{n+1}^2 - X_n^2) + w(X_{n+1}^1 - X_n^1)]X_n^2. \end{aligned}$$

Therefore for all $i = (i_1, i_2)$

$$\begin{aligned} E[h(X_{n+1}) - h(X_n) \mid X_n = i] &= E[h(X_{n+1} - X_n) \mid X_n = i] + [u\alpha(i) + w\beta(i)]i_1 \\ &\quad + [v\beta(i) + w\alpha(i)]i_2. \end{aligned}$$

By the boundedness hypothesis on (Z_{n+1}^1, Z_{n+1}^2) , the first term in the right hand side of the above equality is bounded. Therefore, if one can choose u, v, w such that (1.12) is verified and

$$\begin{aligned} u\alpha + w\beta &< 0, \quad v\beta + w\alpha < 0, \\ u\alpha_1 + w\beta_1 &< 0, \quad v\beta_2 + w\alpha_2 < 0, \end{aligned}$$

the left-hand side can be made smaller than $-\epsilon$ where ϵ is an arbitrary positive number as long as i_1 and i_2 are sufficiently large, and the proof will be completed.

We therefore have to show the existence of u, v, w such that

$$(C1) \quad u\alpha + w\beta < 0$$

$$(C2) \quad u\alpha_1 + w\beta_1 < 0$$

$$(C3) \quad v\beta + w\alpha < 0$$

$$(C4) \quad v\beta_2 + w\alpha_2 < 0$$

with the conditions

$$(C5) \quad u > 0$$

$$(C6) \quad v > 0$$

$$(C7) \quad w^2 > uv$$

For this, one has the hypotheses

$$(h1) \quad \beta_1 > 0$$

$$(h2) \quad \alpha_2 > 0$$

$$(h3) \quad \alpha < 0$$

$$(h4) \quad \beta < 0$$

$$(h5) \quad \alpha\beta_1 - \alpha_1\beta < 0$$

$$(h6) \quad \beta\alpha_2 - \alpha\beta_2 < 0$$

Since $\beta < 0$ and $\beta_1 > 0$, (C1) and (C2) are respectively equivalent to

$$w > -u\alpha/\beta,$$

and

$$-u\alpha_1/\beta_1 > w,$$

that is

$$(C'1) \quad -u\alpha/\beta < w < -u\alpha_1/\beta_1$$

Similarly, (C3) and (C4) are equivalent to

$$(C'3) \quad -v\beta/\alpha < w < -v\beta_2/\alpha_2$$

So far the existence of u, v, w verifying (C'1), (C'3), (C5), (C6) and (C7) remains to be proven. But since $\alpha < 0$ and $\beta < 0$, (C'1) and (C'3) imply (C7) provided $u > 0, v > 0$, so that we have to find u, v, w verifying (C'1), (C'3) and $u > 0, v > 0$. For this we select

$u > 0$ and $v > 0$ arbitrarily, and then observe that in view of (h5) and $\beta < 0, \beta_1 > 0$,

$$-u\alpha/\beta < -u\alpha_1/\beta_1,$$

and in view of (h6) and $\alpha < 0, \alpha_2 > 0$,

$$-v\beta/\alpha < -v\beta_2/\alpha_2.$$

These two inequalities show that either one of (C'1) and (C'3) is feasible, for any $u > 0, v > 0$. They will be simultaneously satisfied by some w if

$$(-u\alpha/\beta, -u\alpha_1/\beta_1) \cap (-v\beta/\alpha, -v\beta_2/\alpha_2) \neq \emptyset.$$

But for this, it suffices to take $u = \beta/\alpha$ and $v = \alpha/\beta$. ◇

1.2 Queuing Applications

Example 1.4. Stabilization of ALOHA

It was proven in Example 3.3, Chapter 3, that the ALOHA protocol with a fixed retransmission probability ν is unstable. It seems natural to try a retransmission probability $\nu = \nu(k)$ depending on the number k of backlogged messages. Fayolle (1976) has shown that there is a choice of the function $\nu(k)$ that achieves stability of the protocol.

The probability that i among the k backlogged messages at the beginning of slot n retransmit in slot n is as in Example 3.3, Chapter 3, except that ν is replaced by $\nu(k)$. The same is true for the transition probabilities.

An elementary computation yields

$$E[X_{n+1} - X_n \mid X_n = i] = \lambda - b_1(i)a_0 - b_0(i)a_1. \tag{1.13}$$

Note that $b_1(i)a_0 + b_0(i)a_1$ is the probability of one successful (re-)transmission in a slot given that the backlog at the beginning of the slot is i . Equivalently, since there is at most one successful (re-)transmission in any slot, this is the average number of successful (re-)transmissions in a slot given the backlog i at the start of the slot.

According to Pakes's lemma, it suffices to find a function $\nu(k)$ guaranteeing that

$$\lambda \leq \lim_{i \uparrow \infty} (b_1(i)a_0 + b_0(i)a_1) - \epsilon, \tag{1.14}$$

for some $\epsilon > 0$. We shall therefore study the function

$$g_k(\nu) = (1 - \nu)^k a_1 + k\nu(1 - \nu)^{k-1} a_0,$$

since condition (1.14) is just $\lambda \leq g_i(\nu(i)) - \epsilon$. The derivative of $g_k(\nu)$ is, for $k \geq 2$,

$$g'_k(\nu) = k(1 - \nu)^{k-2} [(a_0 - a_1) - \nu(ka_0 - a_1)].$$

We first assume that $a_0 > a_1$. In this case, for $k \geq 2$, the derivative is zero for

$$v = v(k) = \frac{a_0 - a_1}{ka_0 - a_1},$$

and the corresponding value of $g_k(v)$ is a maximum equal to

$$g_k(v(k)) = a_0 \left(\frac{k-1}{k - a_1/a_0} \right)^{k-1}.$$

Therefore, $\lim_{k \uparrow \infty} g_k(v(k)) = a_0 \exp \left\{ \frac{a_1}{a_0} - 1 \right\}$, and we see that

$$\lambda < a_0 \exp \left\{ \frac{a_1}{a_0} - 1 \right\} \quad (1.15)$$

is a sufficient condition for stability of the protocol. For instance, with a Poisson distribution of arrivals

$$a_i = e^{-\lambda} \frac{\lambda^i}{i!},$$

condition (1.15) reads

$$\lambda < e^{-1}$$

(in particular, the condition $a_0 > a_1$ is satisfied a posteriori).

If $a_0 \leq a_1$, the protocol can be shown to be unstable, whatever retransmission policy $v(k)$ is adopted (the reader is invited to check this). \diamond

Example 1.5. A Collision-Resolution Protocol

The slotted ALOHA protocol with constant retransmission probability in Example 3.3, Chapter 3, was proved unstable, and it was shown in Example 1.4 that a backlog-dependent retransmission probability could restore stability. The problem then resides in the necessity for each user to know the size of the backlog in order to implement the retransmission policy. This is not practically feasible, and one must devise policies based on the actual information available by just listening to the link: collision, no transmission, or successful transmission. Such policies, which in a sense estimate the backlog, have been found that yield stability. However, we shall not discuss them here, and instead we shall consider a new type of protocol, the so-called *collision-resolution protocol*, or *binary tree protocol* (Tsybahov and Mihailov, Capetanakis; see (Rom and Sidi, 1990) for bibliographical details).

In these protocols, when a collision occurs, all new requests are buffered until all the messages involved in the collision have found their way through the link. When these messages have resolved their collision problem, the buffered messages then try to retransmit, maybe enter a collision, and then resolve their collision.

Time is therefore divided into successive periods, called *collision-resolution intervals* (CRI). Let us examine the fate of the messages arriving in the first slot, which are the messages that arrived during the previous CRI. They all try to retransmit in the first slot of

the CRI, and therefore, if there are two or more messages, a collision occurs (in the other case, the CRI has lasted just one slot, and a new CRI begins in the next slot). An unbiased coin is tossed for each colliding message. If it shows heads, the message joins *layer 0* of a *stack*, whereas if it shows tails, it is placed in layer 1. In the next slot, all messages of layer 0 try the link. If there is no collision (because layer 0 was empty or just contained one message), layer 0 is eliminated, and layer 1 below pops up to become layer 0. If on the contrary there is a collision because layer 0 formed after the first slot contained two or more messages, the colliding messages again flip a coin; those with heads form the new layer 0, those with tails form the new layer 1, and the former layer 1 is pushed bottomwards to form layer 2.

In general, at each step, only layer 0 tries to retransmit. If there is no collision, layer 0 disappears, and the layers 1, 2, 3, . . . become layers 0, 1, 2, . . . If there is a collision, layer 0 splits into layer 0 and layer 1, and layers 1, 2, 3, . . . become layer 2, 3, 4, . . . It should be noted that in this protocol, each user (message) knows at every instant in which layer he or she is, just by listening to the channel that gives the information: collision or no collision. In that sense, the protocol is *distributed*, because there is no central operator broadcasting nonlocally available information, such as the size of the backlog, to all users.

Once a collision is resolved, i.e., when all layers have disappeared, a new CRI begins. The number of customers that are starting this CRI are those that have arrived in the CRI that just ended. Figure 5.2.1 gives an example of what happens in a CRI.




Figure 5.2.1. The binary tree protocol

If we assume that the fresh requests are as in Example 1.4, that is, $\{A_n\}_{n \geq 1}$ is i.i.d., where A_n is the number of new requests in slot n , then the sequence, $\{X_n\}_{n \geq 0}$, where X_n is the *length* of the n th CRI, forms an irreducible HMC. Stability of the protocol is naturally identified with positive recurrence of this chain, which will now be proved with the help of Foster’s theorem.

According to Pakes’s lemma, it suffices to show that

$$\limsup_{i \uparrow \infty} E[X_{n+1} - X_n \mid X_n = i] < 0 \tag{1.16}$$

and for all i ,

$$E[X_{n+1} \mid X_n = i] < \infty. \tag{1.17}$$

For this, let Z_n be the number of fresh arrivals in the n th CRI. We have

$$\begin{aligned} E[X_{n+1} | X_n = i] &= \sum_{k=0}^{\infty} E[X_{n+1} | X_n = i, Z_n = k] P(Z_n = k | X_n = i) \\ &= \sum_{k=0}^{\infty} E[X_{n+1} | Z_n = k] P(Z_n = k | X_n = i). \end{aligned}$$

It will be shown that for all $n \geq 0$,

$$E[X_{n+1} | Z_n = k] \leq \alpha k + 1, \quad (1.18)$$

where $\alpha = 2.886$, and therefore

$$E[X_{n+1} | X_n = i] \leq \sum_{k=0}^{\infty} (\alpha k + 1) P(Z_n = k | X_n = i) = \alpha E[Z_n | X_n = i] + 1.$$

Using Wald's lemma (Theorem 3.2 of Chapter 1), we have

$$E[Z_n | X_n = i] = \lambda i,$$

where λ is the *traffic intensity*, and therefore

$$E[X_{n+1} - X_n | X_n = i] \leq 1 + i(\lambda\alpha - 1).$$

We see that condition (1.17) is always satisfied and that (1.16) is satisfied, provided that

$$\lambda < \frac{1}{\alpha} = 0.346. \quad (1.19)$$

It remains to prove (1.18). Let $E[X_{n+1} | Z_n = k] = L_k$ (it is indeed a quantity independent of n). Clearly,

$$L_0 = L_1 = 1,$$

since with zero or one packet at the beginning of a CRI, there is no collision. When $k \geq 2$, there is a collision, and the k users toss a coin, and depending on the result they split into two sets, layer 0 and layer 1. Among these k users, i obtain heads with probability

$$q_i(k) = \binom{k}{i} \left(\frac{1}{2}\right)^k.$$

The average length of the CRI given that there are $k \geq 2$ customers at the start, and given that the first layer 0 contains i messages, is

$$L_{k,i} = 1 + L_i + L_{k-i}.$$

Indeed, the first slot saw a collision; the i customers in the first layer 0 will take on the average L_i slots to resolve their collision, and L_{k-i} more slots will be needed for the $k - i$ customers in the first-formed layer 1 (these customers are always at the bottom of the stack,

in a layer traveling up and down until it becomes layer 0, at which time they start resolving their collision). Since

$$L_k = \sum_{i=0}^k q_i(k)L_{k,i} ,$$

we have

$$L_k = 1 + \sum_{i=0}^k q_i(k)(L_i + L_{k-i}).$$

Solving for L_k , we obtain

$$L_k = \frac{1 + \sum_{i=0}^{k-1} [q_i(k) + q_{k-i}(k)]L_i}{1 - q_0(k) - q_k(k)} . \tag{1.20}$$

Suppose that for some $m \geq 2$, and α_m satisfying

$$\alpha_m \geq \sup_{j>m} \frac{\sum_{i=0}^{m-1} (L_i + 1)(q_i(j) + q_{j-i}(j))}{\sum_{i=0}^{m-1} i(q_i(j) + q_{j-i}(j))} , \tag{1.21}$$

it holds that $L_m \leq \alpha_m m - 1$. Then we shall prove that for all $n \geq m$,

$$L_n \leq \alpha_m n - 1 . \tag{1.22}$$

We do this by induction, supposing that (1.22) holds true for $n = m, m + 1, \dots, j - 1$, and proving that it holds true for $n = j$. Equality (1.20) gives

$$\begin{aligned} L_j(1 - q_0(j) - q_j(j)) &= 1 + \sum_{i=0}^{j-1} (q_i(j) + q_{j-i}(j))L_i \\ &= 1 + \sum_{i=0}^{m-1} + \sum_{i=m}^{j-1} \\ &\leq 1 + \sum_{i=0}^{m-1} + \sum_{i=m}^{j-1} (q_i(j) + q_{j-i}(j))(\alpha_m i - 1) , \end{aligned}$$

where we used the induction hypothesis. The latter term equals

$$\begin{aligned} &1 + \sum_{i=0}^{m-1} (q_i(j) + q_{j-i}(j))(L_i - \alpha_m i + 1) \\ &+ \sum_{i=0}^j (q_i(j) + q_{j-i}(j))(\alpha_m i - 1) - (q_0(j) + q_j(j))(\alpha_m j - 1) \\ &= 1 + \sum_{i=0}^{m-1} (q_i(j) + q_{j-i}(j))(L_i - \alpha_m i + 1) + \alpha_m j - 2 - (q_0(j) + q_j(j))(\alpha_m j - 1) , \end{aligned}$$

where we used the identities

$$\sum_{i=0}^j q_i(j) = 1, \quad \sum_{i=0}^j i q_i(j) = jp, \quad \sum_{i=0}^j i q_{j-i}(j) = j(1-p).$$

Therefore,

$$L_j \leq (\alpha_m j - 1) + \frac{\sum_{i=0}^{m-1} (q_i(j) + q_{j-i}(j))(L_i - \alpha_m i + 1)}{1 - q_0(j) - q_j(j)}.$$

Therefore, for $L_j \leq \alpha_m j - 1$ to hold, it suffices to have

$$\sum_{i=0}^{m-1} (q_i(j) + q_{j-i}(j))(L_i - \alpha_m i + 1) \leq 0.$$

We require this to be true for all $j > m$, and (1.21) guarantees this. It can be checked *numerically* that for $m = 6$ and $\alpha_6 = 0.286$, (1.21) is satisfied and that equality (1.22) is true for $n = 1, 2, 3, 4, 5, 6$, and this completes the proof.

For more information on multiple-access communications, the reader is directed to the monograph (Rom and Sidi, 1990). \diamond

2 Martingales and Potentials

2.1 Harmonic Functions and Martingales

The concept of martingale is central to the modern theory of stochastic processes, and one of the objectives of the present chapter is to introduce the reader to martingale theory and its applications to Markov chains. These will be developed in Section 3. Closely connected with martingales are the harmonic functions and potentials.

Let $\{X_n\}_{n \geq 0}$ be an HMC on the countable space E with transition matrix \mathbf{P} . In the study of recurrence based on invariant measures or stationary distributions, the principal role is played by the equation $x^T = x^T \mathbf{P}$, where x^T is a row vector. The recurrence/transience criteria of the present chapter are based on the “dual” equation (resp., inequalities)

$$\mathbf{P}h = h, \text{ resp., } \geq h, \leq h, \quad (2.1)$$

where h is a column vector. In developed form, for all $i \in E$,

$$\sum_{j \in E} p_{ij} h(j) = h(i), \text{ resp., } \geq h(i), \leq h(i).$$

Equation (2.1) is equivalent to

$$E[h(X_{n+1}) | X_n = i] = h(i), \text{ resp., } \geq h(i), \leq h(i), \quad (2.2)$$

for all $i \in E$. In view of the Markov property, the left-hand side of the above equality is also equal to

$$E[h(X_{n+1}) \mid X_n = i, X_{n-1} = i_{n-1}, \dots, X_0 = i_0],$$

and therefore (2.1) is equivalent to

$$E[h(X_{n+1} \mid X_0^n)] = h(X_n), \text{ resp., } \leq h(X_n), \geq h(X_n). \quad (2.3)$$

This motivates the two following definitions.

Definition 2.1. *Harmonic, Subharmonic, Superharmonic*

A function $h : E \rightarrow \mathbb{R}$ is called *harmonic* (resp., *subharmonic*, *superharmonic*) iff

$$\mathbf{P}h = h, \text{ resp., } \geq h, \leq h. \quad (2.4)$$

Superharmonic functions are also called *excessive* functions. More generally, let $D \subset E$ be a subset of the state space, called the *domain*, and denote its complement in E , called the *boundary*, by ∂D . If $h : E \rightarrow \mathbb{R}$ satisfies (2.4) on D , then it is called *harmonic* (resp., *subharmonic*, *superharmonic*) on D .

Definition 2.2. *Martingales*

A real-valued stochastic process $\{Y_n\}_{n \geq 0}$ such that for each $n \geq 0$

- (i) Y_n is a function of X_0, \dots, X_n , and
- (ii) $E[|Y_n|] < \infty$ or $Y_n \geq 0$

is called a *martingale* (resp., *submartingale*, *supermartingale*) with respect to $\{X_n\}_{n \geq 0}$ if, moreover,

$$E[Y_{n+1} \mid X_0^n] = Y_n, \text{ resp., } \geq Y_n, \leq Y_n. \quad (2.5)$$

In the above definition, $\{X_n\}_{n \geq 0}$ can be any stochastic process, not necessarily a Markov chain. Also, observe that a martingale is a submartingale *and* a supermartingale.

Example 2.1. *Harmonic Functions Produce Martingales*

Let $\{X_n\}_{n \geq 0}$ be an HMC with transition matrix \mathbf{P} and state space E . The discussion at the beginning of the current subsection shows that if $h : E \rightarrow \mathbb{R}$ is either a function such that $E[|h(X_n)|] < \infty$ for all $n \geq 0$ or a nonnegative function, and if it is harmonic (resp., subharmonic, superharmonic), the process $\{h(X_n)\}_{n \geq 0}$ is, with respect to $\{X_n\}_{n \geq 0}$, a martingale (resp., submartingale, supermartingale). \diamond

Condition (2.5) implies that for all $k \geq 1$,

$$E[Y_{n+k} \mid X_0^n] = Y_n, \text{ resp., } \geq Y_n, \leq Y_n. \quad (2.6)$$

For instance, in the martingale case, with $k = 2$,

$$E[Y_{n+2} \mid X_0^n] = E[E[Y_{n+2} \mid X_0^{n+1}] \mid X_0^n] = E[Y_{n+1} \mid X_0^n] = Y_n. \quad \square$$

Example 2.2. *Lévy's Martingale*

Let $\{X_n\}_{n \geq 0}$ be an HMC with transition matrix \mathbf{P} and state space E , and let $f : E \rightarrow \mathbb{R}$ be a bounded function. The process

$$M_n^f = f(X_n) - f(X_0) - \sum_{k=0}^{n-1} (\mathbf{P} - I)f(X_k) \quad (2.7)$$

is a martingale with respect to $\{X_n\}_{n \geq 0}$. Indeed, since f is a bounded function, say bounded by K ,

$$|(\mathbf{P}f)(i)| = \left| \sum_{j \in E} p_{ij} f(j) \right| \leq K.$$

Therefore, $|M_n^f| \leq 2(n+1)K < \infty$; in particular, M_n is integrable. Also,

$$M_{n+1}^f - M_n^f = f(X_{n+1}) - \mathbf{P}f(X_n),$$

and therefore, since

$$E[f(X_{n+1}) | X_0^n] = E[f(X_{n+1}) | X_n] = \mathbf{P}f(X_n),$$

we have

$$E[M_{n+1}^f - M_n^f | X_0^n] = 0. \quad \diamond$$

This result admits a converse (see Problem 7.2.4).

2.2 The Maximum Principle

Potential theory is connected with martingale theory, Lyapunov functions, etc., and above all to the classical potential theory in analysis. The maximum principle is a big result of potential theory, and we give below one of its avatars when the state space is discrete. It is a good opportunity to show the deep and productive links between probability and analysis.

Let $\{X_n\}_{n \geq 0}$ be an HMC with countable state space E and transition matrix \mathbf{P} . Let D be an arbitrary subset of E , called the *domain*, and denote by ∂D the complement of D in E , which is called the *boundary*. Let $c : D \rightarrow \mathbb{R}$ and $\varphi : \partial D \rightarrow \mathbb{R}$ be nonnegative functions called the *unit time cost* and the *final cost*, respectively. Let T be the hitting time of ∂D .

For each state $i \in E$, define

$$h(i) = E_i \left[\sum_{0 \leq k < T} c(X_k) + \varphi(X_T) \mathbf{1}_{\{T < \infty\}} \right]. \quad (2.8)$$

The function $h : E \rightarrow \overline{\mathbb{R}}$ so defined is nonnegative and possibly infinite. It is called the *average cost*. Note that T is not required to be finite, and that ∂D may be empty.

Theorem 2.1. *The Maximum Principle*

Let $h : E \rightarrow \bar{\mathbb{R}}_+$ be defined by (2.8). Then

(i) h is nonnegative and satisfies

$$h = \begin{cases} \mathbf{P}h + c & \text{on } D, \\ \varphi & \text{on } \partial D. \end{cases} \quad (2.9)$$

(ii) Any nonnegative function $u : E \rightarrow \bar{\mathbb{R}}$ such that

$$u \geq \begin{cases} \mathbf{P}u + c & \text{on } D, \\ \varphi & \text{on } \partial D. \end{cases} \quad (2.10)$$

is a majorant of h , that is,

$$u \geq h. \quad (2.11)$$

(iii) If for all $i \in E$,

$$P_i(T < \infty) = 1, \quad (2.12)$$

then (2.9) has at most one nonnegative bounded solution.

Proof.

(i) Properties $h \geq 0$ and $h = \varphi$ on ∂D are satisfied by definition. First-step analysis gives, for $i \in D$,

$$h(i) = c(i) + \sum_{j \in E} p_{ij} h(j) \quad (2.13)$$

(rely on intuitive arguments or see the details after the proof).

(ii) Define for $n \geq 0$ the nonnegative function $h_n : E \rightarrow \mathbb{R}$ by

$$h_n(i) = E_i \left[\sum_{k=0}^{n-1} c(X_k) \mathbf{1}_{\{k < T\}} + \varphi(X_T) \mathbf{1}_{\{T < n\}} \right]. \quad (2.14)$$

Observe that $h_0 \equiv 0$ and, by monotone convergence, $\lim_{n \uparrow \infty} h_n = h$. Also, with a proof similar to that of (i),

$$h_{n+1} = \begin{cases} \mathbf{P}h_n + c & \text{on } D, \\ \varphi & \text{on } \partial D. \end{cases} \quad (2.15)$$

With u as in (2.10), we have $u \geq h_0$. By induction, $u \geq h_n$ (this is true for $n = 0$, and if this true for some n , it is true for $n + 1$. Indeed $u \geq \mathbf{P}u + c \geq \mathbf{P}h_n + c = h_{n+1}$ on D , and $u \geq \varphi = h_{n+1}$ on ∂D). Therefore, $u \geq \lim_{n \rightarrow \infty} h_n = h$.

(iii) If u is bounded and nonnegative, then by Example 2.2,

$$M_n = u(X_n) - u(X_0) - \sum_{k=0}^{n-1} (\mathbf{P} - I)u(X_k) \quad (2.16)$$

is a Lévy martingale with respect to $\{X_n\}_{n \geq 0}$. By the optional sampling theorem, for all integers $K \geq 0$, $E_i[M_{T \wedge K}] = E_i[M_0] = 0$, and therefore

$$u(i) = E_i[u(X_{T \wedge K})] - E_i\left[\sum_{k=0}^{T \wedge K-1} (\mathbf{P} - I)u(X_k)\right] = E_i[u(X_{T \wedge K})] + \sum_{k=0}^{T \wedge K-1} c(X_k),$$

since by hypothesis $(I - \mathbf{P})u = c$ on D . Since $P_i(T < \infty) = 1$, $\lim_{K \uparrow \infty} E_i[u(X_{T \wedge K})] = E_i[u(X_T)]$ by dominated convergence. But $u(X_T) = \varphi(X_T)$ because $u = \varphi$ on ∂D . Also, $\lim_{K \uparrow \infty} E_i\left[\sum_{k=0}^{T \wedge K-1} c(X_k)\right] = E_i\left[\sum_{k=0}^{T-1} c(X_k)\right]$ by monotone convergence. Finally,

$$u(i) = E_i\left[\sum_{k=0}^{T-1} c(X_k) + \varphi(X_T)\right] = h(i).$$

Proof of (2.13)

Write for $i \in D$,

$$\begin{aligned} v(i) &= E_i\left[c(X_0) + \sum_{1 \leq n < T} c(X_n) + \varphi(X_T)1_{\{T < \infty\}}\right] \\ &= c(i) + E_i\left[\sum_{1 \leq n < T} c(X_n) + \varphi(X_T)1_{\{T < \infty\}}\right], \end{aligned}$$

that is,

$$v(i) = c(i) + \sum_{j \in E} E_i[Z1_{\{X_1=j\}}],$$

where

$$Z = \sum_{1 \leq n < T} c(X_n) + \varphi(X_T)1_{\{T < \infty\}}.$$

Since $X_0 = i \in D$ implies that $T \geq 1$ on $\{X_0 = i\}$, the random variable Z is a function of X_1, X_2, \dots , and therefore, by the Markov property,

$$E_i[Z1_{\{X_1=j\}}] = E[Z | X_1 = j]p_{ij}.$$

Now, since $T \geq 1$ on $\{X_0 = i\}$ when $i \in D$, the quantity Z in the above calculations can be rewritten as

$$Z = \sum_{0 \leq n < T-1} c(Y_n) + \varphi(Y_{T-1})1_{\{T-1 < \infty\}},$$

where $Y_n = X_{n+1}$. Also, for the HMC $\{Y_n\}_{n \geq 0}$, $T' = T - 1$ is the hitting time of ∂D , and therefore

$$E[Z | X_1 = j] = E\left[\sum_{0 \leq n < T'} c(Y_n) + \varphi(Y_{T'})1_{\{T' < \infty\}} \mid Y_0 = j\right],$$

and this quantity is just $v(j)$, since $\{X_n\}_{n \geq 0}$ and $\{Y_n\}_{n \geq 0}$ have the same transition matrix, and therefore have the same distribution when their initial states are the same. \square

Theorem 2.1 can be rephrased as follows.

The function h given by (2.8) is a minorant of all nonnegative solutions of (2.10), and for $u = h$, the inequalities in (2.10) become equalities. Moreover, if h is bounded and $P_i(T < \infty) = 1$ for all $i \in E$, then h is the *unique* solution of (2.9).

Example 2.3. Dirichlet Problem

Let $h : E \rightarrow \mathbb{R}$ be a nonnegative bounded function that is *harmonic* on $D \subset E$, that is,

$$h = \mathbf{P}h \text{ on } D.$$

If $P_i(T < \infty) = 1$ for all $i \in E$, then h is entirely determined by its value on the boundary ∂D . To see this, call φ the restriction of h to ∂D . Then

$$h = \begin{cases} \mathbf{P}h & \text{on } D, \\ \varphi & \text{on } \partial D. \end{cases}$$

Since φ is bounded, the function $i \rightarrow E_i[\varphi(X_T)]$ is bounded, and therefore since $P_i(T < \infty) = 1$ for all $i \in E$, it is the unique nonnegative bounded solution of (2.9) where $c = 0$. Hence

$$h(i) = E_i[\varphi(X_T)]. \quad (2.17)$$

As an illustration, let \mathbf{P} be the transition matrix corresponding to a symmetric random walk on $E = \mathbb{Z}^2$; that is, the one-step transitions allowed are from $i = (i_1, i_2)$ to the four nearest states, and are equiprobable. Then, with $e_1 = (1, 0)$, $e_2 = (0, 1)$,

$$\begin{aligned} 4(\mathbf{P} - I)f(i) &= (f(i + e_1) - f(i)) - (f(i) - f(i - e_1)) \\ &\quad + (f(i + e_2) - f(i)) - (f(i) - f(i - e_2)), \end{aligned}$$

and call $\Delta f(i)$ this quantity. The function Δf is the (discrete) Laplacian of f . With $c = 0$, (2.9) becomes

$$\begin{cases} \Delta^2 h = 0 & \text{on } D, \\ h = \varphi & \text{on } \partial D. \end{cases} \quad (2.18)$$

and the reader will recognize here the Dirichlet problem of potential theory. In two dimensions, the symmetric random walk is irreducible and recurrent, and therefore, the hitting time T of ∂D is finite if ∂D is not empty. Therefore, in this case, (2.9) has at most one bounded solution. If φ is bounded, the solution is given by (2.17). More generally, if ∂D is not empty and $\sup_i E_i[\varphi(X_T)] < \infty$, then (2.17) gives the *unique* nonnegative bounded solution of the Dirichlet problem. \diamond

Example 2.4. Infinite Sojourn

Let $c \equiv 0$ in (2.9), and let $\varphi(k) = 1$ for all $k \in \partial D$. Then for $i \in D$,

$$v(i) = 1 - h(i) = P_i(T = \infty)$$

is the probability of infinite sojourn in D . By application of the maximum principle, we retrieve the results of Theorem 5.1 of Chapter 4. \diamond

Example 2.5. Optimal Control

We consider a stochastic process $\{X_n\}_{n \geq 0}$ with values in E , that is controlled in the following way. Let $\{\mathbf{P}(a); a \in A\}$, where A is some set, the set of *actions*, be a family of transition matrices on E , with the interpretation that, if at time n the controlled process is in state i , and if the controller takes action a , then at time $n + 1$ the state will be j with probability $p_{ij}(a)$. A *control strategy* u is a (measurable) function $u : E \rightarrow A$ which prescribes to take action $u(i)$ when the process is in state i . Therefore, under the strategy u , the controlled process is an HMC with transition matrix \mathbf{P}^u , where

$$p_{ij}^u = p_{ij}(u(i)).$$

There is a cost $V^u(i)$ associated with each strategy u and each initial state i , of the form

$$V^u(i) = E_i^u \left[\sum_{0 \leq k < T} c^u(X_k) + \varphi^u(X_T) 1_{\{T < \infty\}} \right],$$

where c^u , φ^u and T are as in Theorem 2.1, with D fixed, and moreover, $c^u(i) = c(i, u(i))$ and $\varphi^u(i) = \varphi(i, u(i))$, for appropriate functions c and φ . The problem of *optimal control* is that of finding, if it exists, an *optimal strategy* u^* , such that

$$V^{u^*}(i) \geq V^u(i),$$

for all states i , all strategies u .

We have the following result.

Suppose that there exists a function $V : E \rightarrow \mathbb{R}$ such that

$$V(i) = \sup_{a \in A} \left\{ \sum_{j \in E} p_{ij}(a) V(j) + c(i, a) \right\} \text{ for all } i \in D,$$

and

$$V(i) = \sup_{a \in A} \varphi(i, a) \text{ for all } i \in \partial D,$$

and that the suprema above are attained for $a = u^*(i)$, for some (measurable) function $u^* : E \rightarrow A$. Then, u^* is an optimal control and $V = V^{u^*}$.

Proof. Since for all controls u ,

$$V \geq \mathbf{P}^u V + c^u \text{ on } D,$$

and

$$V \geq \varphi^u \text{ on } \partial D,$$

it follows from the maximum principle that

$$V \geq V^u$$

for all controls u . Also, $V = V^{u^*}$ and therefore u^* is an optimal control.

This example shows the link between optimal control and potential theory. Optimal control is a vast and fascinating subject, and the reader is directed, for instance, to the book (Puterman, 1994) for a recent treatment and a bibliography. \diamond

3 Applications of Martingales to HMCs

3.1 The Two Pillars of Martingale Theory

One of the main results of martingale theory, which is the key to the recurrence (resp., transience) criteria of the next subsection, is the probabilistic counterpart of the convergence of a bounded nondecreasing sequence of real numbers to a finite limit.

Theorem 3.1. Martingale Convergence Theorem

Let $\{Y_n\}_{n \geq 0}$ be either a nonnegative supermartingale, or a bounded submartingale, with respect to $\{X_n\}_{n \geq 0}$. Then, almost surely, $\lim_{n \uparrow \infty} Y_n$ exists and is finite.

The proof is omitted; see (Williams, 1991) or (Shiryayev, 1987) for the full version. A first application to Markov chain theory of this result is the following:

Theorem 3.2.

An irreducible recurrent HMC has no nonnegative superharmonic or bounded subharmonic functions besides the constant functions.

Proof. If h is nonnegative superharmonic (resp., bounded subharmonic), then the stochastic sequence $\{h(X_n)\}_{n \geq 0}$ is a nonnegative supermartingale (resp., bounded submartingale), and therefore, by the martingale convergence theorem it converges to a finite limit Y . Since $\{X_n\}_{n \geq 0}$ visits any state $i \in E$ infinitely often, one must have $Y = h(i)$ almost surely for all $i \in E$. In particular, h is a constant. \square

The next result is another pillar of martingale theory. We give a weak version of it, which is sufficient for our purpose, and refer to the books (Williams, 1991) or (Shiryayev, 1987) for the full version.

Theorem 3.3. Optional Sampling Theorem

Let $\{M_n\}_{n \geq 0}$ be a martingale with respect to some process $\{X_n\}_{n \geq 0}$, and let T be a stopping time of $\{X_n\}_{n \geq 0}$. Suppose that at least one of the following condition holds:

(α) P-a.s, $T \leq n_0$ for some $n_0 \geq 0$.

(β) P-a.s, $T < \infty$ and $|M_n| \leq K < \infty$ when $n < T$.

Then

$$E[M_T] = E[M_0]. \tag{3.1}$$

Proof.

(α) Write

$$M_T - M_0 = \sum_{k=0}^{n_0-1} (M_{k+1} - M_k) \mathbf{1}_{\{k < T\}}.$$

Since T is a stopping time of $\{X_n\}_{n \geq 0}$,

$$\mathbf{1}_{\{k < T\}} = \varphi(X_0^k)$$

for some function φ , and therefore

$$E[(M_{k+1} - M_k)1_{\{k < T\}}] = E[(M_{k+1} - M_k)\varphi(X_0^k)] = 0.$$

Therefore,

$$E[M_T - M_0] = \sum_{k=0}^{n_0-1} E[(M_{k+1} - M_k)1_{\{k < T\}}] = 0.$$

(β) Apply the result of (α) to the stopping time $T \wedge n_0$ to obtain

$$E[M_{T \wedge n_0}] = E[M_0].$$

Therefore,

$$|E[M_T] - E[M_0]| = |E[M_T] - E[M_{T \wedge n_0}]| \leq 2K P(T > n_0).$$

Since T is finite, $\lim_{n_0 \rightarrow \infty} P(T > n_0) = 0$, and therefore $E[M_T] = E[M_0]$. \square

3.2 Transience and Recurrence via Martingales

Theorem 3.4. A Transience Criterion

A necessary and sufficient condition for an irreducible HMC to be transient is the existence of some state conventionally called 0 and of a bounded function $h : E \rightarrow \mathbb{R}$, not identically null and satisfying

$$h(j) = \sum_{k \neq 0} p_{jk} h(k), \text{ for all } j \neq 0. \quad (3.2)$$

Proof. Let T_0 be the return time to state 0. First-step analysis shows that the bounded function h defined by

$$h(j) = P_j(T_0 = \infty)$$

satisfies (3.2). If the chain is transient, h is nontrivial. This proves necessity.

Conversely, suppose that (3.2) holds for a not identically null bounded function. Define

$$\tilde{h}(j) = \begin{cases} h(j) & \text{if } j \neq 0, \\ 0 & \text{if } j = 0, \end{cases}$$

and let $\alpha = \sum_{k \in E} p_{0k} \tilde{h}(k)$. Changing signs if necessary, α can be assumed ≥ 0 . Then \tilde{h} is subharmonic. If the chain were recurrent, then by Theorem 3.2, \tilde{h} would be a constant. This constant would be equal to $\tilde{h}(0) = 0$, and this contradicts the assumed nontriviality of h . \square

Example 3.1. *Repair Shop, Once More*

As in Example 1.2, we assume irreducibility. The transition matrix is given in Example 2.2 of Chapter 2. We shall show that, if $E[Z_1] > 1$, the system of equations (3.2) admits a bounded nontrivial solution, and therefore, by Theorem 3.4, the chain is transient. Indeed, trying $y_j = 1 - \zeta^j$ for a solution, we can check that equations (3.2) reduce to a single equation in ζ ,

$$\sum_{k \geq 0} P(Z_1 = k) \zeta^k = \zeta, \quad (3.3)$$

for which (Theorem 5.1 of Chapter 1) there is, under condition $E[Z_1] > 1$ and the irreducibility condition, a solution $\zeta \in (0, 1)$. Therefore, $h(i) = \zeta^i$ is a solution of (3.2) that is nontrivial and bounded. \diamond

The next result is to be compared with Foster's theorem.

Theorem 3.5. *A Recurrence Criterion*

Let the HMC with transition matrix \mathbf{P} be irreducible, and suppose that there exists a function $h : E \rightarrow \mathbb{R}$ such that $\{i ; h(i) < K\}$ is finite for all finite K , and such that

$$\sum_{k \in E} p_{ik} h(k) \leq h(i), \text{ for all } i \notin F, \quad (3.4)$$

for some finite subset $F \subset E$. Then the chain is recurrent.

The conditions of the above result are also necessary (we shall not prove this here), and this is why it is called a criterion. Note that it might then as well be called a transience criterion.

Proof. Since $\{i ; h(i) < 0\}$ is finite, $\inf h(i) > -\infty$, and therefore, adding a constant if necessary, one may assume without loss of generality that $h \geq 0$. Let $\tau = \tau(F)$ be the return time to F , and define $Y_n = h(X_n)1_{\{n < \tau\}}$. The arguments in the proof of Foster's theorem show that for $i \notin F$, P_i -a.s.,

$$E_i[Y_{n+1} | X_0^n] \leq Y_n.$$

Therefore, $\{Y_n\}_{n \geq 0}$ is, under P_i , a nonnegative supermartingale with respect to $\{X_n\}_{n \geq 0}$. By the martingale convergence theorem, $\lim_{n \uparrow \infty} Y_n = Y_\infty$ exists and is finite, P_i -a.s.

Suppose, in view of contradiction, that the chain is transient. It must then visit any finite subset of the state space only a finite number of times. In particular, for arbitrary K , we can have $h(X_n) < K$ only for a finite (random) number of indices n . This implies that $\lim_{n \rightarrow \infty} h(X_n) = +\infty$, P_j -a.s. (for any $j \in E$). For this to be compatible with the fact that $\{1_{\{n < \tau\}} h(X_n)\}$ has P_i -a.s. a finite limit for $i \notin F$, we must have $P_i(\tau < \infty) = 1$.

In summary, $P_i(\tau < \infty) = 1$ for all $i \notin F$. Since F is finite, some state in F must be recurrent, hence the announced contradiction. \square

Example 3.2. *Repair Shop, Conclusion*

We know that the Repair Shop HMC of Examples 1.2 and 3.1 is positive recurrent if and only if $E[Z_1] < 1$, and that it is transient if $E[Z_1] > 1$. It remains to examine the case $E[Z_1] = 1$, for which there are only two possibilities left: transient or null recurrent. It turns out that the chain is null recurrent in this case. Indeed, one easily verifies that Theorem 3.5 applies with $h(i) = i$ and $F = \{0\}$. Therefore, the chain is recurrent. Since it is not positive recurrent, it is null-recurrent.

We are now through with the Repair Shop HMC. We have found that

$$P(Z_1 = 0) > 0 \text{ and } P(Z_1 \geq 2) > 0$$

is a necessary and sufficient condition of irreducibility, and that in this case

if $E[Z_1] < 1$, the chain is positive recurrent,

if $E[Z_1] = 1$, the chain is null recurrent,

if $E[Z_1] > 1$, the chain is transient. ◇

Here is another application of the martingale convergence theorem in the vein of the previous results.

Theorem 3.6. *A Sufficient Condition of Transience*

Let the HMC $\{X_n\}_{n \geq 0}$ with transition matrix \mathbf{P} be irreducible and let $h : E \rightarrow \mathbb{R}$ be a bounded function such that

$$\sum_{k \in E} p_{ik} h(k) \leq h(i), \text{ for all } i \notin F, \quad (3.5)$$

for some set F , not assumed finite. Suppose, moreover, that there exists $i \notin F$ such that

$$h(i) < h(j), \text{ for all } j \in F. \quad (3.6)$$

Then the chain is transient.

Proof. Let τ be the return time in F and let $i \notin F$ satisfy (3.6). Defining $Y_n = h(X_{n \wedge \tau})$, we have

$$E_i[Y_{n+1} | X_0^n] = E_i[1_{\{n < \tau\}} h(X_{n+1}) | X_0^n] + E_i[1_{\{n \geq \tau\}} h(X_\tau) | X_0^n].$$

The second term of the right-hand side of the above equality is $1_{\{n \geq \tau\}} h(X_\tau) = 1_{\{n \geq \tau\}} Y_n$ (observe that $1_{\{n \geq \tau\}} h(X_\tau)$ is a function of X_0^n , and use rule (7.7) of Chapter 1), whereas the first term is, in view of calculations already performed in the proof of Theorem 1.1, less than or equal to $1_{\{n < \tau\}} h(X_n) = 1_{\{n < \tau\}} Y_n$. Therefore, under P_i , $\{Y_n\}_{n \geq 0}$ is a (bounded) supermartingale with respect to $\{X_n\}_{n \geq 0}$. By the martingale convergence theorem, the limit Y of $Y_n = h(X_{n \wedge \tau})$ exists and is finite, P_i -almost surely. By bounded convergence, $E_i[Y] = \lim_{n \uparrow \infty} E_i[Y_n]$, and since $E_i[Y_n] \leq E_i[Y_0] = h(i)$ (supermartingale property), we have $E_i[Y] \leq h(i)$.

If τ were P_i -a.s. finite, then Y_n would eventually be frozen at a value $h(j)$ for $j \in F$, and therefore by (3.6), $E_i[Y] > h(i)$, a contradiction with the last inequality.

Therefore, $P_i(\tau < \infty) < 1$, which means that with a strictly positive probability, the chain starting from $i \notin F$ will not return to F . This is incompatible with irreducibility and recurrence. \square

3.3 Absorption via Martingales

Chapter 4 features a method of computation of absorption probabilities that is essentially algebraic. Other cases may require ingenuity (see Watson's solution of the branching process, Example 2.4, Chapter 2). Martingale theory, and in particular the martingale convergence theorem, can also be usefully applied, as the next two examples demonstrate.

Example 3.3. Branching Processes via Martingales

As in Example 2.4 of Chapter 2, it is assumed that $P(Z = 0) < 1$ and $P(Z \geq 2) > 0$. The stochastic process

$$Y_n = \frac{X_n}{m^n},$$

where m is the average number of sons of a given individual, is a martingale with respect to $\{X_n\}_{n \geq 0}$. Indeed, each of the X_n members of the n th generation gives on the average m sons, and they do this independently. Therefore, $E[X_{n+1}|X_n] = mX_n$, and

$$E\left[\frac{X_{n+1}}{m^{n+1}} \mid X_0^n\right] = E\left[\frac{X_{n+1}}{m^{n+1}} \mid X_n\right] = \frac{X_n}{m^n}.$$

By the martingale convergence theorem, almost surely

$$\lim_{n \uparrow \infty} \frac{X_n}{m^n} = Y < \infty.$$

In particular, if $m < 1$, then $\lim_{n \uparrow \infty} X_n = 0$ almost surely. Since X_n takes integer values, this implies that the branching process eventually becomes extinct.

If $m = 1$, then $\lim_{n \uparrow \infty} X_n = X_\infty < \infty$, and it is easily argued that this limit must be 0. Therefore, in this case as well the process eventually becomes extinct.

For the case $m > 1$, we consider the unique solution in $(0, 1)$ of $x = g(x)$ (Theorem 5.1 of Chapter 1). Suppose we can show that $Z_n = x^{X_n}$ is a martingale. Then, by the martingale convergence theorem, it converges to a finite limit, and therefore X_n has a limit X_∞ , which, however, can be infinite. One can easily argue that this limit cannot be other than 0 (extinction) or ∞ (nonextinction). Since $\{Z_n\}_{n \geq 0}$ is a martingale, $x = E[Z_0] = E[Z_n]$, and therefore, by dominated convergence, $x = E[Z_\infty] = E[x^{X_\infty}] = P(X_\infty = 0)$. Therefore, x is the probability of extinction.

It remains to show that $\{Z_n\}_{n \geq 0}$ is a martingale. We have

$$E[x^{X_{n+1}} \mid X_n = i] = x^i.$$

This is obvious if $i = 0$, and if $i > 0$, X_{n+1} is the sum of i independent random variables with the same generating function g . Therefore, $E[x^{X_{n+1}}|X_n = i] = g(x)^i = x^i$. From this last result and the Markov property,

$$E[x^{X_{n+1}}|X_0^n] = E[x^{X_{n+1}}|X_n] = x^{X_n}. \quad \diamond$$

Example 3.4. A Cellular Automaton

Consider a chessboard of size $N \times N$, on which are placed stones, exactly one on each square. Each stone has one among k possible colors. The state X_n of the process at time n is the $N \times N$ matrix with elements in $\{1, \dots, k\}$ describing the chessboard and the color of the stone in each square. The evolution of $\{X_n\}_{n \geq 0}$ is that of a homogeneous Markov chain, where the transition from X_n to X_{n+1} is as follows. Select one case of the chessboard at random, and change the color of the stone there, the new color being the color of a stone chosen at random among the 4 neighboring stones. To avoid boundary effects, we shall consider that the chessboard is a bi-torus in the sense of Figure 5.3.1.



Figure 5.3.1. Neighbors in the cellular automaton model

This chain has $2^k - 1$ communication classes. A communication class corresponds to a given nonempty subset of ℓ different colors. For instance, with $\ell = 3$, we consider all the configurations of the chessboard with a combination of three given colors, say blue, white and red. It is straightforward to verify that one can pass (in several steps) from a configuration with at least one stone of each color, blue, white or red, to any other such configuration. Any monochromatic state is of course closed.

Denote by Y_n the proportion of red stones at stage n . The process $\{Y_n\}_{n \geq 0}$ is a martingale with respect to $\{X_n\}_{n \geq 0}$. Indeed, Y_n is a function of X_n and is integrable, since it is bounded by 1. Also, $E[Y_{n+1}|X_0^n] = Y_n$, as the following exchange argument shows.

Let α_{n+1} be the box selected at time $n + 1$ and let β_{n+1} be the selected neighbor of α_{n+1} . Then, for any pair (α, β) of boxes, $P(\alpha_{n+1} = \alpha, \beta_{n+1} = \beta|X_0^n) = P(\alpha_{n+1} = \beta, \beta_{n+1} = \alpha|X_0^n) = \frac{1}{8N^2}$. Clearly, if the choice $\alpha_{n+1} = \alpha, \beta_{n+1} = \beta$ changes Y_n to $Y_{n+1} = Y_n + \Delta Y_{n+1}$, the choice $\alpha_{n+1} = \beta, \beta_{n+1} = \alpha$ changes Y_n to $Y_{n+1} = Y_n - \Delta Y_{n+1}$. Since these two situations are equiprobable, the martingale property easily follows.

By the martingale convergence theorem, $\lim_{n \uparrow \infty} Y_n = Y$ exists, and by dominated convergence $E[Y] = \lim_{n \uparrow \infty} E[Y_n]$. Therefore, since $E[Y_n] = E[Y_0]$, we have $E[Y] =$

$E[Y_0] = y_0$, where y_0 is the initial proportion of red stones. Because $|\Delta Y_n| = 0$ or $\frac{1}{N^2}$ for all n , $\{Y_n\}$ can converge only if it remains constant after some (random) time, and this constant is either 0 or 1. Since the limit 1 corresponds to absorption by the “all-red” state, we see that the probability of being absorbed by the “all-red” state is equal to the initial proportion of red stones. This analysis being true for all colors, we see that the monochromatic states are the only absorbing states. \diamond

Next example illustrates how the optional sampling theorem can be used to compute absorption probabilities.

Example 3.5. *Gambler’s Ruin Revisited*

Consider the symmetric random walk $\{X_n\}_{n \geq 0}$ on \mathbb{Z} with $X_0 = 0$. It is a martingale (with respect to itself). Let T be the first time n for which $X_n = -a$ or $+b$, where $a, b > 0$. This is a stopping time, and moreover $T < \infty$ (use Remark 5.2 of Chapter 4, for instance). We can apply the optional sampling theorem, part (β) , with $K = \sup(a, b)$, to obtain

$$0 = E[X_0] = E[X_T].$$

Writing

$$v = P(-a \text{ is hit before } b),$$

we have

$$E[X_T] = -av + b(1 - v),$$

and therefore

$$v = \frac{b}{a + b}. \quad \diamond$$

Example 3.6. *A Counterexample*

Consider the symmetric random walk with $X_0 = 0$, as in the previous example, but now define T to be the hitting time of $b > 0$. We know that $T < \infty$, since the symmetric walk on \mathbb{Z} is recurrent. If the optional sampling theorem applied, we would have

$$0 = E[X_0] = E[X_T] = b,$$

an obvious contradiction. The optional sampling theorem does not apply because neither condition (α) nor (β) thereof is satisfied. \diamond

Problems

5.1.1 Consider the following model of a waiting line. At the beginning of each time period, exactly one customer enters the system. In period $n \geq 1$, the service capacity is Z_{n+1} , that is to say, up to Z_{n+1} customers can be served. This means that if there are X_n customers at the beginning of period n , there will be $(X_n - Z_{n+1})^+ + 1$ at the beginning of period

$n + 1$. It is assumed that the sequence $\{Z_n\}_{n \geq 1}$ is i.i.d and independent of the initial number of customers, so that $\{X_n\}_{n \geq 0}$ is an HMC. Give a necessary and sufficient condition of irreducibility. Assuming irreducibility, show that if $E[Z_1] > 1$, then the HMC is positive recurrent.

5.1.2 Consider the HMC with integer values and transition matrix

$$\begin{aligned} p_{01} &= 1, \\ p_{i,i+k} &= \frac{1}{2^{i+1}} \frac{i!}{k!(i-k)!}, \quad i \geq 1, \quad 0 \leq k \leq i, \\ p_{i0} &= \frac{1}{2}, \quad \text{for } i \geq 1. \end{aligned}$$

Show that it is ergodic (give *two* proofs). Compute the mean time between two visits to 0.

5.1.3 The purpose of this exercise is to show that condition (1.1) is important. Apply Foster's theorem to the *forward recurrence chain* of Problem 2.6.1.

5.1.4 Show that in Example 1.1,

$$E_i[\tau] \leq \frac{i}{-E[Z_1]}.$$

Give similar estimates for Example 1.2 and Problem 5.1.1.

5.2.1 (α) Let $\{X_n\}_{n \geq 0}$ be an HMC with state space E and transition matrix \mathbf{P} . Let $f : \mathbb{N} \times E \rightarrow \mathbb{R}$ be a function such that for all $n \geq 0$ and $i \in E$,

$$E[|f(n, X_n)|] < \infty$$

and

$$(\mathbf{P}f)(n+1, i) \stackrel{\text{def}}{=} \sum_{j \in E} p_{ij} f(n+1, j) = f(n, i).$$

Show that

$$M_n = f(n, X_n)$$

defines a martingale $\{M_n\}_{n \geq 0}$ with respect to $\{X_n\}_{n \geq 0}$.

(β) Let $\{X_n\}_{n \geq 0}$ be a symmetric random walk on \mathbb{Z} . Show that $\{X_n\}_{n \geq 0}$ and $\{X_n^2 - n\}_{n \geq 0}$ are martingales with respect to $\{X_n\}_{n \geq 0}$.

5.2.2 Let $\{X_n\}_{n \geq 0}$ be an HMC with state space E , and let B be a closed subset of states, that is,

$$i \in B \Rightarrow \sum_{j \in B} p_{ij} = 1.$$

Let T be the hitting time of B , and define for $i \in E$,

$$h(i) = P_i(T < \infty).$$

Show that $\{h(X_n)\}_{n \geq 0}$ is a martingale with respect to $\{X_n\}_{n \geq 0}$.

5.2.3 Let $\{X_n\}_{n \geq 0}$ be a stochastic process with values in E , which is not assumed to be an HMC. Let \mathbf{P} be some transition matrix on E . Prove that if for all bounded $f : E \rightarrow \mathbb{R}$, $\{M_n^f\}_{n \geq 0}$ defined by (2.7) is a martingale with respect to $\{X_n\}_{n \geq 0}$, then $\{X_n\}_{n \geq 0}$ is an HMC with transition matrix \mathbf{P} .

5.2.4 Let $\{X_n\}_{n \geq 1}$ be an HMC on E , with transition matrix \mathbf{P} . Let $D \subset E$, and suppose that the boundary ∂D is accessible in finite time T from all states. Show that any nonnegative bounded function $h : \partial D \rightarrow \mathbb{R}$ that is harmonic on D is of the form

$$h(i) = \sum_{k \in \partial D} h(k) P_i(\text{the chain enters } \partial D \text{ through } k).$$

5.3.1 Show that the function $h(i) = \left(\frac{q}{p}\right)^i$ is harmonic for the random walk on \mathbb{Z} with $p_{i,i+1} = p$, $p_{i,i-1} = q = 1 - p$, where $p \in (0, 1)$, $p \neq \frac{1}{2}$. Apply the martingale convergence theorem to obtain the ruin probability in the ruin problem of Example 3.1, Chapter 2.

5.3.2 Consider the homogeneous Markov chain with state space $E = \{0, 1, \dots, m\}$ and transition probabilities

$$p_{ij} = \binom{m}{j} \left(\frac{i}{m}\right)^j \left(1 - \frac{i}{m}\right)^{m-j}.$$

In particular, 0 and m are absorbing states. Compute the probability of absorption by 0.

5.3.3 In Example 3.5, compute $E[T]$ (Hint: use (β) of Problem 5.2.1).

Eigenvalues and Nonhomogeneous Markov Chains

1 Finite Transition Matrices

1.1 Perron–Frobenius Theorem

When the state space is finite, we can rely on the standard results of linear algebra to study the asymptotic behavior of homogeneous Markov chains. Indeed, the asymptotic behavior of the distribution at time n of the chain is entirely described by the asymptotic behavior of the n -step transition matrix \mathbf{P}^n , and the latter depends on the eigenstructure of \mathbf{P} . The Perron–Frobenius theorem detailing the eigenstructure of nonnegative matrices is therefore all that is needed, at least in the theory.

The principal result of Perron and Frobenius is that convergence to steady state of an ergodic finite state space HMC is geometric, with relative speed equal to the second-largest eigenvalue modulus (SLEM). It is true that there are a number of interesting models, especially in biology, where the eigenstructure of the transition matrix can be extracted; see, for instance, (Iosifescu, 1980) or (Karlin and Taylor, 1975). Nevertheless, this situation remains exceptional. It is therefore important to find estimates, more precisely, upper and lower bounds, of the SLEM. This is one of the topics of the present chapter. Most methods presented in this chapter to achieve this goal use matrix algebra and graph theory, with the notable exception of a probabilistic method based on the notion of strong stationary times.

We shall also consider the problem of computing the variance of ergodic estimates, in terms of the eigenstructure of the transition matrix. The object of interest here is the fundamental matrix of finite irreducible transition matrices.

All the results collected in the present chapter are relevant to Monte Carlo simulation, a topic treated in Chapter 7, where the basic theory of simulated annealing will be given. In preparation for the latter, we give in the present chapter the essentials of the theory of nonhomogeneous Markov chains, based on Dobrushin's ergodic coefficient, another upper bound of the SLEM.

The basic results of the theory of matrices relative to eigenvalues and eigenvectors are reviewed in the appendix, from which we quote the following result, relative to a square matrix A of dimension r with *distinct* eigenvalues.

Let $\lambda_1, \dots, \lambda_r$ be the r distinct eigenvalues and let u_1, \dots, u_r and v_1, \dots, v_r be the associated sequences of left and right eigenvectors, respectively. Then, u_1, \dots, u_r form an independent collection of vectors, and so do v_1, \dots, v_r . Also, $u_i^T v_j = 0$ if $i \neq j$. Since eigenvectors are determined up to multiplication by an arbitrary nonnull scalar, one can choose them in such a way that $u_i^T v_i = 1$ for all $i \in [1, r]$. We then have the spectral decomposition

$$A^n = \sum_{i=1}^r \lambda_i^n v_i u_i^T. \quad (1.1)$$

Example 1.1. *Two-State Chain*

Consider the transition matrix on $E = \{1, 2\}$

$$\mathbf{P} = \begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix},$$

where $\alpha, \beta \in (0, 1)$. Its characteristic polynomial $(1 - \alpha - \lambda)(1 - \beta - \lambda) - \alpha\beta$ admits the roots $\lambda_1 = 1$ and

$$\lambda_2 = 1 - \alpha - \beta.$$

Observe at this point that $\lambda = 1$ is always an eigenvalue of a stochastic $r \times r$ matrix \mathbf{P} , associated with the right eigenvector $v = \mathbf{1}$ with all entries equal to 1, since $\mathbf{P}\mathbf{1} = \mathbf{1}$. Also, the stationary distribution

$$\pi^T = \left(\frac{\beta}{\alpha + \beta}, \frac{\alpha}{\alpha + \beta} \right)$$

is the left eigenvector corresponding to the eigenvalue 1. In this example, the representation (1.1) takes the form

$$\mathbf{P}^n = \frac{1}{\alpha + \beta} \begin{pmatrix} \beta & \alpha \\ \beta & \alpha \end{pmatrix} + \frac{(1 - \alpha - \beta)^n}{\alpha + \beta} \begin{pmatrix} \alpha & -\alpha \\ -\beta & -\beta \end{pmatrix},$$

and therefore, since $|1 - \alpha - \beta| < 1$,

$$\lim_{n \uparrow \infty} \mathbf{P}^n = \frac{1}{\alpha + \beta} \begin{pmatrix} \beta & \alpha \\ \beta & \alpha \end{pmatrix}.$$

In particular, the result of convergence to steady state,

$$\lim_{n \uparrow \infty} \mathbf{P}^n = \mathbf{1}\pi^T = \mathbf{P}^\infty,$$

was recovered for this special case in a purely algebraic way. In addition, this algebraic method gives the convergence speed, which is exponential and determined by the second-largest eigenvalue modulus:

$$(\mathbf{P}^n - \mathbf{P}^\infty) = \frac{(1 - \alpha - \beta)^n}{\alpha + \beta} \begin{pmatrix} \alpha & -\alpha \\ -\beta & -\beta \end{pmatrix}.$$

This is a general fact, which follows from the Perron–Frobenius theory of nonnegative matrices below. \diamond

Definition 1.1. Nonnegative Matrices

A matrix $A = \{a_{ij}\}_{1 \leq i, j \leq r}$ with real coefficients is called *nonnegative* (resp., *positive*) if all its entries are nonnegative (resp., positive). A nonnegative matrix A is called *stochastic* if $\sum_{j=1}^r a_{ij} = 1$ for all i , and *substochastic* if $\sum_{j=1}^r a_{ij} \leq 1$ for all i , with strict inequality for at least one i .

Nonnegativity (resp., positivity) of A is denoted by $A \geq 0$ (resp., $A > 0$). If A and B are two matrices of the same dimensions with real coefficients, the notation $A \geq B$ (resp., $A > B$) means that $A - B \geq 0$ (resp., $A - B > 0$).

The *communication graph* of a square nonnegative matrix A is the oriented graph with the state space $E = \{1, \dots, r\}$ as its set of vertices and an oriented edge from vertex i to vertex j if and only if $a_{ij} > 0$.

Definition 1.2. Primitive Matrices

A nonnegative square matrix A is called *irreducible* (resp., *irreducible aperiodic*) if it has the same communication graph as an irreducible (resp., irreducible aperiodic) stochastic matrix. It is called *primitive* if there exists an integer k such that $A^k > 0$.

Example 1.2.

A nonnegative matrix is primitive if and only if it is irreducible and aperiodic (Problem 2.4.4). \diamond

Theorem 1.1. Perron–Frobenius Theorem

Let A be a *nonnegative primitive* $r \times r$ matrix. There exists a real eigenvalue λ_1 with algebraic as well as geometric multiplicity one such that $\lambda_1 > 0$, and $\lambda_1 > |\lambda_j|$ for any other eigenvalue λ_j . Moreover, the left eigenvector u_1 and the right eigenvector v_1 associated with λ_1 can be chosen positive and such that $u_1^T v_1 = 1$.

Let $\lambda_2, \lambda_3, \dots, \lambda_r$ be the eigenvalues of A other than λ_1 ordered in such a way that

$$\lambda_1 > |\lambda_2| \geq \dots \geq |\lambda_r| \tag{1.2}$$

and if $|\lambda_2| = |\lambda_j|$ for some $j \geq 3$, then $m_2 \geq m_j$, where m_j is the algebraic multiplicity of λ_j . Then

$$A^n = \lambda_1^n v_1 u_1^T + O(n^{m_2-1} |\lambda_2|^n), \quad (1.3)$$

where $O(f(n))$ represents a function of n such that there exists $\alpha, \beta \in \mathbb{R}$, $0 < \alpha \leq \beta < \infty$, such that $\alpha f(n) \leq O(f(n)) \leq \beta f(n)$ for all n sufficiently large.

If in addition, A is stochastic (resp., substochastic), then $\lambda_1 = 1$ (resp., $\lambda_1 < 1$).

If A is stochastic but not irreducible, then the algebraic and geometric multiplicities of the eigenvalue 1 are equal to the number of communication classes.

If A is stochastic and irreducible with period $d > 1$, then there are exactly d distinct eigenvalues of modulus 1, namely the d th roots of unity, and all other eigenvalues have modulus strictly less than 1.

For the proof, see (Seneta, 1981) or (Gantmacher, 1959). □

Example 1.3. *Rates of Convergence via the Perron–Frobenius Theorem*

If \mathbf{P} is a transition matrix on $E = \{1, \dots, r\}$ that is irreducible and aperiodic, and therefore primitive, then

$$v_1 = \mathbf{1}, \quad u_1 = \pi,$$

where π is the unique stationary distribution. Therefore

$$\mathbf{P}^n = \mathbf{1}\pi^T + O(n^{m_2-1} |\lambda_2|^n), \quad (1.4)$$

which generalizes the observation in Example 1.1. ◇

Example 1.4.

The characteristic polynomial of the doubly stochastic matrix

$$\mathbf{P} = \frac{1}{12} \begin{pmatrix} 0 & 6 & 6 \\ 4 & 3 & 5 \\ 8 & 3 & 1 \end{pmatrix}$$

is

$$\det(\lambda I - \mathbf{P}) = (\lambda - 1)\left(\lambda + \frac{1}{6}\right)\left(\lambda + \frac{1}{2}\right).$$

Since this matrix is doubly stochastic, $v_1 = \frac{1}{3}(1, 1, 1)^T$ and $u_1 = (1, 1, 1)^T$ are a right eigenvector and a left eigenvector, respectively, corresponding to the eigenvalue $\lambda_1 = 1$ and such that $u_1^T v_1 = 1$. Elementary computations yield for the right and left eigenvectors corresponding to $\lambda_2 = -\frac{1}{2}$ and $\lambda_3 = -\frac{1}{6}$, that

$$u_2 = \frac{1}{12}(2, -1, -1)^T, \quad v_2 = (4, 1, -5)^T, \quad u_3 = \frac{1}{4}(-2, 3, -1)^T, \quad v_3 = (0, 1, -1)^T,$$

where, here again, $u_2^T v_2 = u_3^T v_3 = 1$. Therefore,

$$P^n = \frac{1}{3} \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix} + \left(-\frac{1}{2}\right)^n \frac{1}{12} \begin{pmatrix} 8 & -4 & -4 \\ 2 & -1 & -1 \\ -10 & 5 & 5 \end{pmatrix} + \left(-\frac{1}{6}\right)^n \frac{1}{4} \begin{pmatrix} 0 & 0 & 0 \\ -2 & 3 & -1 \\ 2 & -3 & 1 \end{pmatrix}.$$

The convergence to the steady state is geometric with relative speed $\frac{1}{2}$. \diamond

1.2 Quasi-stationary Distributions

Let $\{X_n\}_{n \geq 0}$ be an HMC with finite state space E . Suppose that the set of recurrent states R and the set of transient states T are both nonempty. In the block decomposition of P with respect to the partition $R + T = E$,

$$P = \begin{pmatrix} D & 0 \\ B & Q \end{pmatrix}$$

the matrix Q is substochastic, since B is not identically null (otherwise, the transient set would be closed, and therefore recurrent, being finite). We assume, in addition, that Q is irreducible and aperiodic. Therefore,

$$Q^n = \lambda_1^n v_1 u_1^T + O(n^{m_2-1} |\lambda_2|^n), \tag{1.5}$$

where λ_1, v_1, u_1, m_2 , and λ_2 are as in Theorem 1.1 with $A = Q$. In particular, $\lambda_1 \in (0, 1)$ and $|\lambda_2| < \lambda_1$.

Let $\nu = \inf \{n \geq 0; X_n \in R\}$ be the entrance time into R . Recall that ν is almost surely finite, since T is a finite set (see Remark 5.1 of Chapter 4). We want to find the distribution of X_n for large n , conditioned by the fact that X_n is still in T . For this we compute for $i, j \in T$,

$$P_i(X_n = j \mid \nu > n) = \frac{P_i(X_n = j, \nu > n)}{P_i(\nu > n)} = \frac{P_i(X_n = j)}{P_i(X_n \in T)}.$$

Therefore,

$$P_i(X_n = j \mid \nu > n) = \frac{p_{ij}(n)}{\sum_{k \in T} p_{ik}(n)}.$$

In view of (1.5),

$$p_{ik}(n) = \lambda_1^n v_1(i) u_1(k) + O(n^{m_2-1} |\lambda_2|^n).$$

Therefore,

$$P_i(X_n = j \mid \nu > n) = \frac{u_1(j)}{\sum_{k \in T} u_1(k)} + O\left(n^{m_2-1} \left|\frac{\lambda_2}{\lambda_1}\right|^n\right), \tag{1.6}$$

and in particular,

$$\lim_{n \uparrow \infty} P_i(X_n = j \mid \nu > n) = \frac{u_1(j)}{\sum_{k \in T} u_1(k)}. \tag{1.7}$$

The probability distribution $\{u_1(i) / \sum_{k \in T} u_1(k)\}_{i \in T}$ is called the *quasi-stationary distribution* of the chain relative to T .

Quasi-stationary distributions were introduced by Bartlett (1957). The reader is referred to (Iosifescu, 1980) or (Seneta, 1981) for additional information on this topic. We shall only give a small illustrative example.

Example 1.5. Sibmating

Let us return to Example 6.3 of Chapter 4. The characteristic equation of \mathbf{Q} is

$$\left(\lambda - \frac{1}{2}\right)\left(\lambda^3 - \frac{3}{4}\lambda^2 - \frac{1}{8}\lambda + \frac{1}{16}\right) = 0,$$

and its eigenvalues are, properly ordered in descending order.

$$\lambda_1 = \frac{1 + \sqrt{5}}{4}, \quad \lambda_2 = \frac{1}{2}, \quad \lambda_3 = \frac{1 - \sqrt{5}}{4}, \quad \lambda_4 = \frac{1}{4}.$$

If we define U (resp., V) to be the matrix whose columns are left eigenvectors (resp., right eigenvectors) of \mathbf{Q} , placed in the same order as the corresponding eigenvalues, one then has

$$U^T = \begin{pmatrix} 6 + 2\sqrt{5} & 4 + 4\sqrt{5} & 6 + 2\sqrt{5} & 2 \\ 2 & 0 & -2 & 0 \\ 6 - 2\sqrt{5} & 4 - 4\sqrt{5} & 6 - 2\sqrt{5} & 2 \\ 2 & -2 & 2 & -1 \end{pmatrix}$$

and

$$V = \begin{pmatrix} 1 & 1 & 1 & 1 \\ \sqrt{5} - 1 & 0 & -\sqrt{5} - 1 & -1 \\ 1 & -1 & 1 & 1 \\ 6 - 2\sqrt{5} & 0 & 6 + 2\sqrt{5} & -4 \end{pmatrix}.$$

Note that we have not normalized these vectors in order to have $u_i^T v_i = 1$ for all $i \in [1, 4]$. If we do this for the eigenvectors corresponding to λ_1 , we then have the new eigenvectors

$$u_1^T = \frac{1}{9 + 4\sqrt{5}}(3 + \sqrt{5}, 2 + 2\sqrt{5}, 3 + \sqrt{5}, 1), \quad v_1^T = \frac{9 + 4\sqrt{5}}{20}(1, \sqrt{5} - 1, 1, 6 - 2\sqrt{5}).$$

Note that in addition we have $\sum_{k=1}^4 u_i(k) = 1$, so that u_1^T is in fact the quasi-stationary distribution. Here the multiplicity of λ_2 is $m_2 = 1$, and therefore the rate of convergence to the quasi-stationary distribution is geometric, with the relative speed

$$\frac{\lambda_2}{\lambda_1} = \frac{2}{1 + \sqrt{5}}.$$

◇

2 Reversible Transition Matrices

2.1 Eigenstructure and Diagonalization

As we know from the Perron–Frobenius theorem, the second-largest eigenvalue modulus (SLEM) of an ergodic matrix gives the rate of convergence to equilibrium. Unfortunately, the eigenstructure of a transition matrix is in general difficult to obtain. However, for reversible transition matrices, the added structure allows one to push the analysis much further.

For convenience, we recall the definition of reversibility, and introduce a slight change in the terminology.

Definition 2.1. *Reversibility*

Let \mathbf{P} be a transition matrix and π a strictly positive probability vector on E . The pair (\mathbf{P}, π) is called *reversible* if the detailed balance equations

$$\pi(i)p_{ij} = \pi(j)p_{ji} \quad (2.1)$$

hold for all $i, j \in E$.

This implies, in particular, that π is a stationary distribution of \mathbf{P} , the unique one if \mathbf{P} is irreducible.

We shall assume in the sequel that E is a finite state space, say $\{1, 2, \dots, r\}$, and that \mathbf{P} is irreducible, and therefore positive recurrent. For short we shall sometimes say, as we used to, that \mathbf{P} is reversible.

We now introduce two convenient norms on \mathbb{R}^r that are linked to the stationary distribution.

Let $\ell^2(\pi)$, where π is a strictly positive probability distribution on E , be the real vector space \mathbb{R}^r endowed with the scalar product

$$\langle x, y \rangle_\pi \stackrel{\text{def}}{=} \sum_{i \in E} x(i)y(i)\pi(i) \quad (2.2)$$

and the corresponding norm

$$\|x\|_\pi \stackrel{\text{def}}{=} \left(\sum_{i \in E} x(i)^2 \pi(i) \right)^{\frac{1}{2}}.$$

We shall write

$$\langle x \rangle_\pi \stackrel{\text{def}}{=} \langle x, 1 \rangle_\pi$$

the *mean* of x with respect to π . Also,

$$\text{Var}_\pi(x) \stackrel{\text{def}}{=} \|x\|_\pi^2 - \langle x \rangle_\pi^2$$

is the *variance* of x with respect to π .

We shall also need $\ell^2(\frac{1}{\pi})$, the real vector space \mathbb{R}^E endowed with the scalar product

$$\langle x, y \rangle_{\frac{1}{\pi}} \stackrel{\text{def}}{=} \sum_{i \in E} x(i)y(i) \frac{1}{\pi(i)}. \quad (2.3)$$

Theorem 2.1. *Necessary and Sufficient Condition of Reversibility*

The pair (\mathbf{P}, π) is reversible if and only if \mathbf{P} is self-adjoint in $\ell^2(\pi)$, that is,

$$\langle \mathbf{P}x, y \rangle_{\pi} = \langle x, \mathbf{P}y \rangle_{\pi} \quad (2.4)$$

for all $x, y \in \ell^2(\pi)$.

Proof. Suppose (\mathbf{P}, π) is reversible. Then

$$\begin{aligned} \langle \mathbf{P}x, y \rangle_{\pi} &= \sum_{i \in E} \left\{ \left(\sum_{j \in E} p_{ij}x(j) \right) y(i)\pi(i) \right\} \\ &= \sum_{i, j \in E} \pi(i)p_{ij}x(j)y(i) \\ &= \sum_{i, j \in E} \pi(j)p_{ji}y(i)x(j) \\ &= \sum_{j \in E} \left\{ x(j) \left(\sum_{i \in E} p_{ji}y(i) \right) \pi(j) \right\} = \langle x, \mathbf{P}y \rangle_{\pi}. \end{aligned}$$

Conversely, suppose \mathbf{P} self-adjoint in $\ell^2(\pi)$. Define for all $k \in E$, $\delta_k = (0, \dots, 0, 1, 0, \dots, 0)^T$, where the entry 1 is in the k th position. Then (2.1) follows from the choice $x = \delta_i, y = \delta_j$ in (2.4). \square

Although $\ell^2(\pi)$ is a convenient framework, we can also express the reversibility of (\mathbf{P}, π) more directly, by saying that

$$\mathbf{P}^* = D^{\frac{1}{2}} \mathbf{P} D^{-\frac{1}{2}}$$

is a symmetric matrix, where

$$D = \text{diag} \{ \pi(1), \dots, \pi(r) \}.$$

Note that

$$x^T D y = \langle x, y \rangle_{\pi}. \quad (2.5)$$

Since \mathbf{P}^* is symmetric, its eigenvalues are real, it is diagonalizable, and the sets of right and left eigenvectors are the same.

Choose an *orthonormal* basis of \mathbb{R}^r formed of right eigenvectors w_1, \dots, w_r associated, respectively, with the eigenvalues $\lambda_1, \dots, \lambda_r$. Define u and v by

$$w = D^{-\frac{1}{2}}u, \quad w = D^{\frac{1}{2}}v,$$

where w is a right (and therefore left) eigenvector of \mathbf{P}^* , corresponding to the eigenvalue λ . In particular,

$$u = Dv. \quad (2.6)$$

The matrices \mathbf{P} and \mathbf{P}^* have the same eigenvalues, and moreover, v (resp., u) is a right eigenvector (resp., left eigenvector) of \mathbf{P} corresponding to the eigenvalue λ .

Orthonormality of $\{w_1, \dots, w_r\}$ is with respect to the usual Euclidean norm, and it is equivalent to orthonormality in $\ell^2(\pi)$ of $\{v_1, \dots, v_r\}$, where $w_i = D^{\frac{1}{2}}v_i$, that is,

$$\langle v_i, v_j \rangle_\pi = \delta_{ij}. \quad (2.7)$$

Similarly,

$$\langle u_i, u_j \rangle_{\frac{1}{\pi}} = \delta_{ij}. \quad (2.8)$$

Recall that $u_1 = \pi$ and $v_1 = \mathbf{1}$.

Since $\{v_1, \dots, v_r\}$ is also a basis of \mathbb{R}^r , any vector $x \in \mathbb{R}^r$ can be expressed as $x = \sum_{i \in E} \alpha_i v_i$. Taking the scalar product in $\ell^2(\pi)$ with v_j gives $\langle x, v_j \rangle_\pi = \alpha_j$, and therefore

$$x = \sum_{j=1}^r \langle x, v_j \rangle_\pi v_j. \quad (2.9)$$

Similarly,

$$x^T = \sum_{j=1}^r \langle x, u_j \rangle_{\frac{1}{\pi}} u_j^T. \quad (2.10)$$

Moreover, for all n , $\mathbf{P}^n v_j = \lambda_j^n v_j$, and therefore

$$\mathbf{P}^n x = \sum_{j=1}^r \lambda_j^n \langle x, v_j \rangle_\pi v_j. \quad (2.11)$$

Similarly,

$$x^T \mathbf{P}^n = \sum_{j=1}^r \lambda_j^n \langle x, u_j \rangle_{\frac{1}{\pi}} u_j^T. \quad (2.12)$$

From (2.11), (2.5), and (2.6), we obtain

$$\mathbf{P}^n x = \sum_{j=1}^r \lambda_j^n v_j u_j^T x,$$

and therefore we retrieve representation (1.1) for $A = \mathbf{P}$.

Another way of writing (1.1) in this case is

$$\mathbf{P}^n = V \Lambda^n U^T = D^{-\frac{1}{2}} W \Lambda^n W^T D^{\frac{1}{2}}, \quad (2.13)$$

where $\Lambda = \text{diag} \{\lambda_1, \dots, \lambda_r\}$, $W = [w_1, \dots, w_r]$, and with a similar definition for U and V . Also, recall that $W W^T = I$.

The first equality in (2.13) is a particular case of a general result recalled in the appendix without proof.

2.2 Spectral Theorem

Coming back to the eigenvalues of \mathbf{P} , we know that $\lambda_1 = 1$ is one of them, with multiplicity 1 if and only if \mathbf{P} is irreducible. This eigenvalue corresponds to the unique right eigenvector v_1 such that $\|v_1\|_\pi = 1$, namely $v_1 = \mathbf{1}$. Moreover, the eigenvalues of \mathbf{P} are all in the closed unit disk of \mathbb{C} , and in the reversible case of interest in this section, they are real. Therefore, with proper ordering,

$$1 = \lambda_1 > \lambda_2 \geq \dots \geq \lambda_r \geq -1. \quad (2.14)$$

Note that this order is different from the one adopted in (1.2) for the statement of the Perron–Frobenius theorem. In (2.14), λ_2 is the second-largest eigenvalue (SLE), whereas in (1.2) it was the eigenvalue with the second-largest *modulus*.

The strict inequality $\lambda_1 > \lambda_2$ expresses the fact that λ_1 is the unique eigenvalue equal to 1. We also know from the Perron–Frobenius theorem that the only eigenvalue of modulus 1 and not equal to 1, in this case -1 , occurs if and only if the chain is periodic of period $d = 2$. In particular, in the reversible case, the period cannot exceed 2.

It will be convenient to consider the matrix $I - \mathbf{P}$, also called the *Laplacian* of the HMC. It has for eigenvalues $\beta_i = 1 - \lambda_i$, $i \in [1, r]$, and therefore

$$0 = \beta_1 < \beta_2 \leq \dots \leq \beta_r \leq 2.$$

Clearly, the right eigenvector of $I - \mathbf{P}$ corresponding to $\beta_i = 1 - \lambda_i$ is v_i .

Definition 2.2. The Dirichlet Form

The *Dirichlet form* $\mathcal{E}_\pi(x, x)$ associated with a reversible pair (\mathbf{P}, π) is defined by

$$\mathcal{E}_\pi(x, x) = \langle (I - \mathbf{P})x, x \rangle_\pi. \quad (2.15)$$

The reader will keep in mind that $\mathcal{E}_\pi(x, x)$ also depends on \mathbf{P} .

Theorem 2.2.

We have

$$\mathcal{E}_\pi(x, x) = \frac{1}{2} \sum_{i,j \in E} \pi(i) p_{ij} (x(j) - x(i))^2. \quad (2.16)$$

Proof.

$$\begin{aligned}
\langle (I - \mathbf{P})x, x \rangle_\pi &= \sum_{i, j \in E} \pi(i) p_{ij} x(i)(x(i) - x(j)) \\
&= \sum_{i, j \in E} \pi(j) p_{ji} x(j)(x(j) - x(i)) \\
&= \sum_{i, j \in E} \pi(i) p_{ij} x(j)(x(j) - x(i)),
\end{aligned}$$

where the second equality is obtained by a change of indexation, and the third uses reversibility of (\mathbf{P}, π) . Expressing $\mathcal{E}_\pi(x, x)$ as the half sum of the second and last terms in the above chain of equalities yields (2.16). \square

Observe from (2.16) that

$$\mathcal{E}_\pi(x, x) = \mathcal{E}_\pi(x - c\mathbf{1}, x - c\mathbf{1}) \quad (2.17)$$

for any real number c .

We can now state the characterization of the second-largest eigenvalue λ_2 , or equivalently of $\beta_2 = 1 - \lambda_2$.

Theorem 2.3. *Rayleigh's Theorem*

Let \mathbf{P} be an irreducible transition matrix on a finite state space, and let π be its stationary distribution. If (\mathbf{P}, π) is reversible, then for $j \geq 2$,

$$\beta_j = \inf \left\{ \frac{\mathcal{E}_\pi(x, x)}{\text{Var}_\pi(x)}; \langle x, v_i \rangle_\pi = 0 \text{ for } i \in [1, j-1], \quad x \neq 0 \right\}. \quad (2.18)$$

Any vector x achieving the infimum in (2.18) is an eigenvector of \mathbf{P} corresponding to the eigenvalue $\lambda_j = 1 - \beta_j$.

In particular,

$$\beta_2 = \inf \left\{ \frac{\mathcal{E}_\pi(x, x)}{\text{Var}_\pi(x)}; \langle x \rangle_\pi = 0; x \neq 0 \right\},$$

which is, in view of (2.17), equivalent to

$$\beta_2 = \inf \left\{ \frac{\mathcal{E}_\pi(x, x)}{\text{Var}_\pi(x)}; x \text{ nonconstant} \right\}. \quad (2.19)$$

Proof. From (2.11),

$$(I - \mathbf{P})x = \sum_{j=1}^r \beta_j \langle x, v_j \rangle_\pi v_j,$$

and therefore

$$\mathcal{E}_\pi(x, x) = \sum_{j=1}^r \beta_j |\langle x, v_j \rangle_\pi|^2.$$

Also from (2.9),

$$\langle x, x \rangle_\pi = \sum_{j=1}^r |\langle x, v_j \rangle_\pi|^2.$$

Therefore, if

$$\langle x, v_1 \rangle_\pi = 0, \dots, \langle x, v_{j-1} \rangle_\pi = 0$$

(in particular, $\langle x \rangle_\pi = 0$ and $\text{Var}_\pi(x) = \|x\|_\pi^2$), then

$$\frac{\mathcal{E}_\pi(x, x)}{\text{Var}_\pi(x)} = \frac{\sum_{\ell=j}^r \beta_\ell |\langle x, v_\ell \rangle_\pi|^2}{\sum_{\ell=j}^r |\langle x, v_\ell \rangle_\pi|^2}.$$

Writing $\alpha_\ell = |\langle x, v_\ell \rangle_\pi|^2$, and supposing that $\beta_j = \dots = \beta_{j+k} \neq \beta_{j+k+1}$, we have

$$\frac{\mathcal{E}_\pi(x, x)}{\text{Var}_\pi(x)} = \beta_j + \frac{\alpha_{j+k+1}(\beta_{j+k+1} - \beta_j) + \dots + \alpha_r(\beta_r - \beta_j)}{\alpha_j + \dots + \alpha_r} \geq \beta_j,$$

since $\beta_\ell - \beta_j \geq 0$ for $\ell > j$. From this, (2.18) follows by the choice $x = v_j$. Also, if x realizes the infimum, then for all $m \geq 1$, $\langle x, v_{j+k+m} \rangle_\pi = 0$. Taking into account $\langle x, v_1 \rangle_\pi = 0, \dots, \langle x, v_{j-1} \rangle_\pi = 0$, we see that x is an eigenvector corresponding to β_j , since it is orthogonal to all the eigensubspaces not corresponding to β_j . \square

Theorem 2.3 can be used to obtain an *upper bound* of the SLE.

Corollary 2.1.

If $A > 0$ is such that for all $x \in \mathbb{R}^r$,

$$\text{Var}_\pi(x) \leq A \mathcal{E}_\pi(x, x), \tag{2.20}$$

then, denoting by λ_2 the SLE of \mathbf{P} ,

$$\lambda_2 \leq 1 - \frac{1}{A}. \tag{2.21}$$

Proof. It follows from (2.19) that $\beta_2 \geq 1/A$. \square

We shall also need a *lower bound* for the smallest eigenvalue λ_r . If one can exhibit $B > 0$ such that for all $x \in \mathbb{R}^r$,

$$\langle \mathbf{P}x, x \rangle_\pi + \|x\|_\pi^2 \geq B \|x\|_\pi^2, \tag{2.22}$$

then we will have, with $x = v_j$ in (2.22),

$$\lambda_j \geq -1 + B. \tag{2.23}$$

The reason why an upper bound of λ_2 and a lower bound of λ_r are needed is that the speed of convergence to equilibrium in the aperiodic case is determined by the second largest eigenvalue modulus (SLEM)

$$\rho = \sup(\lambda_2, |\lambda_r|). \tag{2.24}$$

We already know this from the Perron–Frobenius theorem. For the reversible case, a direct proof is easy. Indeed, from (2.11), we have

$$\mathbf{P}^n x - \langle x \rangle_\pi \mathbf{1} = \sum_{j=2}^r \lambda_j^n \langle x, v_j \rangle_\pi v_j. \tag{2.25}$$

Writing the above equality for $x = \delta_k$, we obtain

$$p_{ik}(n) - \pi(k) = \sum_{j=2}^r \lambda_j^n v_j(i) v_j(k) \pi(k) \tag{2.26}$$

for all i , and therefore, denoting $\sum_{i=1}^r |x(i)|$ by $|x(\cdot)|$,

$$|p_i(\cdot) - \pi(\cdot)| \leq \sum_{j=2}^r |\lambda_j|^n |v_j(i)| \left(\sum_{k=1}^r \pi(k) |v_j(k)| \right).$$

Using the Cauchy–Schwarz inequality and $\|v_j\|_\pi^2 = 1$,

$$\sum_{k=1}^r \pi(k) |v_j(k)| \leq \left(\sum_{k=1}^r \pi(k) \right)^{\frac{1}{2}} \left(\sum_{k=1}^r \pi(k) |v_j(k)|^2 \right)^{\frac{1}{2}} = 1,$$

and therefore, by definition of ρ ,

$$|p_i(\cdot) - \pi(\cdot)| \leq r \left(\sup_{2 \leq j \leq r} |v_j(i)| \right) \rho^n.$$

3 Convergence Bounds Without Eigenvectors

3.1 Basic Bounds, Reversible Case

One disadvantage of the latter bound is the difficulty of obtaining the eigenvectors on the right-hand side. These eigenvectors do not appear in the bounds below.

We first recall a definition that was introduced in Chapter 4.

The distance in variation $d_V(\alpha, \beta)$ between two probability distributions α and β on the countable set E is defined by

$$d_V(\alpha, \beta) = \frac{1}{2} |\alpha - \beta| = \frac{1}{2} \sum_{i \in E} |\alpha(i) - \beta(i)|. \tag{3.1}$$

(Notational hazards will sometimes force us to use nonsymmetric notation such as $d_V(\alpha, \beta^t)$.)

For future reference, we recall the alternative definition of the distance in variation (Problem 4.1.1):

$$d_V(\alpha, \beta) = \sup \left(\sum_{i=1}^r \alpha(i)y(i) - \sum_{i=1}^r \beta(i)y(i); \sup |y(i)| = 1 \right). \quad (3.2)$$

The strange names of the next two results are not standard. They refer to the proofs, which use right eigenvectors of the transition matrix in the $\ell^2(\pi)$ setting for the first one, and left eigenvectors in the $\ell^2(\frac{1}{\pi})$ setting for the second one.

Theorem 3.1. *The π -Bound*

Let \mathbf{P} be a reversible irreducible transition matrix on the finite state space $E = \{1, \dots, r\}$, with the stationary distribution π . Then for all $n \geq 1$ and all $i \in E$,

$$d_V(\delta_i^T \mathbf{P}^n, \pi)^2 \leq \frac{p_{ii}(2)}{\pi(i)} \rho^{2n-2}, \quad (3.3)$$

where ρ is the SLEM of \mathbf{P} .

Proof. From (2.25) and $\text{Var}_\pi(x) = \sum_{j=2}^r |\langle x, v_j \rangle_\pi|^2$, we have

$$\|\mathbf{P}^n x - \langle x \rangle_\pi \mathbf{1}\|_\pi^2 = \sum_{j=2}^r |\lambda_j|^{2n} |\langle x, v_j \rangle_\pi|^2 \leq \rho^{2n} \text{Var}_\pi(x). \quad (3.4)$$

Now, using reversibility and the Cauchy–Schwarz inequality,

$$\begin{aligned} \left| \sum_{j \in E} p_{ij} x(j) \right|^2 &= \left| \sum_{j \in E} p_{ji} \frac{\pi(j)}{\pi(i)} x(j) \right|^2 \leq \left(\sum_{j \in E} \frac{p_{ji}}{\pi(i)} |x(j)| \pi(j) \right)^2 \\ &\leq \left(\sum_{j \in E} x(j)^2 \pi(j) \right) \left(\sum_{j \in E} \left(\frac{p_{ji}}{\pi(i)} \right)^2 \pi(j) \right) \\ &= \left(\sum_{j \in E} x(j)^2 \pi(j) \right) \left(\sum_{j \in E} (p_{ji} p_{ij}) \frac{1}{\pi(i)} \right) = \left(\sum_{j \in E} x(j)^2 \pi(j) \right) \frac{p_{ii}(2)}{\pi(i)}, \end{aligned}$$

that is,

$$\left| \sum_{j \in E} p_{ij} x(j) \right|^2 \leq \frac{p_{ii}(2)}{\pi(i)} \|x\|_\pi^2.$$

With $x = \mathbf{P}^{n-1} y - \langle y \rangle_\pi \mathbf{1}$, this gives, in view of (3.4),

$$\begin{aligned} \left| \sum_{j=1}^r p_{ij}(n) y(j) - \sum_{i=1}^r \pi(j) y(j) \right|^2 &\leq \frac{p_{ii}(2)}{\pi(i)} \|\mathbf{P}^{n-1} y - \langle y \rangle_\pi \mathbf{1}\|_\pi^2 \\ &\leq \frac{p_{ii}(2)}{\pi(i)} \text{Var}_\pi(y) \rho^{2n-2}. \end{aligned}$$

The result then follows from (3.2) by observing that if y is such that $\sup |y(i)| \leq 1$, then $\text{Var}_\pi(y) \leq 1$. \square

For the second basic bound, we need a definition.

Definition 3.1. χ^2 -Contrast

The χ^2 -contrast $\chi^2(\alpha; \beta)$ of α with respect to β is defined by

$$\chi^2(\alpha; \beta) = \sum_{i \in E} \frac{(\alpha(i) - \beta(i))^2}{\beta(i)}. \tag{3.5}$$

Note that

$$\chi^2(\alpha; \pi) = \|\alpha - \pi\|_{\frac{1}{\pi}}^2. \tag{3.6}$$

Theorem 3.2.

We have

$$4d_V(\alpha, \beta)^2 \leq \chi^2(\alpha; \beta). \tag{3.7}$$

Proof. The result follows from the Cauchy–Schwarz inequality:

$$\begin{aligned} \left(\sum_{i \in E} |\alpha(i) - \beta(i)| \right)^2 &= \left(\sum_{i \in E} \left| \frac{\alpha(i)}{\beta(i)} - 1 \right| \beta(i)^{\frac{1}{2}} \beta(i)^{\frac{1}{2}} \right)^2 \\ &\leq \sum_{i \in E} \left(\frac{\alpha(i)}{\beta(i)} - 1 \right)^2 \beta(i) \\ &= \sum_{i \in E} \frac{1}{\beta(i)} (\alpha(i) - \beta(i))^2. \end{aligned} \quad \square$$

Theorem 3.3. *The $\frac{1}{\pi}$ -Bound*

Let \mathbf{P} be a reversible irreducible transition matrix on the finite state space $E = \{1, \dots, r\}$, with the stationary distribution π . Then for any probability distribution μ on E , and for all $n \geq 1$,

$$\|\mu^T \mathbf{P}^n - \pi^T\|_{\frac{1}{\pi}} \leq \rho^n \|\mu - \pi\|_{\frac{1}{\pi}}. \tag{3.8}$$

Also, for $n \geq 1$, all $i \in E$, and all $A \subset E$,

$$|\delta_i^T \mathbf{P}^n(A) - \pi^T(A)| \leq \left(\frac{1 - \pi(i)}{\pi(i)} \right)^{\frac{1}{2}} \min \left(\pi(A)^{\frac{1}{2}}, \frac{1}{2} \right) \rho^n, \tag{3.9}$$

where ρ is the SLEM of \mathbf{P} . In particular,

$$4d_V(\delta_i^T \mathbf{P}^n, \pi)^2 \leq \frac{1 - \pi(i)}{\pi(i)} \rho^{2n}. \tag{3.10}$$

Proof. Recall that $u_1 = \pi$, and therefore $\langle \mu - \pi, u_1 \rangle_{\frac{1}{\pi}} = \sum_{i \in E} (\mu(i) - \pi(i)) = 0$. Therefore, by (2.12), and denoting by α_j the quantity $\langle \mu - \pi, u_j \rangle_{\frac{1}{\pi}}$, we obtain

$$\begin{aligned} \|(\mu - \pi)^T \mathbf{P}^n\|_{\frac{1}{\pi}}^2 &= \sum_{j=2}^r \alpha_j^2 \lambda_j^{2n} \|u_j\|_{\frac{1}{\pi}}^2 = \sum_{j=2}^r \alpha_j^2 \lambda_j^{2n} \\ &\leq \rho^{2n} \sum_{j=2}^r \alpha_j^2 = \rho^{2n} \|\mu - \pi\|_{\frac{1}{\pi}}^2, \end{aligned}$$

and (3.8) follows, since $\pi^T \mathbf{P}^n = \pi^T$.

Write $\delta_i^T \mathbf{P}^n = \mu_n^T$. Then, using the Cauchy–Schwarz inequality,

$$\begin{aligned} |\mu_n(A) - \pi(A)|^2 &= \left| \sum_{i \in A} \left(\frac{\mu_n(i)}{\pi(i)} - 1 \right) \pi(i) \right|^2 \\ &\leq \left(\sum_{i \in A} \left(\frac{\mu_n(i)}{\pi(i)} - 1 \right)^2 \pi(i) \right) \pi(A) \leq \left(\sum_{i \in E} \left(\frac{\mu_n(i)}{\pi(i)} - 1 \right)^2 \pi(i) \right) \pi(A) \\ &= \|\delta_i^T \mathbf{P}^n - \pi^T\|_{\frac{1}{\pi}}^2 \pi(A) \leq \rho^{2n} \|\delta_i - \pi\|_{\frac{1}{\pi}}^2 \pi(A), \end{aligned}$$

where the last inequality uses (3.8). But, as simple calculations reveal,

$$\|\delta_i - \pi\|_{\frac{1}{\pi}}^2 = \frac{1 - \pi(i)}{\pi(i)}, \quad (3.11)$$

and therefore we have obtained

$$|\delta_i^T \mathbf{P}^n(A) - \pi^T(A)| \leq \left(\frac{1 - \pi(i)}{\pi(i)} \right)^{\frac{1}{2}} \pi(A)^{\frac{1}{2}} \rho^n. \quad (3.12)$$

Now,

$$|\mu_n(A) - \pi(A)|^2 \leq d_V(\mu_n, \pi)^2 \leq \frac{1}{4} \chi^2(\mu_n; \pi).$$

But, by (3.8), (3.6), and (3.11)

$$\chi^2(\mu_n; \pi) = \|\delta_i^T \mathbf{P}^n - \pi^T\|_{\frac{1}{\pi}}^2 \leq \rho^{2n} \|\delta_i - \pi\|_{\frac{1}{\pi}}^2 = \rho^{2n} \frac{1 - \pi(i)}{\pi(i)}.$$

Therefore,

$$|\delta_i^T \mathbf{P}^n(A) - \pi^T(A)| \leq \left(\frac{1 - \pi(i)}{\pi(i)} \right)^{\frac{1}{2}} \frac{1}{2} \rho^n. \quad (3.13)$$

Combining (3.12) and (3.13) gives (3.9). Inequality (3.10) then follows since $d_V(\alpha, \beta) = \sup_{A \subseteq E} |\alpha(A) - \beta(A)|$ (see Lemma 1.1 of Chapter 4). \square

3.2 Nonreversible Case

In the nonreversible case, we can still obtain similar bounds, but they feature the SLEM (actually the SLE) of a reversible matrix intimately connected to the original matrix.

Suppose \mathbf{P} is an ergodic transition matrix on the finite state space $E = \{1, 2, \dots, r\}$, with stationary distribution π . This time do not assume that (\mathbf{P}, π) is reversible. Consider the transition matrix $\tilde{\mathbf{P}}$ of the time-reversed chain, defined by

$$\tilde{p}_{ij} = \frac{\pi(j)p_{ji}}{\pi(i)},$$

or, in compact form,

$$\tilde{\mathbf{P}} = D^{-1}\mathbf{P}^T D. \tag{3.14}$$

From (3.14) we see that if \mathbf{P} and \mathbf{R} are two ergodic transition matrices with the same stationary distribution π , then

$$\tilde{\mathbf{P}}\mathbf{R} = \tilde{\mathbf{R}}\tilde{\mathbf{P}}.$$

In particular, the matrix

$$M = M(\mathbf{P}) = \mathbf{P}\tilde{\mathbf{P}} \tag{3.15}$$

is reversible with respect to its stationary distribution π . Its eigenvalues are real, and all belong to $[-1, +1]$. In fact, they all belong to the interval $[0, 1]$. To see this, observe that M has the same eigenvalues as $D^{\frac{1}{2}}MD^{-\frac{1}{2}}$, and that the latter matrix is $(D^{\frac{1}{2}}\mathbf{P}D^{-\frac{1}{2}})(D^{\frac{1}{2}}\mathbf{P}D^{-\frac{1}{2}})^T$, a symmetric definite nonnegative matrix.

The matrix M given by (3.15) is the *multiplicative reversibilization* of \mathbf{P} . The *additive reversibilization* of \mathbf{P} is, by definition, the reversible matrix

$$A = A(\mathbf{P}) = \frac{1}{2}(\mathbf{P} + \tilde{\mathbf{P}}). \tag{3.16}$$

The following bound is due to Fill (1991).

Theorem 3.4. Contrast Bound

Let $\gamma_1 = \gamma_1(M)$ be the second-largest eigenvalue of $M = \mathbf{P}\tilde{\mathbf{P}}$, where \mathbf{P} is an ergodic transition matrix on the finite state space E . Then for any probability distribution ν on E ,

$$|\nu^T \mathbf{P}^n - \pi^T|^2 \leq \gamma_1(M)^n \chi^2(\nu; \pi). \tag{3.17}$$

Proof. We shall use the following identity, due to Mihaïl (1989):

$$\text{Var}_\pi(x) = \text{Var}_\pi(\tilde{\mathbf{P}}x) + \langle (I - M)x, x \rangle_\pi, \tag{3.18}$$

which is proven as follows. First, from (2.17), if we let $\hat{x} = x - \langle x \rangle_\pi \mathbf{1}$, then

$$\begin{aligned} \langle (I - M)x, x \rangle_\pi &= \langle (I - M)\hat{x}, \hat{x} \rangle_\pi = \|\hat{x}\|_\pi^2 - \langle M\hat{x}, \hat{x} \rangle_\pi \\ &= \|\hat{x}\|_\pi^2 - \langle \mathbf{P}\tilde{\mathbf{P}}\hat{x}, \hat{x} \rangle_\pi = \|\hat{x}\|_\pi^2 - \|\tilde{\mathbf{P}}\hat{x}\|_\pi^2, \end{aligned}$$

where we use the fact that $\tilde{\mathbf{P}}$ is the adjoint of \mathbf{P} in $\ell^2(\pi)$. The identity (3.18) follows, since $\|\hat{x}\|_\pi^2 = \text{Var}_\pi(x)$ and $\|\tilde{\mathbf{P}}\hat{x}\|_\pi^2 = \text{Var}_\pi(\tilde{\mathbf{P}}x)$.

Now set $\chi_n^2 = \chi^2(v^T \mathbf{P}^n; \pi)$, and let $\rho_n(i) = \frac{(v^T \mathbf{P}^n)(i)}{\pi(i)}$. One verifies by inspection that

$$\text{Var}_\pi(\rho_n) = \chi_n^2$$

and

$$\tilde{\mathbf{P}}\rho_n = \rho_{n+1}.$$

Therefore, from Mihail's identity,

$$\chi_n^2 = \chi_{n+1}^2 + \langle (1 - M)\rho_n, \rho_n \rangle_\pi.$$

By the spectral theorem,

$$\langle (1 - M)\rho_n, \rho_n \rangle_\pi \geq (1 - \gamma_1(M)) \text{Var}_\pi(\rho_n) = (1 - \gamma_1(M))\chi_n^2,$$

and therefore $\chi_{n-1}^2 \leq \gamma_1 \chi_n^2$, from which it follows that $\chi_n^2 \leq \gamma_1^n \chi_0^2$. But by Theorem 3.2, $d_V(v, \pi)^2 \leq \chi^2(v; \pi)$, and this finishes the proof. \square

Fill's inequality is interesting in several respects. It gives a bound that is proportional to the χ^2 -contrast between the initial distribution and the *target* distribution π . Secondly, this bound is in terms of the SLEM of a reversible matrix, for which some powerful bounds will be obtained in the next section. Also, the SLEM $\gamma_1(M)$ is, in fact, the SLE, and good bounds for the SLE are in general easier than good bounds for the SLEM.

4 Geometric Bounds

4.1 Weighted Paths

In order to use the results of the previous section, it remains to obtain good upper bounds of the SLEM. For this one has to take a closer look at the transition matrix. The next two theorems give an upper bound and a lower bound in terms of the geometry of the transition graph. They concern *reversible* transition matrices. The two theorems and the two examples below are due to Diaconis and Stroock (1991).

In the transition graph associated with \mathbf{P} , we shall denote an oriented edge $i \rightarrow j$ by e , and call $e^- = i$ and $e^+ = j$ its initial vertex and end vertex, respectively. Define for any such oriented edge e ,

$$Q(e) = \pi(i)p_{ij}. \tag{4.1}$$

For each ordered pair of *distinct* states (i, j) , select arbitrarily one and only one path from i to j , that is a sequence i, i_1, \dots, i_m, j such that $p_{ii_1} p_{i_1 i_2} \cdots p_{i_m j} > 0$, which does not use the same edge twice. Let Γ be the collection of paths so selected. For a path $\gamma_{ij} \in \Gamma$, let

$$|\gamma_{ij}|_Q \stackrel{\text{def}}{=} \sum_{e \in \gamma_{ij}} \frac{1}{Q(e)} = \frac{1}{\pi(i)p_{ii_1}} + \frac{1}{\pi(i_1)p_{i_1 i_2}} + \cdots + \frac{1}{\pi(i_m)p_{i_m j}}. \tag{4.2}$$

Define the *Poincaré coefficient*

$$\kappa = \kappa(\Gamma) = \max_e \sum_{\gamma_{ij} \ni e} |\gamma_{ij}|_Q \pi(i)\pi(j). \tag{4.3}$$

Theorem 4.1. *Weighted Path Upper Bound*

Let \mathbf{P} be an irreducible transition matrix on the finite state space E , with stationary distribution π , and assume (\mathbf{P}, π) to be reversible. Denoting by λ_2 its SLE,

$$\lambda_2 \leq 1 - \frac{1}{\kappa}. \tag{4.4}$$

Proof. It suffices to show that (2.20) holds for $A = \kappa$. For this, write

$$\begin{aligned} \text{Var}_\pi(x) &= \frac{1}{2} \sum_{i,j \in E} (x(i) - x(j))^2 \pi(i)\pi(j) \\ &= \frac{1}{2} \sum_{i,j \in E} \left\{ \sum_{e \in \gamma_{ij}} \frac{1}{Q(e)^{\frac{1}{2}}} Q(e)^{\frac{1}{2}} (x(e^-) - x(e^+)) \right\}^2 \pi(i)\pi(j). \end{aligned}$$

This quantity is, by the Cauchy–Schwarz inequality, bounded above by

$$\begin{aligned} &\frac{1}{2} \sum_{i,j \in E} \left\{ |\gamma_{ij}|_Q \sum_{e \in \gamma_{ij}} Q(e) (x(e^-) - x(e^+))^2 \right\} \pi(i)\pi(j) \\ &= \frac{1}{2} \sum_e \left\{ Q(e) (x(e^-) - x(e^+))^2 \left[\sum_{\gamma_{ij} \ni e} |\gamma_{ij}|_Q \pi(i)\pi(j) \right] \right\} \leq \mathcal{E}_\pi(x, x) \kappa(\Gamma). \quad \square \end{aligned}$$

Similar bounds in the same vein can be obtained, with a different coefficient (see Problem 6.4.1, or Example 7.7.6).

We now turn to the lower bound. For each state i , select exactly one closed path σ_i from i to i that does not pass twice through the same edge, and with an *odd* number of edges (for this to be possible, we assume that \mathbf{P} is aperiodic), and let Σ be the collection of paths so selected. For a path $\sigma_i \in \Sigma$, let

$$|\sigma_i|_Q = \sum_{e \in \sigma_i} \frac{1}{Q(e)}. \tag{4.5}$$

Define

$$\alpha = \alpha(\Sigma) = \max_e \sum_{\sigma_i \ni e} |\sigma_i|_Q \pi(i). \tag{4.6}$$

Theorem 4.2. *Weighted Path Lower Bound*

Let \mathbf{P} be irreducible and *aperiodic* transition matrix on the finite state space E , with stationary distribution π , and assume (\mathbf{P}, π) to be reversible. Then

$$\lambda_r \geq -1 + \frac{2}{\alpha}. \tag{4.7}$$

Proof. It suffices to prove (2.22) with $B = \frac{2}{\alpha}$. For this, we use the easily established identity

$$\frac{1}{2} \sum_{i,j \in E} (x(i) + x(j))^2 \pi(i) p_{ij} = \langle \mathbf{P}x, x \rangle_\pi + \|x\|_\pi^2. \tag{4.8}$$

If σ_i is a path from i to i with an odd number of edges, of the form $\sigma_i = (i_0 = i, i_1, i_2, \dots, i_{2m}, i)$, then

$$\begin{aligned} x(i) &= \frac{1}{2} \{ (x(i_0) + x(i_1)) - (x(i_1) + x(i_2)) + \dots + (x(i_{2m}) + x(i)) \} \\ &= \frac{1}{2} \sum_{e \in \sigma_i} (-1)^{n(e)} (x(e^+) + x(e^-)), \end{aligned}$$

where $n(e) = k$ if $e = (i_k, i_{k+1}) \in \sigma_i$. Therefore,

$$\|x\|_\pi^2 = \sum_{i \in E} \frac{\pi(i)}{4} \left\{ \sum_{e \in \sigma_i} \frac{1}{Q(e)^{\frac{1}{2}}} Q(e)^{\frac{1}{2}} (-1)^{n(e)} (x(e^+) + x(e^-)) \right\}^2,$$

and by the Cauchy–Schwarz inequality, this quantity is lesser than or equal to

$$\begin{aligned} &\sum_{i \in E} \left\{ \frac{\pi(i)}{4} |\sigma_i|_Q \sum_{e \in \sigma_i} (x(e^+) + x(e^-))^2 Q(e) \right\} \\ &= \frac{1}{4} \sum_e \left\{ (x(e^+) + x(e^-))^2 Q(e) \sum_{\sigma_i \ni e} |\sigma_i|_Q \pi(i) \right\} \\ &\leq \frac{\alpha}{4} \sum_e (x(e^-) + x(e^+))^2 Q(e). \end{aligned}$$

Therefore, in view of (4.8),

$$\|x\|_\pi^2 \leq \frac{\alpha}{2} \{ \|x\|_\pi^2 + \langle \mathbf{P}x, x \rangle_\pi \},$$

and this is the inequality sought. □

Example 4.1. *The Random Walk on a Graph*

For the *random walk on a graph* of Example 6.3 of Chapter 2, recall that the stationary distribution is $\pi(i) = \frac{d_i}{2|G|}$, where d_i is the degree of vertex i and $|G|$ is the number of edges. Hence $Q(e)^{-1} = 2|G|$ and

$$\kappa(\Gamma) = \max_e \frac{1}{2|G|} \sum_{\gamma_{ij} \ni e} |\gamma_{ij}| d_i d_j,$$

where $|\gamma_{ij}|$ is the length of path γ_{ij} . Let $d = \max d_i$ and $|\gamma| = \max |\gamma_{ij}|$. We have

$$\kappa(\Gamma) \leq \frac{1}{2|G|} |\gamma| d^2 B,$$

where B is the *bottleneck coefficient*

$$B = \max_e |\{\gamma \in \Gamma; e \in \gamma\}|.$$

Therefore,

$$\lambda_2 \leq 1 - \frac{2|G|}{d^2|\gamma|B}. \quad (4.9)$$

Similar calculations give

$$\lambda_r \geq -1 + \frac{2}{d|\sigma|b}, \quad (4.10)$$

where $|\sigma| = \max |\sigma_i|$, and

$$b = \max_e |\{\sigma \in \Sigma; e \in \sigma\}|. \quad \diamond$$

4.2 Conductance

This type of bound concerns, as before, finite state space irreducible reversible transition matrices \mathbf{P} .

For a nonempty set $B \subset E = \{1, \dots, r\}$, the state space, one defines the *capacity* of B ,

$$\pi(B) \stackrel{\text{def}}{=} \sum_{i \in B} \pi(i), \quad (4.11)$$

and the *ergodic flow* out of B ,

$$F(B) \stackrel{\text{def}}{=} \sum_{i \in B, j \in \bar{B}} \pi(i) p_{ij}. \quad (4.12)$$

Note that $0 \leq F(B) \leq \pi(B) \leq 1$. Define for B not empty,

$$\psi(B, \mathbf{P}) \stackrel{\text{def}}{=} \frac{F(B)}{\pi(B)}. \quad (4.13)$$

The *conductance* of the pair (\mathbf{P}, π) is

$$\varphi(\mathbf{P}) \stackrel{\text{def}}{=} \inf \left(\psi(B, \mathbf{P}); \quad 0 < |B| < r, \pi(B) \leq \frac{1}{2} \right). \quad (4.14)$$

The following result is due to Jerrum and Sinclair (1989).

Theorem 4.3. *Conductance Bounds*

$$1 - 2\varphi(\mathbf{P}) \leq \lambda_2 \leq 1 - \frac{\varphi(\mathbf{P})^2}{2}. \quad (4.15)$$

Proof. (a) We start with the rightmost inequality. Let u be a left eigenvector of \mathbf{P} associated with an eigenvalue $\lambda \neq 1$. In particular, u is orthogonal to π , the left eigenvector associated with λ_1 , and therefore u has positive as well as negative entries. The same is true for x defined by

$$x(i) = \frac{u(i)}{\pi(i)}.$$

Assume without loss of generality that for some $k \in [1, r]$,

$$x(i) \geq \cdots \geq x(k) > 0 \geq x(k+1) \geq \cdots \geq x(r),$$

and that, writing $S = \{1, \dots, k\}$, we have $\pi(S) \leq \frac{1}{2}$ (if necessary, change the order of the states, and for the last assumption, change u into $-u$). Define

$$y(i) = \begin{cases} \frac{u(i)}{\pi(i)} & \text{if } u(i) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

We have

$$u^T(I - \mathbf{P}) = u^T(1 - \lambda),$$

and therefore

$$u^T(I - \mathbf{P})y = (1 - \lambda)u^T y = (1 - \lambda) \sum_{i \in S} \pi(i) y(i)^2. \quad (4.16)$$

Also,

$$\begin{aligned} u^T(I - \mathbf{P})y &= \sum_{i \in S} \sum_{j=1}^r (\delta_{ji} - p_{ji}) u(j) y(i) \\ &\geq \sum_{i \in S} \sum_{j \in S} (\delta_{ji} - p_{ji}) u(j) y(i), \end{aligned}$$

since the missing terms $-p_{ji}u(j)y(i)$ corresponding to $i \in S$ and $j \notin S$ are negative or null. Therefore,

$$u^T(I - \mathbf{P})y \geq \langle y, (I - \mathbf{P})y \rangle_{\pi},$$

and using (4.16) and Theorem 2.1,

$$1 - \lambda \geq \frac{\sum_{i < j} \pi(i) p_{ij} (y(i) - y(j))^2}{\sum_{i \in S} \pi(i) y(i)^2}.$$

From $(a + b)^2 \leq 2(a^2 + b^2)$, we obtain

$$\sum_{i < j} \pi(i) p_{ij} (y(i) + y(j))^2 \leq 2 \sum_{i < j} \pi(i) p_{ij} (y(i)^2 + y(j)^2),$$

and, by reversibility,

$$\begin{aligned} \sum_{i < j} \pi(i) p_{ij} (y(i)^2 + y(j)^2) &= \sum_{i < j} \pi(i) p_{ij} y(i)^2 + \sum_{i < j} \pi(j) p_{ji} y(j)^2 \\ &= \sum_{i \neq j} \pi(i) p_{ij} y(i)^2 \\ &\leq \sum_{i \in S} \pi(i) y(i)^2. \end{aligned}$$

Thus

$$1 - \lambda \geq \frac{\sum_{i < j} \pi(i) p_{ij} (y(i) - y(j))^2}{\sum_{i \in S} \pi(i) y(i)^2} \frac{\sum_{i < j} \pi(i) p_{ij} (y(i) + y(j))^2}{2 \sum_{i \in S} \pi(i) y(i)^2}.$$

By the Cauchy–Schwarz inequality (and using the identity $a^2 - b^2 = (a - b)(a + b)$),

$$\begin{aligned} &\left(\sum_{i < j} \pi(i) p_{ij} (y(i)^2 - y(j)^2) \right)^2 \\ &\leq \left(\sum_{i < j} \pi(i) p_{ij} (y(i) - y(j))^2 \right) \left(\sum_{i < j} \pi(i) p_{ij} (y(i) + y(j))^2 \right), \end{aligned}$$

and therefore

$$1 - \lambda \geq \frac{1}{2} \left(\frac{\sum_{i < j} \pi(i) p_{ij} (y(i)^2 - y(j)^2)}{\sum_{i \in S} \pi(i) y(i)^2} \right)^2.$$

Define $S_l = \{1, \dots, l\}$. We have

$$\begin{aligned} \sum_{i < j} \pi(i) p_{ij} (y(i)^2 - y(j)^2) &= \sum_{i < j} \pi(i) p_{ij} \left(\sum_{i \leq l < j} (y(l)^2 - y(l+1)^2) \right) \\ &= \sum_{l=1}^k (y(l)^2 - y(l+1)^2) \sum_{i \in S_l, j \notin S_l} \pi(i) p_{ij} \\ &= \sum_{l=1}^k (y(l)^2 - y(l+1)^2) F(S_l). \end{aligned}$$

Since for $l \in [1, k]$, $\pi(S_l) \leq \pi(S) \leq \frac{1}{2}$, we have

$$F(S_l) \geq \varphi(\mathbf{P}) \pi(S_l).$$

Therefore,

$$\begin{aligned}
 \sum_{i < j} \pi(i) p_{ij} (y(i)^2 - y(j)^2) &\geq \varphi(\mathbf{P}) \sum_{l=1}^k (y(l)^2 - y(l+1)^2) \pi(S_l) \\
 &= \varphi(\mathbf{P}) \sum_{l=1}^k (y(l)^2 - y(l+1)^2) \sum_{i=1}^l \pi(i) \\
 &= \varphi(\mathbf{P}) \sum_{i=1}^k \pi(i) \left(\sum_{l=1}^k (y(l)^2 - y(l+1)^2) \right) \\
 &= \varphi(\mathbf{P}) \sum_{i \in S} \pi(i) y(i)^2.
 \end{aligned}$$

Therefore,

$$1 - \lambda \geq \frac{\varphi(\mathbf{P})^2}{2}.$$

(b) Proof of the leftmost inequality in (4.15). By the spectral theorem,

$$1 - \lambda_2 \leq \frac{\mathcal{E}_\pi(x, x)}{\|x\|_\pi^2}$$

for any vector nontrivial vector x such that $\langle x \rangle_\pi = 0$. Select $B \subset E$ such that $\pi(B) \leq \frac{1}{2}$, and define

$$x(i) = \begin{cases} 1 - \pi(B) & \text{if } i \in B, \\ -\pi(B) & \text{if } i \notin B. \end{cases}$$

Then $\langle x \rangle_\pi = 0$ and $\|x\|_\pi^2 = \pi(B)(1 - \pi(B))$. Also,

$$\begin{aligned}
 \mathcal{E}_\pi(x, x) &= \frac{1}{2} \sum_{ij} (x(i) - x(j))^2 \pi(i) p_{ij} \\
 &= \frac{1}{2} \sum_{i \in B} \sum_{j \notin B} + \frac{1}{2} \sum_{i \notin B} \sum_{j \in B} \\
 &= \frac{1}{2} F(B) + \frac{1}{2} F(\overline{B}) \\
 &= F(B).
 \end{aligned}$$

Therefore,

$$1 - \lambda_2 \leq \frac{F(B)}{\pi(B)(1 - \pi(B))} \leq 2 \frac{F(B)}{\pi(B)}.$$

This being true for all B such that $\pi(B) \leq \frac{1}{2}$, we have, by definition of $\varphi(\mathbf{P})$,

$$1 - \lambda_2 \leq 2 \varphi(\mathbf{P}).$$

□

5 Probabilistic Bounds

5.1 Separation and Strong Stationary Times

In the present section we shall review the basic properties of a special class of coupling times introduced and studied by Aldous and Diaconis (1981) and Diaconis and Fill (1991), in view of obtaining sharp bounds for the rate of convergence to steady state. Another theoretical advantage of these stopping times is the fact that the chain observed at such times has the stationary distribution; see (Fill, 1998) in connection with perfect sampling.

Definition 5.1. *Separation*

Let α and β be two probability distributions on the denumerable space E . The *separation* of α from β , denoted by $s(\alpha; \beta)$, is defined by

$$s(\alpha; \beta) = \max_{i \in E} \left(1 - \frac{\alpha(i)}{\beta(i)} \right). \quad (5.1)$$

Note that

$$0 \leq s(\alpha; \beta) \leq 1 \quad (5.2)$$

(for the lower bound, observe that one cannot have $\alpha(i) > \beta(i)$ for all i).

Theorem 5.1. *Separation vs. Distance in Variation*

Let α and β be two probability distributions on the denumerable space E . Then

$$d_V(\alpha; \beta) \leq s(\alpha; \beta). \quad (5.3)$$

Proof. From the proof of Theorem 1.1 of Chapter 4, we extract the identity

$$d_V(\alpha; \beta) = \sum_{i; \beta(i) > \alpha(i)} (\beta(i) - \alpha(i)).$$

But the latter sum equals

$$\sum_{i; \beta(i) > \alpha(i)} \beta(i) \left(1 - \frac{\alpha(i)}{\beta(i)} \right) \leq \left(\sum_{i; \beta(i) > \alpha(i)} \beta(i) \right) s(\alpha; \beta) \leq s(\alpha; \beta). \quad \square$$

We need a slight extension of the notion of stopping time.

Definition 5.2. *Randomized Stopping Time*

Let $\{X_n\}_{n \geq 0}$ be some E -valued stochastic process, and let $\{Z_n\}_{n \geq 0}$ be another process, F -valued and independent of $\{X_n\}_{n \geq 0}$. A random time T with values in $\bar{\mathbb{N}} = \mathbb{N} \cup \{+\infty\}$ is called a *randomized stopping time* with respect to $\{X_n\}_{n \geq 0}$ using the *side information* $\{Z_n\}_{n \geq 0}$, if for all $k \in \mathbb{N}$, the event $\{T = k\}$ is expressible in terms of X_0, \dots, X_k and $\{Z_n\}_{n \geq 0}$.

In the definition, $\{Z_n\}_{n \geq 0}$ is an arbitrary given stochastic process. In the sequel, we shall generally omit the phrase *using the side information* $\{Z_n\}_{n \geq 0}$, and even let the context dictate the choice of $\{Z_n\}_{n \geq 0}$.

Suppose that in the above definition $\{X_n\}_{n \geq 0}$ is an HMC. Then for all $m, n \geq 0$ and for all $i, j \in E$,

$$P(X_{m+n} = j | X_n = i, T \leq n) = p_{ij}(m).$$

Indeed, $\{T \leq n\}$ is expressible in terms of X_0, \dots, X_n and $\{Z_n\}_{n \geq 0}$, and is therefore independent of X_{m+n} given $X_n = j$.

Similar formulas, formally identical to the case where T is a usual, nonrandomized, stopping time of $\{X_n\}_{n \geq 0}$, hold true and will be used in the calculations below. Their validity essentially depends on the independence of $\{X_n\}_{n \geq 0}$ and $\{Z_n\}_{n \geq 0}$.

We introduce the main concept of this section.

Definition 5.3. *Strong Stationary Time*

A randomized stopping time T with respect to the HMC $\{X_n\}_{n \geq 0}$ admitting a unique stationary distribution π is called a *strong stationary time* of this HMC iff it is almost surely finite and

(α) X_T has the distribution π and is independent of T .

In the above definition, condition (α) is equivalent to either one of the following two conditions:

(β) For all $i \in E$ and all $n \geq 0$,

$$P(X_n = i | T = n) = \pi(i).$$

(γ) For all $i \in E$ and all $n \geq 0$,

$$P(X_n = i | T \leq n) = \pi(i).$$

The reader is invited to provide the proof (Problem 6.5.2).

Also, if either (α), or (β), or (γ) holds, then $\{X_{T+n}\}_{n \geq 0}$ is a stationary HMC with the transition matrix \mathbf{P} and is independent of T . To check this, just write

$$\begin{aligned} & P(T = k, X_T = i_0, X_{T+1} = i_1, \dots, X_{T+n} = i_n) \\ &= P(T = k, X_k = i_0, X_{k+1} = i_1, \dots, X_{k+n} = i_n) \\ &= P(T = k, X_k = i_0)P(X_{k+1} = i_1, \dots, X_{k+n} = i_n | T = k, X_k = i_0) \\ &= P(T = k)\pi(i_0)P(X_{k+1} = i_1, \dots, X_{k+n} = i_n | X_k = i_0) \\ &= P(T = k)P_\pi(X_k = i_0, X_{k+1} = i_1, \dots, X_{k+n} = i_n) \\ &= P(T = k)P_\pi(X_0 = i_0, X_1 = i_1, \dots, X_n = i_n). \end{aligned}$$

The announced result then follows.

Theorem 5.2. *Minimal Strong Stationary Time*

Let \mathbf{P} be an irreducible transition matrix with the stationary distribution π . If T is a strong stationary time of the HMC $\{X_n\}_{n \geq 0}$ with transition matrix \mathbf{P} and initial distribution μ , then

$$s(\mu^T \mathbf{P}^n; \pi) \leq P(T > n). \quad (5.4)$$

Conversely, if $\{X_n\}_{n \geq 0}$ is an HMC with transition matrix \mathbf{P} , initial distribution μ , and stationary distribution π , there exists a strong stationary time T such that

$$P(T > n) = s(\mu^T \mathbf{P}^n; \pi), \quad (5.5)$$

and it is called a *minimal strong stationary time*.

Proof. Direct part:

$$P(X_n = i) \geq P(X_n = i, T \leq n) = (1 - P(T > n))\pi(i).$$

Therefore, for all i ,

$$P(T > n) \geq 1 - \frac{P(X_n = i)}{\pi(i)}.$$

Converse part: The strong stationary time T is obtained by the following construction. Let

$$a_n = \min_{i \in E} \frac{\mu_n(i)}{\pi(i)},$$

where $\mu_n^T = \mu^T \mathbf{P}^n$, and let k be the smallest integer such that $a_k > 0$. Define T such that $P(T < k) = 0$ and

$$P(T = k | X_k = i) = \frac{a_k \pi(i)}{P(X_k = i)},$$

and, by induction,

$$P(T = n | X_n = i, T > n - 1) = \frac{a_n - a_{n-1}}{\frac{\mu_n(i)}{\pi(i)} - a_{n-1}}.$$

Note that the right-hand side of the above identity is in $[0, 1]$, by definition of a_n .

Suppose that for all $m \geq k$ and all $i \in E$,

$$P(X_m = i, T = m) = \pi(i)(a_m - a_{m-1}). \quad (5.6)$$

Then T is a strong stationary time, since by summing (5.6) with respect to i , we have $P(T = m) = a_m - a_{m-1}$, and therefore

$$P(X_m = i, T = m) = \pi(i)P(T = m).$$

Also,

$$\begin{aligned} P(T > n) &= 1 - P(T \leq n) \\ &= 1 - [a_k + (a_{k+1} - a_k) + \cdots + (a_n - a_{n-1})] \\ &= 1 - a_n \\ &= s(\mu_n; \pi), \end{aligned}$$

and this proves (5.5).

It now remains to check (5.6). We do this by induction, and for this, we assume that (5.6) is true for all $m \leq n - 1$ and all $i \in E$. This implies

$$P(X_n = i, T \leq n - 1) = \pi(i) a_{n-1}, \tag{5.7}$$

because for $m \leq n - 1$,

$$\begin{aligned} P(X_n = i, T = k) &= \sum_{j \in E} P(X_n = i, X_k = j, T = k) \\ &= \sum_{j \in E} P(X_n = i | X_k = j, T = k) P(X_k = j, T = k) \\ &= \sum_{j \in E} p_{ji} \pi(i) (a_k - a_{k-1}) \\ &= \pi(i) (a_k - a_{k-1}), \end{aligned}$$

from which (5.7) follows.

To obtain (5.6) for $m = n$,

$$\begin{aligned} P(X_n = i, T = n) &= P(T = n | X_n = i, T > n - 1) P(X_n = i, T > n - 1) \\ &= \frac{a_n - a_{n-1}}{\frac{\mu_n(i)}{\pi(i)} - a_{n-1}} (P(X_n = i) - P(X_n = i, T \leq n - 1)) \\ &= \frac{a_n - a_{n-1}}{\frac{\mu_n(i)}{\pi(i)} - a_{n-1}} (\mu_n(i) - \pi(i) a_{n-1}) \\ &= \pi(i) (a_n - a_{n-1}). \end{aligned}$$

□

Thus, if T is a minimal strong stationary time, then for any other strong stationary time T' ,

$$P_\mu(T' > n) \geq P_\mu(T > n) \tag{5.8}$$

for all $n \geq 0$, and therefore a time to stationarity is stochastically smaller than any strong stationary time (the random variable X is said to be *stochastically larger* than Y , and this is denoted by $X \geq_{st} Y$, iff $P(X > x) \geq P(Y > x)$ for all $x \in \mathbb{R}$).

We proceed to compare strong stationary times with coupling times.

The construction of a minimal strong stationary time in Theorem 5.2 requires the knowledge of $\mu^T \mathbf{P}^n$ and π , and is therefore mainly of theoretical interest. It shows, however, that the notion of separation fits perfectly with that of strong stationarity. Similarly, the notion of distance in variation fits with that of coupling, as the following analogue of Theorem 5.2 shows.

Theorem 5.3. Minimal Stationary Coupling Time

Let $\{X_n\}_{n \geq 0}$ be an HMC with initial distribution μ , ergodic transition matrix \mathbf{P} , and stationary distribution π . Let T be a finite random time such that there exists an HMC $\{Y_n\}_{n \geq 0}$ with transition matrix \mathbf{P} and initial distribution π such that

$$X_n = Y_n \text{ if } n \geq T. \tag{5.9}$$

Then T is called a *stationary coupling time* of $\{X_n\}_{n \geq 0}$, and

$$d_V(\mu^T \mathbf{P}^n, \pi) \leq P(T > n). \tag{5.10}$$

There exists a stationary coupling time T such that equality holds in (5.10), and it is called a *minimal stationary coupling time*.

Proof. Taking into account the results of Chapter 4, only the existence of a minimal stationary coupling time remains to be proven. We refer to (Lindvall, 1992) for the proof. \square

Set

$$d_\mu(n) = d_V(\mu^T \mathbf{P}^n, \pi), \quad s_\mu(n) = s(\mu^T \mathbf{P}^n; \pi). \tag{5.11}$$

We know that $d_\mu(n) \leq s_\mu(n)$, and therefore the existence of a strong stationary time T implies a bound for convergence in variation, since then

$$d_\mu(n) \leq s_\mu(n) \leq P(T > n).$$

Also, denoting by $T_{s,min}$ and $T_{c,min}$ the minimal strong stationary time and the minimal stationary coupling time, respectively, we see that

$$T_{c,min} \geq_{st} T_{s,min}. \tag{5.12}$$

Theorem 5.4. *Strong Stationary Times Are Stationary Coupling Times*

Let T be a strong stationary time of the HMC $\{X_n\}_{n \geq 0}$. Then T is also a stationary coupling time of the same chain.

Proof. For each $m \geq 0$, define on $\{T = m\}$ the process $\{Y_n^{(m)}\}_{n \geq m}$ by

$$Y_n^{(m)} = X_n \text{ if } n \geq m.$$

Since for $n \geq m$, by definition of a strong stationary time, $P(X_n = i, T = m) = \pi(i) P(T = m)$, we see that, conditionally on $\{T = m\}$, $\{Y_n^{(m)}\}_{n \geq m}$ is a stationary HMC. It can be extended to a stationary HMC $\{Y_n^{(m)}\}_{n \geq 0}$ (see Example 6.1 of Chapter 2). Letting

$$Y_n = Y_n^{(m)} \text{ on } \{T = m\},$$

we obtain an HMC $\{Y_n\}_{n \geq 0}$ that is stationary and such that

$$X_n = Y_n \text{ for } n \geq T. \tag{5.13} \quad \square$$

5.2 Convergence Rates via Strong Stationary Times

Rather than developing an abstract framework, we shall give two examples that clearly show how strong stationary times lead to convergence rates.




Figure 6.5.1. Strong stationary time for the symmetric walk on the circle

Example 5.1. Symmetric Walk on the Circle

(Diaconis and Fill, 1991) Let $\{X_n\}_{n \geq 0}$ be a symmetric random walk on $E = \mathbb{Z}_d$, the integers modulo d , identified with d points on the circle (see Figure 6.5.1). One moves one step in either direction or remains still, each with probability $\frac{1}{3}$.

This chain is clearly ergodic with the uniform probability on E . A strong stationary time can be constructed as follows in the case $d = 2^a$. We do the case $d = 16$ for definiteness.

Starting from 0, let T_1 be the first time either state 4 or 12 is visited. Clearly, X_{T_1} is uniformly distributed on $\{4, 12\}$ and is independent of T_1 . Next, let T_2 be the first time after T_1 when the chain visits the states at distance 2 from X_{T_1} . Then X_{T_2} is uniformly distributed on $\{2, 6, 10, 14\}$ and is independent of T_2 . Time T_3 is now the first time after T_2 when the chain hits a state at distance 1 from X_{T_2} . Then X_{T_3} is uniformly distributed on the odd numbers $\{1, 3, 5, 7, 9, 11, 13, 15\}$ and is independent of T_3 . Finally, let T be the first time after T_3 where the chain makes a clockwise move or stays still. We can take as T the desired strong stationary time, since it is independent of X_{T_3} , and X_T is uniform on E .

For $d = 2^a$, the random walk travels successively the distances $\pm 2^{a-2}, \pm 2^{a-3}, \dots, \pm 1 = \pm 2^{a-a}$. The mean time to travel at distance b of this symmetric walk is (see Example 3.4 of Chapter 2) $\frac{3}{2}b^2$. The last step from T_{a-1} to $T_a = T$ takes $\frac{3}{2}$ time units on the average. Therefore

$$E_0[T_a] = \frac{3}{2}(2^{2a-4} + \dots + 1 + 1) = \frac{3}{2}2^{2a}(2^{-4} + 2^{-6} + \dots + 2^{-2(a-1)} + 2 \cdot 2^{-2a}).$$

Therefore, for $a \geq 2$,

$$E_0[T_a] \leq \frac{3}{16}2^{2a} = \frac{3}{16}d^2.$$

By Markov's inequality,

$$P_0(T_a > n) \leq \frac{E_0[T_a]}{n} \leq \frac{3}{16} \frac{d^2}{n},$$

and therefore, since the result would be the same for any state,

$$d_V(\mu^T \mathbf{P}^n, \pi) \leq \frac{3}{16} \frac{d^2}{n}. \quad \diamond$$

Example 5.2. Top to Random Card Shuffling

(Aldous and Diaconis, 1981) The title refers to a method of shuffling cards whereby the top card of the deck is removed and placed at random in the deck, and the procedure is repeated. This defines an irreducible HMC, where a state is an arrangement of the deck. Its stationary distribution is the uniform distribution (use symmetry, and to make symmetry more apparent, arrange the cards in a circle rather than in a deck). Denote by \star the card originally at the bottom. Let T_j be the j th time a card is inserted below \star . If there are N cards, then at time T_{N-1} card \star has reached the top. Let $T = T_N$. An inductive argument shows that for $j \leq N - 1$, at time T_j all the $j!$ arrangements of the j cards below \star are equally likely, and that T is therefore a stationary time. Also,

$$T = T_1 + (T_2 - T_1) + \dots + (T - T_{N-1}),$$

where $T - T_{N-1} = 1$. At time T_i , card \star has i cards below it, and the probability that the current top card is inserted below \star is therefore $\frac{i+1}{N}$. Therefore, $T_{i+1} - T_i$ is geometric:

$$P(T_{i+1} - T_i = k) = \frac{i+1}{N} \left(1 - \frac{i+1}{N}\right)^{k-1}.$$

Consider now the following problem: Sample uniformly with replacement an urn containing N balls, and denote by V the number of draws until each ball has been sampled at least once. Let V_i be the number of draws until i distinct balls have been sampled at least once. We have the identity

$$V = (V - V_{N-1}) + \dots + (V_2 - V_1) + V_1.$$

Once i distinct balls have been drawn at least once, there is probability $\frac{N-i}{N}$ sampling a ball not previously sampled. Therefore, $V_i - V_{i-1}$ is geometric:

$$P(V_i - V_{i-1} = k) = \frac{N-i}{N} \left(1 - \frac{N-i}{N}\right)^{k-1}.$$

In particular, T and V have the same distribution. For each ball b , let A_b be the event that ball b was not drawn in the first $m = N \log(N) + cN$ draws, $c \geq 0$. We have

$$P(V > m) = P(\cup A_b) = N \left(1 - \frac{1}{N}\right)^m \leq N e^{-\frac{m}{N}} = N e^{-\log(N) - c} = e^{-c}.$$

Therefore

$$d(N \log(N) + cN) \leq (P(T > N \log(N) + cN)) \leq e^{-c},$$

where $d(k) = d_V(\mu^T \mathbf{P}^k, \pi)$. In particular, taking $c = \epsilon \log(N)$,

$$d(N \log(N) (1 + \epsilon)) < \frac{1}{N^\epsilon},$$

for all $\epsilon > 0$. Aldous and Diaconis (1981) show that for all $\epsilon > 0$

$$d(N \log(N) (1 - \epsilon)) \rightarrow 1$$

as $N \rightarrow \infty$. In this sense, for large N , $N \log(N)$ shuffles suffice, but no less. ◇

6 Fundamental Matrix of Recurrent Chains

6.1 Definition of the Fundamental Matrix

The *fundamental matrix* has been defined for absorbing chains in Chapter 4. There is also a fundamental matrix for ergodic chains. These two types of fundamental matrices are different objects. The one for absorbing chains is the restriction of the potential matrix to the transient set. For ergodic recurrent chains, the potential matrix is uninteresting, since all its entries are infinite. However, if the term in the definition of the potential matrix is properly modified, as in the formula (6.2) below, the corresponding series does converge. The fundamental matrix for ergodic chains is interesting because it gives access to a number of quantities such as, for instance, the mean time $E_i[T_j]$ to return to j from state i , or the variance of the ergodic estimate $\frac{1}{n} \sum_{k=1}^n f(X_k)$. For a finite state space ergodic HMC, the fundamental matrix is the matrix

$$\mathbf{Z} = (\mathbf{I} - (\mathbf{P} - \mathbf{\Pi}))^{-1}, \quad (6.1)$$

where

$$\mathbf{\Pi} = \mathbf{1}\pi^T = \begin{pmatrix} \pi(1) & \cdots & \pi(r) \\ \pi(1) & \cdots & \pi(r) \\ \vdots & & \vdots \\ \pi(1) & \cdots & \pi(r) \end{pmatrix}.$$

Here the state space is $E = \{1, \dots, r\}$, and π is the stationary distribution.

Theorem 6.1. Fundamental Matrix of Ergodic Chains

For any ergodic transition matrix \mathbf{P} on a finite state space, the right-hand side of (6.1) is well defined and

$$\mathbf{Z} = \mathbf{I} + \sum_{n \geq 1} (\mathbf{P}^n - \mathbf{\Pi}). \quad (6.2)$$

Proof. First observe that

$$\mathbf{\Pi}\mathbf{P} = \mathbf{\Pi} \text{ (since } \pi^T \mathbf{P} = \pi^T \text{ and } \mathbf{\Pi} = \mathbf{1}\pi^T),$$

$$\mathbf{P}\mathbf{\Pi} = \mathbf{\Pi} \text{ (since } \mathbf{P}\mathbf{1} = \mathbf{1} \text{ and } \mathbf{\Pi} = \mathbf{1}\pi^T),$$

$$\mathbf{\Pi}^2 = \mathbf{\Pi} \text{ (since } \mathbf{\Pi} = \mathbf{1}\pi^T \text{ and } \pi^T \mathbf{1} = 1).$$

In particular, for all $k \geq 1$,

$$\mathbf{P}\mathbf{\Pi}^k = \mathbf{\Pi} = \mathbf{\Pi}^k \mathbf{P}.$$

Therefore,

$$\begin{aligned} (\mathbf{P} - \mathbf{\Pi})^n &= \sum_{k=0}^n \binom{n}{k} (-1)^{n-k} \mathbf{P}^k \mathbf{\Pi}^{n-k} \\ &= \mathbf{P}^n + \left(\sum_{k=0}^{n-1} \binom{n}{k} (-1)^{n-k} \right) \mathbf{\Pi} = \mathbf{P}^n - \mathbf{\Pi}. \end{aligned}$$

Therefore, with $A = \mathbf{P} - \Pi$,

$$(I - A)(I + A + \cdots + A^{n-1}) = I - A^n = I + \mathbf{P}^n - \Pi.$$

Passing to the limit $n \rightarrow \infty$,

$$(I - A)(I + \sum_{n \geq 1} A^n) = I,$$

which shows that $I - (\mathbf{P} - \Pi)$ is invertible, with inverse

$$I + \sum_{n \geq 1} (\mathbf{P} - \Pi)^n = I + \sum_{n \geq 1} (\mathbf{P}^n - \Pi). \quad \square$$

Example 6.1. *The Diagonal of the Fundamental Matrix*

For fixed $m \geq 1$, let

$$S = m + \inf \{k \geq 0; X_{m+k} = i\}.$$

Then by Theorem 4.4 of Chapter 3,

$$E_i \left[\sum_{n=0}^{S-1} 1_{X_n=j} \right] = \pi(i) E_i [S],$$

that is,

$$E_i \left[\sum_{n=0}^{m-1} 1_{X_n=j} \right] = \pi(i) (m + E_{\nu_m} [S_i]),$$

where S_i is the hitting time of i , and ν_m is the distribution of the chain at time m . Therefore,

$$\sum_{n=0}^{m-1} (p_{ii}(n) - \pi(i)) = \pi(i) E_{\nu_m} [S_i].$$

Since $\lim_{m \uparrow \infty} \nu_m = \pi$, we have

$$\sum_{n=0}^{\infty} (p_{ii}(n) - \pi(i)) = \pi(i) E_{\pi} [S_i],$$

which is easily converted into

$$z_{ii} = \pi(i) E_{\pi} [T_i]. \quad (6.3)$$

Example 6.2. *Patterns in Coin Tossing*

(Aldous and Fill, 1998) Let $\{Y_n\}_{n \geq 0}$ be an i.i.d sequence of Bernoulli variables, with $P(Y_1 = 1) = P(Y_1 = 0) = \frac{1}{2}$ and let $\{X_n\}_{n \geq 0}$ be the *snake chain* defined by

$$X_n = (Y_n, Y_{n+1}, \dots, Y_{n+L-1})$$

for some $L \geq 1$. Note that both $\{Y_n\}_{n \geq 0}$ and $\{X_n\}_{n \geq 0}$ are irreducible ergodic chains, starting in the steady state. Define

$$\tilde{z}_{ij} = \sum_{n=0}^{\infty} (p_{ij}(n) - \pi(j)) \tag{6.4}$$

($= z_{ij} - \pi(j)$), where \mathbf{P} is the transition matrix of $\{X_n\}_{n \geq 0}$ and $\pi(j) = \frac{1}{2^L}$ is its stationary distribution.

For $n \geq L$, X_0 and X_n are independent, and therefore $p_{ij}(n) - \pi(j) = 0$, so that only the first L terms of (6.4) are nonzero.

For $n < L$, $p_{ij}(n) > 0$ if and only if the pattern $j = (j_0, \dots, j_{L-1})$ shifted n to the right and the pattern $i = (i_0, \dots, i_{L-1})$ agree where they overlap (see Figure 6.6.1).




Figure 6.6.1.

In this case $p_{ij}(n)$ equals $\frac{1}{2^n}$. Therefore, defining

$$c(i, j) = \sum_{n=0}^{L-1} \frac{1}{2^n} \chi(i, j, n), \tag{6.5}$$

where $\chi(i, j, n) = 1$ if and only if the situation depicted in Figure 6.6.1 is realized,

$$\tilde{z}_{ij} = c(i, j) - L2^{-L}.$$

In view of the result of the previous example,

$$E_{\pi} [S_i] = 2^L c(i, i) - L.$$

But remember that X_0 is always distributed as π , and that to generate the first pattern, L coin tosses are necessary. Therefore, $2^L c(i, i)$ is the average number of coin tosses needed to see pattern i for the first time.

To illustrate this, consider the pattern $i = \text{HTTTHT}$. We have $c(i, i) = 68$ (see Fig. 6.6.2).

An Extension of the Fundamental Matrix

Expression (6.2) is meaningful only if the chain is ergodic. In particular, if the chain is only recurrent positive, but periodic, the series on the right-hand side of (6.2) oscillates. This does not mean, however, that in the periodic case the inverse in (6.2) does not exist. As a matter of fact, it does exist, but it is not given by formula (6.2).

H	T	T	T	H	T	2^{-0}	
	H	T	T	T	H	0	
		H	T	T	T	0	$c(\text{HTTTHT}) = 2^6 \left(1 + \frac{1}{2^4}\right)$
			H	T	T	0	
				H	T	2^{-4}	
				H		0	
				H		0	

Figure 6.6.2.

The following is an extension of the definition of the fundamental matrix due to Kemeny (1991); see (Grinstead and Snell, 1997) for additional details). This extension does not require, in principle, knowledge of the stationary distribution.

Let b be any vector such that

$$b^T \mathbf{1} \neq 0, \tag{6.6}$$

and define

$$\mathbf{Z} = (I - \mathbf{P} + \mathbf{1}b^T)^{-1}, \tag{6.7}$$

where \mathbf{P} is an ergodic matrix on the finite space E , with the stationary distribution π . The matrix differs from the usual fundamental matrix in that π is replaced by b .

Theorem 6.2.

The inverse matrix in (6.7) exists and

$$\pi^T = b^T \mathbf{Z}. \tag{6.8}$$

Proof. Since $\pi^T \mathbf{1} = 1$ and $\pi^T (I - \mathbf{P}) = 0$,

$$\pi^T (I - \mathbf{P} + \mathbf{1}b^T) = \pi^T \mathbf{1}b^T = b^T, \tag{6.9}$$

and therefore, for any vector x such that

$$(I - \mathbf{P} + \mathbf{1}b^T)x = 0, \tag{6.10}$$

we have

$$b^T x = 0$$

and

$$(I - \mathbf{P})x = 0.$$

Therefore, x must be a right eigenvector associated with the eigenvalue $\lambda_1 = 1$, and consequently, x is a multiple of $\mathbf{1}$. But this is compatible with $b^T x = 0$ and $b^T \mathbf{1} \neq 0$ only if $x = 0$. Therefore (6.10) implies $x = 0$, which implies that $(I - \mathbf{P} + \mathbf{1}b^T)$ is invertible; and (6.9) proves (6.8). □

6.2 Mutual Time–Distance Matrix

We introduce some notation: For any square matrix B , $d(B)$ is the diagonal matrix which has the same diagonal as B . In particular $d(\Pi)^{-1}$ is the diagonal matrix for which the (i, i) -th entry is $\pi(i)^{-1}$. Recall also that $\mathbf{1}\mathbf{1}^T$ is the matrix with all entries equal to 1.

Theorem 6.3. Mutual Time–Distance Matrix

The mutual time distance matrix $M = \{m_{ij}\}_{1 \leq i, j \leq r}$, defined by $m_{ij} = E_i[T_j]$, is given by the formula

$$M = (I - \mathbf{Z} + \mathbf{1}\mathbf{1}^T d(\mathbf{Z}))d(\Pi)^{-1}. \quad (6.11)$$

Proof. We first observe that M has finite entries. Indeed, we already know that $m_{ii} = E_i[T_i] = 1/\pi(i)$ and that $\pi(i) > 0$. As for m_{ij} when $i \neq j$, it is the mean time to absorption in the modified chain where j is made absorbing. By Remark 5.2 of Chapter 4, it is finite.

First-step analysis gives

$$m_{ij} = 1 + \sum_{k, k \neq j} p_{ik} m_{kj}, \quad (6.12)$$

that is,

$$M = \mathbf{P}(M - d(M)) + \mathbf{1}\mathbf{1}^T. \quad (6.13)$$

We now prove that there is but one finite solution of (6.13). To do this, we first show that for any solution M of (6.13), $d(M)$ is necessarily equal to $d(\Pi)^{-1}$. (We know this to be true when M is the mutual distance matrix, but not yet for a general solution of (6.13).) Indeed, premultiplying (6.13) by π^T yields

$$\begin{aligned} \pi^T M &= \pi^T \mathbf{P}(M - d(M)) + (\pi^T \mathbf{1})\mathbf{1}^T \\ &= \pi^T (M - d(M)) + \mathbf{1}^T, \end{aligned}$$

and therefore $\pi^T d(M) = \mathbf{1}^T$, which implies the announced result.

Now suppose that (6.13) has two finite solutions M_1 and M_2 . Since $d(M_1) = d(M_2)$, it follows that

$$M_1 - M_2 = \mathbf{P}(M_1 - M_2).$$

Therefore, any column v of $M_1 - M_2$ is a righteigenvector of \mathbf{P} corresponding to the eigenvalue 1. We know that the righteigenspace R_λ and the lefteigenspace L_λ corresponding to any given eigenvalue λ have the same dimension. For $\lambda = 1$, we know that the dimension of L_λ is one, since there is only one stationary distribution for an ergodic chain. Therefore, R_λ has dimension 1 for $\lambda = 1$. Thus any righteigenvector is a scalar multiple of $\mathbf{1}$. Therefore, $M_1 - M_2$ has columns of the type $\alpha\mathbf{1}$ for some α (α may a priori depend on the column). Since $d(M_1) = d(M_2)$, each column contains a zero, and therefore $\alpha = 0$ for all columns, i.e., $M_1 - M_2 \equiv 0$.

At this point we have proven that M is the unique finite solution of (6.13). It remains to show that M defined by (6.11) satisfies equation (6.13). Indeed, from (6.11) and $d(M) = d(\Pi)^{-1}$,

$$M - d(\Pi)^{-1} = (-\mathbf{Z} + \mathbf{1}\mathbf{1}^T d(\mathbf{Z}))d(\Pi)^{-1}.$$

Therefore,

$$\begin{aligned} \mathbf{P}(M - d(\Pi)^{-1}) &= (-\mathbf{P}\mathbf{Z} + \mathbf{P}\mathbf{1}\mathbf{1}^T d(\mathbf{Z}))d(\Pi)^{-1} \\ &= (-\mathbf{P}\mathbf{Z} + \mathbf{1}\mathbf{1}^T d(\mathbf{Z}))d(\Pi)^{-1} \\ &= M + (-\mathbf{P}\mathbf{Z} - \mathbf{I} + \mathbf{Z})d(\Pi)^{-1}, \end{aligned}$$

where we have used the identity $\mathbf{P}\mathbf{1} = \mathbf{1}$ for the second equality and (6.11) again for the third. Using now (6.1), i.e. $\mathbf{I} - \mathbf{Z} = \Pi - \mathbf{P}\mathbf{Z}$, we see that

$$\mathbf{P}(M - d(\Pi)^{-1}) = M - \Pi d(\Pi)^{-1} = M - \mathbf{1}\mathbf{1}^T,$$

and (6.13) follows, since $d(M) = d(\Pi)^{-1}$. \square

Theorem 6.4. *Time-Distance and the Fundamental Matrix*

Let \mathbf{Z} be the fundamental matrix as in (6.7). Then for all $i \neq j$,

$$E_i [T_j] = \frac{z_{jj} - z_{ij}}{\pi(j)}. \quad (6.14)$$

Proof. We shall need two preliminary formulas. First,

$$\mathbf{Z}\mathbf{1} = \theta\mathbf{1}, \quad (6.15)$$

where $\theta^{-1} = b^T\mathbf{1}$. Indeed, from the definition of \mathbf{Z} ,

$$\mathbf{Z}(\mathbf{I} - \mathbf{P} + \mathbf{1}b^T)\mathbf{1} = \mathbf{1}. \quad (6.16)$$

But $(\mathbf{I} - \mathbf{P})\mathbf{1} = 0$, and therefore (6.15) follows.

Next we need the formula

$$\mathbf{Z}(\mathbf{I} - \mathbf{P}) = \mathbf{I} - \theta\mathbf{1}b^T, \quad (6.17)$$

which follows from (6.7) and (6.15).

We now proceed to the main part of the proof. Call N the mutual distance matrix M in which the diagonal elements have been replaced by 0's. From (6.13), we obtain

$$(\mathbf{I} - \mathbf{P})N = \mathbf{1}\mathbf{1}^T - D^{-1},$$

where $D = \text{diag} \{ \pi(1), \dots, \pi(n) \}$. Multiplying both sides by \mathbf{Z} , and using (6.15), we obtain

$$\mathbf{Z}(\mathbf{I} - \mathbf{P})N = \theta\mathbf{1}\mathbf{1}^T - \mathbf{Z}D^{-1}.$$

Using (6.17),

$$\mathbf{Z}(I - \mathbf{P})N = N - \theta \mathbf{1}b^T N.$$

Therefore,

$$N = \theta \mathbf{1}\mathbf{1}^T - \mathbf{Z}D^{-1} + \theta \mathbf{1}b^T N.$$

Thus, for all $i, j \in E$,

$$n_{ij} = \theta - \frac{z_{ij}}{\pi(j)} + \theta (b^T N)_j.$$

For $i = j$, $n_{ij} = \theta - \frac{z_{jj}}{\pi(j)} + \theta (b^T N)_j = 0$, which gives $(b^T N)_j$. Finally, for $i \neq j$,

$$n_{ij} = \frac{z_{jj} - z_{ij}}{\pi(j)}. \quad \square$$

6.3 Variance of Ergodic Estimates

For a positive recurrent Markov chain, we know from the ergodic theorem that the estimate $\frac{1}{n} \sum_{k=1}^n f(X_k)$ of $E_\pi[f(X_0)] = \langle f \rangle_\pi$ (where $f : E \rightarrow \mathbb{R}$ is such that $\langle |f| \rangle_\pi < \infty$, a condition that is always satisfied when the state space is finite) is asymptotically unbiased, in the sense that it converges to $\langle f \rangle_\pi$ as $n \rightarrow \infty$.

Theorem 6.5. Asymptotic Variance of Ergodic Estimates

Let $\{X_n\}_{n \geq 0}$ be an ergodic Markov chain with finite state space. For any function $f : E \rightarrow \mathbb{R}$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} V_{P_\mu} \left(\sum_{k=1}^n f(X_k) \right) = 2 \langle f, \mathbf{Z}f \rangle_\pi - \langle f, (I + \Pi)f \rangle_\pi \quad (6.18)$$

for any initial distribution μ .

In (6.18), $f^T = (f(1), \dots, f(r))$, where r is the number of states. The notation V_{P_μ} indicates that the variance is computed with respect to P_μ .

The quantity (6.18) will be denoted by $v(f, \mathbf{P}, \pi)$.

Proof. We first treat the case where $\mu = \pi$, the stationary distribution. To simplify, we write V_π for V_{P_π} . Then

$$\begin{aligned} \frac{1}{n} V_\pi \left(\sum_{k=1}^n f(X_k) \right) &= \frac{1}{n} \left\{ \sum_{k=1}^n V_\pi(f(X_k)) + 2 \sum_{\substack{k,j=1 \\ k < j}}^n \text{cov}_\pi(f(X_k), f(X_j)) \right\} \\ &= V_\pi(f(X_0)) + \sum_{\ell=1}^{n-1} \frac{n-\ell}{n} \text{cov}_\pi(f(X_0), f(X_\ell)) \end{aligned}$$

where we have used the fact that when the initial distribution is π , the chain is stationary, and in particular, $\text{cov}_\pi(f(X_k), f(X_j)) = \text{cov}_\pi(f(X_0), f(X_{j-k}))$ for $k < j$. Now,

$$\begin{aligned} V_\pi(f(X_0)) &= E_\pi[f(X_0)^2] - E_\pi[f(X_0)]^2 \\ &= \sum_{i \in E} \pi(i) f(i)^2 - \left(\sum_{i \in E} \pi(i) f(i) \right)^2 \\ &= \langle f, f \rangle_\pi - \langle f, \Pi f \rangle_\pi. \end{aligned}$$

Also,

$$\begin{aligned} \text{cov}_\pi(f(X_0), f(X_\ell)) &= E_\pi[f(X_0)f(X_\ell)] - E_\pi[f(X_0)]^2 \\ &= \sum_{i \in E} \sum_{j \in E} \pi(i) p_{ij}(\ell) f(i) f(j) - E_\pi[f(X_0)]^2 \\ &= \langle f, \mathbf{P}^\ell f \rangle_\pi - \langle f, \Pi f \rangle_\pi \\ &= \langle f, (\mathbf{P}^\ell - \Pi) f \rangle_\pi. \end{aligned}$$

Since $\lim_{n \rightarrow \infty} \sum_{\ell=1}^n (\mathbf{P}^\ell - \Pi) = \mathbf{Z} - I$, we have

$$\lim_{n \rightarrow \infty} \sum_{\ell=1}^{n-1} \frac{n-\ell}{n} (\mathbf{P}^\ell - \Pi) = \mathbf{Z} - I.$$

(This is Cesaro's lemma (Theorem 1.5 of the appendix): If $A_n = \sum_{\ell=1}^n \alpha_\ell$ tends to A as $n \rightarrow \infty$, then $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{\ell=1}^{n-1} A_\ell = A$. But $\frac{1}{n} \sum_{\ell=1}^{n-1} A_\ell = \frac{1}{n}(\alpha_1 + (\alpha_1 + \alpha_2) + \dots + (\alpha_1 + \dots + \alpha_{n-1})) = \sum_{\ell=1}^{n-1} \frac{n-\ell}{n} \alpha_\ell$.) Therefore,

$$\lim_{n \rightarrow \infty} \frac{1}{n} V_\pi \left(\sum_{k=1}^n f(X_k) \right) = \langle f, f \rangle_\pi - \langle f, \Pi f \rangle_\pi + 2 \langle f, (\mathbf{Z} - I) f \rangle_\pi,$$

which is the announced result (for $\mu = \pi$).

To prove the result in the general case where the initial distribution is arbitrary, it suffices to show that for two chains $\{X_n^{(1)}\}_{n \geq 0}$ and $\{X_n^{(2)}\}_{n \geq 0}$ with transition matrix \mathbf{P} , and arbitrary initial distributions μ and ν , respectively, that couple at a time τ such that $E[\tau^2] < \infty$ (this is the case here, see Theorem 2.6 of Chapter 4),

$$\lim_{n \rightarrow \infty} \frac{1}{n} V \left(\sum_{k=1}^{\infty} f(X_k^{(1)}) \right) = \lim_{n \rightarrow \infty} \frac{1}{n} V \left(\sum_{k=1}^{\infty} f(X_k^{(2)}) \right).$$

But with $X_n = X_n^{(1)}$ or $X_n^{(2)}$,

$$V \left(\sum_{k=1}^n f(X_k) \right) = E \left[\left(\sum_{k=1}^n f(X_k) \right)^2 \right] - E \left[\sum_{k=1}^n f(X_k) \right]^2$$

$$\begin{aligned}
&= E \left[\left(\sum_{k=1}^{\tau \wedge n} + \sum_{k=\tau+1}^n \right)^2 \right] - \left(E \left[\sum_{k=1}^{\tau \wedge n} \right] + E \left[\sum_{k=\tau+1}^n \right] \right)^2 \\
&= E \left[\left(\sum_{k=1}^{\tau \wedge n} \right)^2 \right] + E \left[\left(\sum_{k=\tau+1}^n \right)^2 + 2E \left[\left(\sum_{k=1}^{\tau \wedge n} \right) \left(\sum_{k=\tau+1}^n \right) \right] \right] \\
&\quad - E \left[\sum_{k=1}^{\tau \wedge n} \right]^2 - E \left[\sum_{k=\tau+1}^n \right]^2 - 2E \left[\sum_{k=1}^{\tau \wedge n} \right] E \left[\sum_{k=\tau+1}^n \right].
\end{aligned}$$

Since $\sum_{k=\tau+1}^n f(X_k^{(1)}) = \sum_{k=\tau+1}^n f(X_k^{(2)})$, it follows (with obvious shorthand notations) that

$$\frac{1}{n} \left\{ V \left(\sum_{k=1}^n f(X_k^{(1)}) \right) - \frac{1}{n} V \left(\sum_{k=1}^n f(X_k^{(2)}) \right) \right\} = \frac{1}{n} A_n + \frac{2}{n} B_n - \frac{2}{n} C_n,$$

where

$$\begin{aligned}
A_n &= \left\{ E \left[\left(\sum_{k=1}^{\tau \wedge n} (1) \right)^2 \right] - E \left[\left(\sum_{k=1}^{\tau \wedge n} (2) \right)^2 \right] - E \left[\sum_{k=1}^{\tau \wedge n} (1) \right]^2 + E \left[\sum_{k=1}^{\tau \wedge n} (2) \right]^2 \right\}, \\
B_n &= \left\{ E \left[\left(\sum_{k=\tau+1}^n (1, 2) \right) \left(\sum_{k=1}^{\tau \wedge n} (1) - \sum_{k=1}^{\tau \wedge n} (2) \right) \right] \right\}, \\
C_n &= \left\{ E \left[\sum_{k=\tau+1}^n (1, 2) \right] E \left[\sum_{k=1}^{\tau \wedge n} (1) - \sum_{k=1}^{\tau \wedge n} (2) \right] \right\}.
\end{aligned}$$

Write

$$\frac{2}{n} B_n = 2E \left[\frac{\sum_{k=\tau+1}^n (1, 2)}{n} \left(\sum_{k=1}^{\tau \wedge n} (1) - \sum_{k=1}^{\tau \wedge n} (2) \right) \right]$$

and observe that the quantity under the expectation converges, as $n \rightarrow \infty$, towards $E_\pi[f(X_0)]$ ($\sum_{k=1}^{\tau} (f(X_k^{(1)}) - f(X_k^{(2)}))$) and is for fixed n bounded in absolute value by $2(\sup |f|)\tau$, an integrable random variable. Therefore, by dominated convergence,

$$\lim_{n \rightarrow \infty} \frac{2}{n} B_n = 2E_\pi[f(X_0)] E \left[\sum_{k=1}^{\tau} (f(X_k^{(1)}) - f(X_k^{(2)})) \right].$$

A similar argument shows that $\frac{2}{n} C_n$ has the same limit. Therefore, $\lim_{n \rightarrow \infty} \frac{2}{n} (B_n - C_n) = 0$. As for A_n , it is bounded by $4(\sup |f|)^2 E[\tau^2] < \infty$, and therefore $\lim_{n \rightarrow \infty} \frac{1}{n} A_n = 0$. \square

We shall now give an expression of the asymptotic variance in terms of the eigenvalues, when \mathbf{P} has r distinct eigenvalues. We have, in view of (1.1),

$$(\mathbf{P}^n - \Pi) = \sum_{i=2}^r \lambda_i^n v_i u_i^T,$$

and therefore

$$\mathbf{Z} = I + \sum_{n \geq 1} (\mathbf{P}^n - \Pi) = I + \sum_{i=2}^r \frac{\lambda_i}{1 - \lambda_i} v_i u_i^T. \tag{6.19}$$

Also, from (6.18),

$$v(f, \mathbf{P}, \pi) = V_\pi(f(X_0)) + 2 \sum_{i=2}^r \frac{\lambda_i}{1 - \lambda_i} \langle f, v_i \rangle_\pi (f^T u_i). \tag{6.20}$$

For a reversible pair (\mathbf{P}, π) , we have $u_i = Dv_i$, and therefore $f^T u_i = \langle f, v_i \rangle_\pi$. Using this observation and (2.9), we obtain from (6.20),

$$v(f, \mathbf{P}, \pi) = \sum_{i=2}^r \frac{1 + \lambda_i}{1 - \lambda_i} |\langle f, v_i \rangle_\pi|^2. \tag{6.21}$$

Recall that when one is interested in the speed of convergence to equilibrium, it is the second-largest eigenvalue modulus that is important. However, if one is interested in simulation, i.e., the computation of $E_\pi[f(X_0)]$ as the ergodic mean $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n f(X_k)$, all eigenvalues play a role if we measure the quality of the ergodic estimator by the asymptotic variance, as the above formulas show.

7 The Ergodic Coefficient

7.1 Dobrushin’s Inequality

The Perron–Frobenius theorem shows that the SLEM determines the rate of convergence to equilibrium of a finite state space ergodic HMC. We have seen in previous sections methods for obtaining upperbounds of the SLEM, mainly for reversible matrices. For *theoretical* purposes, there is an upper bound of the SLEM that is often very tractable, due to Dobrushin (1956); see (Isaacson and Madsen, 1976) or (Seneta, 1981) for additional references. When the state space E is finite and the chain is ergodic, it provides a computable geometric rate of convergence to steady state. It will be especially useful in the present chapter to obtain a necessary and sufficient condition of weak ergodicity (yet to be defined) of *nonhomogeneous* Markov chains (also to be defined).

In the sequel, E, F, G denote countable sets, finite when so indicated.

Definition 7.1. Ergodic Coefficient

Let \mathbf{Q} be a stochastic matrix indexed by $F \times E$. Its Dobrushin’s ergodic coefficient $\delta(\mathbf{Q})$ is defined by

$$\delta(\mathbf{Q}) = \frac{1}{2} \sup_{i, j \in F} \sum_{k \in E} |q_{ik} - q_{jk}| = \sup_{i, j \in F} d_V(q_i, q_j). \tag{7.1}$$

Observe that

$$0 \leq \delta(\mathbf{Q}) \leq 1, \tag{7.2}$$

and, using (1.5) of Chapter 4,

$$\delta(\mathbf{Q}) = 1 - \inf_{i, j \in F} \sum_{k \in E} q_{ik} \wedge q_{jk}. \quad (7.3)$$

The ergodic coefficient is most of the time useless if F is infinite. Indeed, if the stochastic matrix \mathbf{Q} has two orthogonal rows (that is, rows i, j such that $q_{ik}q_{jk} = 0$ for all $k \in E$, and this is the most frequent case with an infinite state space), then $\delta(\mathbf{Q}) = 1$. However, for finite state spaces, this notion becomes very powerful.

Example 7.1.

The ergodic coefficient of

$$\mathbf{Q} = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{3} \\ 0 & \frac{1}{3} & \frac{2}{3} \\ \frac{3}{4} & \frac{1}{4} & 0 \end{pmatrix}$$

is

$$\frac{1}{2} \sup \left(\frac{1}{2} + \frac{1}{3} + \left(\frac{2}{3} - \frac{1}{2} \right), \frac{3}{4} + \left(\frac{1}{3} - \frac{1}{4} \right) + \frac{2}{3}, \left(\frac{3}{4} - \frac{1}{2} \right) + \frac{1}{4} + \frac{1}{2} \right) = \frac{3}{4}. \quad \diamond$$

Example 7.2. 2-Row Matrix

An important special case is when \mathbf{Q} has only two rows:

$$\mathbf{Q} = \begin{pmatrix} \mu_1 & \mu_2 & \mu_3 & \cdots \\ \nu_1 & \nu_2 & \nu_3 & \cdots \end{pmatrix}.$$

The ergodic coefficient is then just the distance in variation

$$\delta(\mathbf{Q}) = d_V(\mu, \nu). \quad (7.4)$$

\diamond

Let $\mathbf{Q}_1 = \{a_{ij}\}$ and $\mathbf{Q}_2 = \{b_{ij}\}$ be two stochastic matrices indexed by $F \times G$ and $G \times E$, respectively, and define the stochastic matrix $\mathbf{Q} = \{q_{ij}\}$ indexed by $F \times E$ by

$$q_{ij} = \sum_{k \in G} a_{ik} b_{kj}.$$

Theorem 7.1. Dobrushin's Inequality

$$\delta(\mathbf{Q}_1 \mathbf{Q}_2) \leq \delta(\mathbf{Q}_1) \delta(\mathbf{Q}_2). \quad (7.5)$$

Proof. By Lemma 1.1 of Chapter 4,

$$\frac{1}{2} \sum_{k \in E} |q_{ik} - q_{jk}| = \sup_{A \subset E} \sum_{k \in A} (q_{ik} - q_{jk}),$$

and therefore

$$\delta(\mathbf{Q}_1 \mathbf{Q}_2) = \frac{1}{2} \sup_{i,j \in F} \sup_{A \subseteq E} \sum_{k \in A} \left(\sum_{\ell \in G} (a_{i\ell} - a_{j\ell}) b_{\ell k} \right).$$

But using (1.5) of Chapter 4,

$$\sum_{\ell \in G} (a_{i\ell} - a_{j\ell})^+ = \sum_{\ell \in G} (a_{i\ell} - a_{j\ell})^- = \frac{1}{2} \sum_{\ell \in G} |a_{i\ell} - a_{j\ell}|,$$

and therefore

$$\begin{aligned} & \sum_{k \in A} \sum_{\ell \in G} (a_{i\ell} - a_{j\ell}) b_{\ell k} \\ &= \sum_{\ell \in G} (a_{i\ell} - a_{j\ell})^+ \sum_{k \in A} b_{\ell k} - \sum_{\ell \in G} (a_{i\ell} - a_{j\ell})^- \sum_{k \in A} b_{\ell k} \\ &\leq \left(\frac{1}{2} \sum_{\ell \in G} |a_{i\ell} - a_{j\ell}| \right) \sup_{\ell \in G} \sum_{k \in A} b_{\ell k} - \left(\frac{1}{2} \sum_{\ell \in G} |a_{i\ell} - a_{j\ell}| \right) \inf_{\ell \in G} \sum_{k \in A} b_{\ell k} \\ &\leq \left(\frac{1}{2} \sum_{\ell \in G} |a_{i\ell} - a_{j\ell}| \right) \left(\sup_{\ell, \ell' \in G} \sum_{k \in A} (b_{\ell k} - b_{\ell' k}) \right). \end{aligned}$$

The announced inequality follows, since

$$\sup_{A \subseteq E} \sup_{\ell, \ell' \in G} \sum_{k \in A} (b_{\ell k} - b_{\ell' k}) = \frac{1}{2} \sup_{\ell, \ell' \in G} \sum_{k \in E} |b_{\ell k} - b_{\ell' k}| = \delta(\mathbf{Q}_2). \quad \square$$

Theorem 7.2. *Convergence Rate in Terms of Dobrushin's Coefficient* Let \mathbf{P} be a stochastic matrix indexed by E , and let μ and ν be two probability distributions on E . Then

$$d_V(\mu^T \mathbf{P}^n, \nu^T \mathbf{P}^n) \leq d_V(\mu, \nu) \delta(\mathbf{P})^n. \quad (7.6)$$

Proof. Let

$$\mathbf{Q} = \begin{pmatrix} \mu(1) & \mu(2) & \cdots \\ \nu(1) & \nu(2) & \cdots \end{pmatrix}, \quad \mathbf{Q}_2 = \mathbf{P},$$

and therefore

$$\mathbf{Q} = \mathbf{Q}_1 \mathbf{Q}_2 = \begin{pmatrix} (\mu^T \mathbf{P})_1 & (\mu^T \mathbf{P})_2 & \cdots \\ (\nu^T \mathbf{P})_1 & (\nu^T \mathbf{P})_2 & \cdots \end{pmatrix}.$$

Inequality (7.6) with $n = 1$ then follows from Dobrushin's inequality (7.5), because as observed in Example 7.2,

$$\delta(\mathbf{Q}_1) = d_V(\mu, \nu), \quad \delta(\mathbf{Q}) = d_V(\mu^T \mathbf{P}, \nu^T \mathbf{P}).$$

For $n > 1$, replace \mathbf{P} by its n th iterate, and use Dobrushin's inequality: $\delta(\mathbf{P}^n) \leq \delta(\mathbf{P})^n$. \square

We shall give a variant of Theorem 7.2 that will be needed later. For this, we shall use the fact that

$$|A| = \sup_{i \in E} \sum_{j \in E} |a_{ij}|$$

defines a norm on the vector space of square matrices.

Corollary 7.1.

Let \mathbf{Q}_1 , \mathbf{Q}_2 , and \mathbf{P} be stochastic matrices indexed by $E \times E$. Then

$$|(\mathbf{Q}_1 - \mathbf{Q}_2)\mathbf{P}| \leq |\mathbf{Q}_1 - \mathbf{Q}_2|\delta(\mathbf{P}). \quad (7.7)$$

Proof. Call μ_k and ν_k the k th rows of \mathbf{Q}_1 and \mathbf{Q}_2 , respectively. Inequality (7.7) reads

$$\sup_{k \in E} |\mu_k^T \mathbf{P} - \nu_k^T \mathbf{P}| \leq \left(\sup_{k \in E} |\mu_k - \nu_k| \right) \delta(\mathbf{P})$$

and is therefore a direct consequence of Theorem 7.2. \square

Theorem 7.2 implies that Dobrushin's coefficient is an upper bound of the SLEM, as a comparison of (7.6) above and (1.3) shows.

7.2 Interaction Coefficients and Coincidence

Recall that the distance in variation between two probability distributions α and β on a set E has the following interpretation (see Theorem 1.1 of Chapter 4): $1 - d_V(\alpha, \beta)$ is the maximum value of $P(X = Y)$ among all couples (X, Y) of E -valued random variables having the prescribed marginals α and β . Therefore, clearly, Dobrushin's coefficient can be interpreted in terms of coupling. We shall use this interpretation in a slightly more general setting.

Definition 7.2. *Attraction and Binding*

Let \mathbf{P} and \mathbf{Q} be two transition matrices on the same state space E . Define the *attraction coefficient*

$$\alpha(\mathbf{P}, \mathbf{Q}) = 1 - \frac{1}{2} \sup_{i, j \in E, i \neq j} \sum_{k \in E} |p_{ik} - q_{jk}| \quad (7.8)$$

and the *binding coefficient*

$$\beta(\mathbf{P}, \mathbf{Q}) = 1 - \frac{1}{2} \sup_{i \in E} \sum_{k \in E} |p_{ik} - q_{ik}|. \quad (7.9)$$

Clearly, $\alpha(\mathbf{P}, \mathbf{P}) = 1 - \delta(\mathbf{P}, \mathbf{P})$ and $\beta(\mathbf{P}, \mathbf{P}) = 1$.

We shall now interpret the above coefficients in terms of "coincidence" of the two HMCs governed by \mathbf{P} and \mathbf{Q} , respectively. For this we construct a Markov chain $\{Z_n\}_{n \geq 0} =$

$\{(X_n, Y_n)\}_{n \geq 0}$ such that $\{X_n\}_{n \geq 0}$ and $\{Y_n\}_{n \geq 0}$ are HMCs with the respective transition matrices \mathbf{P} and \mathbf{Q} , as follows.

If $X_n = i, Y_n = j$, construct X_{n+1} and Y_{n+1} in such a way that

$$P(X_{n+1} = k | X_n = i, Y_n = j) = p_{ik} \text{ and } P(Y_{n+1} = k | X_n = i, Y_n = j) = q_{jk},$$

with maximal coupling. In particular, for $i \neq j$,

$$P(X_{n+1} = Y_{n+1} | X_n = i, Y_n = j) = 1 - d_V(p_i, q_j) \geq \alpha(\mathbf{P}, \mathbf{Q}),$$

and for $i = j$,

$$P(X_{n+1} = Y_{n+1} | X_n = i, Y_n = j) = 1 - d_V(p_i, q_j) \geq \beta(\mathbf{P}, \mathbf{Q}).$$

Thus, if the chains do not coincide at time n , they will at time $n + 1$ with a probability at least $\alpha(\mathbf{P}, \mathbf{Q})$, whence the appellation *attraction coefficient*. Similarly, if the chains coincide at time n , they will still coincide at time $n + 1$ with a probability at least $\beta(\mathbf{P}, \mathbf{Q})$, whence the appellation *binding coefficient*.

Theorem 7.3. Coincidence Rate Lower Bound

The long-range coincidence rate satisfies

$$\lim_{N \uparrow \infty} \frac{1}{N} \sum_{k=1}^N 1_{\{X_k=Y_k\}} \geq \frac{\alpha(\mathbf{P}, \mathbf{Q})}{1 + \alpha(\mathbf{P}, \mathbf{Q}) - \beta(\mathbf{P}, \mathbf{Q})}. \tag{7.10}$$

Proof. Suppose (for instance) that $X_0 \neq Y_0$, and call D_0, D_1, \dots the lengths of the successive time intervals where the two chains are different and C_1, C_2, \dots the lengths of the successive intervals where they coincide. Then for some i.i.d. sequences $\hat{D}_0, \hat{D}_1, \dots$ and $\hat{C}_1, \hat{C}_2, \dots$ that are independent, and where \hat{D}_k is a geometric random variable with parameter $\alpha(\mathbf{P}, \mathbf{Q})$ and \hat{C}_l is geometric with parameter $1 - \beta(\mathbf{P}, \mathbf{Q})$, we have for all $k \geq 0, l \geq 1$,

$$D_k \leq \hat{D}_k, \quad C_l \geq \hat{C}_l. \tag{7.11}$$

The proof follows readily from these observations (Problem 6.7.4). □

8 Nonhomogeneous Markov Chains

8.1 Ergodicity of Nonhomogeneous Markov Chains

Let $\{X_n\}_{n \geq 0}$ be a stochastic process with values in a countable set E , and such that

$$P(X_{n+1} = j | X_n = i, X_0 = i_0, \dots, X_{n-1} = i_{n-1}) = P(X_{n+1} = j | X_n = i) \tag{8.1}$$

for all $n \geq 0$ and all states $i, j, i_0, i_1, \dots, i_{n-1}$. In other words, $\{X_n\}_{n \geq 0}$ is a Markov chain. In this section, we consider *nonhomogeneous* Markov chains, that is, the quantity

$$p_{n,i,j} = P(X_{n+1} = j | X_n = i), \tag{8.2}$$

in general depends on n . The matrix

$$\mathbf{P}(n) = \{p_{n,i,j}\}_{i,j \in E} \quad (8.3)$$

is called the transition matrix at time n . We introduce the notation

$$\mathbf{P}(m, k) = \mathbf{P}(m)\mathbf{P}(m+1) \cdots \mathbf{P}(k-1) \quad (k > m \geq 0). \quad (8.4)$$

Using Bayes's sequential rule, we see that if the distribution of X_m is μ_m , the distribution of X_k is $\mu_m^T \mathbf{P}(m, k)$.

Definition 8.1. *Weak Ergodicity*

The above nonhomogeneous Markov chain is called *weakly ergodic* if for all $m \geq 0$,

$$\limsup_{k \uparrow \infty} d_V(\mu^T \mathbf{P}(m, k), \nu^T \mathbf{P}(m, k)) = 0, \quad (8.5)$$

where the supremum is taken over all the probability distributions μ, ν on E .

Definition 8.2. *Strong Ergodicity*

The chain is called *strongly ergodic* if there exists a probability distribution π on E such that for all $m \geq 0$,

$$\limsup_{k \uparrow \infty} d_V(\mu^T \mathbf{P}(m, k), \pi) = 0, \quad (8.6)$$

where the supremum is taken over all the probability distributions μ on E .

Sometimes one says that the family of transition matrices $\{\mathbf{P}(n)\}_{n \geq 0}$ (rather than the chain) is weakly ergodic (resp., strongly ergodic).

Strong ergodicity implies weak ergodicity, since

$$d_V(\mu^T \mathbf{P}(m, k), \nu^T \mathbf{P}(m, k)) \leq d_V(\mu^T \mathbf{P}(m, k), \pi) + d_V(\nu^T \mathbf{P}(m, k), \pi).$$

However, there are weakly ergodic chains that are not strongly ergodic, as the following example shows.

Example 8.1. *Weakly but Not Strongly Ergodic*

Here the state space has 2 elements, $\mathbf{P}(0) = I$, the identity, and for $n \geq 1$,

$$\mathbf{P}(2n) = \begin{pmatrix} 1/2n & 1 - 1/2n \\ 1/2n & 1 - 1/2n \end{pmatrix}, \quad \mathbf{P}(2n+1) = \begin{pmatrix} 1 - 1/(2n+1) & 1/(2n+1) \\ 1 - 1/(2n+1) & 1/(2n+1) \end{pmatrix}.$$

Elementary computations show that for any probability distribution μ on $E = \{1, 2\}$,

$$\mu^T \mathbf{P}(m, 2k+1) = \left(1 - \frac{1}{2k+1}, \frac{1}{2k+1}\right), \quad \mu^T \mathbf{P}(m, 2k) = \left(\frac{1}{2k}, 1 - \frac{1}{2k}\right),$$

and therefore, for all $k \geq m$,

$$\mu^T \mathbf{P}(m, k) - \nu^T \mathbf{P}(m, k) = 0.$$

Thus, the chain is weakly ergodic. But it cannot be strongly ergodic, since $\mu^T \mathbf{P}(m, k)$ has, as $k \rightarrow \infty$, two limit vectors: $(1, 0)$ and $(0, 1)$. \diamond

Remark 8.1. If the finite state chain is homogeneous and ergodic, then it is strongly ergodic in the sense of Definition 8.2, by the theorem of convergence to steady state for ergodic HMCs. \square

8.2 Block Criterion of Weak Ergodicity

As one might guess, weak ergodicity is in general not easy to check directly from the definition. Fortunately, there is a somewhat useable criterion in terms of Dobrushin's coefficient of ergodicity.

Theorem 8.1. Basic Weak Ergodicity Criterion

The chain is weakly ergodic if and only if for all $m \geq 0$,

$$\lim_{k \uparrow \infty} \delta(\mathbf{P}(m, k)) = 0. \quad (8.7)$$

Proof. Using Corollary 1.1 and the observation $d_V(\mu, \nu) \leq 1$,

$$d_V(\mu^T \mathbf{P}(m, k), \nu^T \mathbf{P}(m, k)) \leq d_V(\mu, \nu) \delta(\mathbf{P}(m, k)) \leq \delta(\mathbf{P}(m, k)).$$

Therefore (8.7) implies weak ergodicity. Now we have

$$\begin{aligned} \delta(\mathbf{P}(m, k)) &= \frac{1}{2} \sup_{i, j \in E} \sum_{\ell \in E} |p_{i\ell}(m, k) - p_{j\ell}(m, k)| \\ &\leq \frac{1}{2} \sup_{\mu, \nu} |\mu^T \mathbf{P}(m, k) - \nu^T \mathbf{P}(m, k)|. \end{aligned}$$

It follows that weak ergodicity implies (8.7). \square

By Dobrushin's inequality, $\delta(\mathbf{P}(m, k)) \leq \prod_{r=m}^{k-1} \delta(\mathbf{P}(r))$, and therefore nullity of the infinite product $\prod_{r \geq 1} \delta(\mathbf{P}(r))$ is enough to guarantee weak ergodicity. However, in many applications, weak ergodicity occurs without the above infinite product diverging to zero. It turns out that the consideration of blocks gives a useful necessary and sufficient condition due to Hajnal(1958).

Theorem 8.2. Block Criterion of Weak Ergodicity

The chain is weakly ergodic if and only if there exists a strictly increasing sequence of integers $\{n_s\}_{s \geq 0}$ such that

$$\sum_{s=0}^{\infty} (1 - \delta(\mathbf{P}(n_s, n_{s+1}))) = \infty. \quad (8.8)$$

Proof. Since $0 \leq \delta(\mathbf{P}(n_s, n_{s+1})) \leq 1$, (8.8) is equivalent to nullity of the infinite product $\prod_{s \geq 0} \delta(\mathbf{P}(n_s, n_{s+1}))$. But denoting by i the first integer s such that $n_s \geq m$, and by j the

last integer s such that $n_s \leq k - 1$, Dobrushin's inequality gives

$$\begin{aligned} \delta(\mathbf{P}(m, k)) &\leq \delta(\mathbf{P}(m, n_i)) \left\{ \prod_{s=i}^{j-1} \delta(\mathbf{P}(n_s, n_{s+1})) \right\} \delta(\mathbf{P}(n_j, k)) \\ &\leq \prod_{s=i}^{j-1} \delta(\mathbf{P}(n_s, n_{s+1})). \end{aligned}$$

Since $j \rightarrow \infty$ as $k \rightarrow \infty$, we see that (8.8) implies weak ergodicity, by Theorem 8.1.

Conversely, if we suppose weak ergodicity, then by Theorem 8.1 we can inductively construct for any $\gamma \in (0, 1)$ a strictly increasing sequence of integers $\{n_s\}_{s \geq 0}$ by

$$n_0 = 1, n_{s+1} = \inf \{k > n_s; \delta(\mathbf{P}(n_s, k)) \leq 1 - \gamma\}.$$

For such sequences, the product $\prod_{s \geq 0} \delta(\mathbf{P}(n_s, n_{s+1}))$ is null, and this is equivalent to (8.8). \square

8.3 Sufficient Condition of Strong Ergodicity

Knowing that a given nonhomogeneous Markov chain is weakly ergodic, it remains to decide whether it is strongly ergodic. No useful criterion of strong ergodicity is available, and we have to resort to sufficient conditions.

Theorem 8.3. *Sufficient Condition of Strong Ergodicity*

Suppose the chain is weakly ergodic, and assume that for all $n \geq 0$, there exists a probability distribution $\pi(n)$ on E such that

$$\pi^T(n) = \pi^T(n)\mathbf{P}(n) \tag{8.9}$$

and

$$\sum_{n=0}^{\infty} |\pi(n+1) - \pi(n)| < \infty. \tag{8.10}$$

The chain is then strongly ergodic.

Proof. Condition (8.10) implies the existence of a probability distribution π such that

$$\lim_{n \uparrow \infty} |\pi(n) - \pi| = 0. \tag{8.11}$$

Define Π_n (resp., Π) to be the matrix with all rows equal to $\pi(n)$ (resp., π). In particular, $|\Pi_n - \Pi| = |\pi(n) - \pi|$ and similarly $|\Pi_{n+1} - \Pi_n| = |\pi(n+1) - \pi(n)|$.

Also, for any probability distribution μ on E , $\mu^T \Pi = \pi^T$, and therefore (8.6) is equivalent to

$$\limsup_{k \uparrow \infty} \liminf_{\mu} |\mu^T (\mathbf{P}(m, k) - \Pi)| = 0,$$

which is in turn implied by

$$\lim_{k \uparrow \infty} |\mathbf{P}(m, k) - \Pi| = 0. \quad (8.12)$$

We therefore proceed to the proof of (8.12), writing for ℓ between m and k ,

$$\begin{aligned} \mathbf{P}(m, k) - \Pi &= \mathbf{P}(m, \ell)\mathbf{P}(\ell, k) - \Pi_{\ell+1}\mathbf{P}(\ell, k) \\ &\quad + \Pi_{\ell+1}\mathbf{P}(\ell, k) - \Pi_k + \Pi_k - \Pi. \end{aligned}$$

The triangle inequality for matrix norms gives

$$\begin{aligned} |\mathbf{P}(m, k) - \Pi| &\leq |\mathbf{P}(m, \ell)\mathbf{P}(\ell, k) - \Pi_{\ell}\mathbf{P}(\ell, k)| \\ &\quad + |\Pi_{\ell}\mathbf{P}(\ell, k) - \Pi_{k-1}| + |\Pi_{k-1} - \Pi| = A + B + C. \end{aligned} \quad (8.13)$$

An upper bound for A is given by Corollary 7.1:

$$A \leq |\mathbf{P}(m, \ell) - \Pi_{\ell}|\delta(\mathbf{P}(\ell, k)) \leq 2\delta(\mathbf{P}(\ell, k)), \quad (8.14)$$

where the last inequality follows from the definition of the matrix norm used here.

In order to bound B from above, we first observe that $\Pi_{\ell}\mathbf{P}(\ell) = \Pi_{\ell}$ and therefore

$$\Pi_{\ell}\mathbf{P}(\ell, k) = (\Pi_{\ell} - \Pi_{\ell+1})\mathbf{P}(\ell + 1, k) + \Pi_{\ell+1}\mathbf{P}(\ell + 1, k).$$

Iterating the process, we obtain

$$\Pi_{\ell}\mathbf{P}(\ell, k) = \sum_{j=\ell+1}^{k-1} (\Pi_{j-1} - \Pi_j)\mathbf{P}(j, k) + \Pi_{k-1},$$

and therefore

$$B \leq \sum_{j=\ell+1}^{k-1} |\Pi_{j-1} - \Pi_j|\delta(\mathbf{P}(j, k)) \leq \sum_{j=\ell+1}^{k-1} |\pi(j-1) - \pi(j)|, \quad (8.15)$$

where we have used the triangle inequality, Corollary 1.1, (8.10), and the observation $|\Pi_{j-1} - \Pi_j| = |\pi(j-1) - \pi(j)|$. As for C , we have, using the last observation,

$$C = |\pi(k-1) - \pi|. \quad (8.16)$$

The rest of the proof is now clear: For a given $\epsilon > 0$, fix ℓ such that $B \leq \frac{\epsilon}{3}$ for all $k \geq \ell$ (use (8.10)), and take k large enough so that $A \leq \frac{\epsilon}{3}$ (use Dobrushin's inequality) and $C \leq \frac{\epsilon}{3}$ (use (8.11)). \square

Remark 8.2. It is *not* required that $\mathbf{P}(n)$ be an ergodic stochastic matrix, or that $\pi(n)$ be a unique stationary probability of $\mathbf{P}(n)$.

The question is: How useful is Theorem 8.1? It seems that in order to satisfy (8.10), one has to obtain a closed-form expression for $\pi(n)$, or at least sufficient information about $\pi(n)$. How much information? It turns out that very little is needed in practice. More

precisely, a qualitative property of $\{\pi(n)\}_{n \geq 0}$ in terms of bounded variation extensions (to be defined) is sufficient to guarantee (8.10), and therefore strong ergodicity, if the chain is weakly ergodic.

The results below are due to Anily and Federgruen (1987). We first recall a definition:

A function $f : (0, 1] \rightarrow \mathbb{R}$ is said to be of *bounded variation* (BV) if

$$\sup \left\{ \sum_{i=1}^{\infty} |f(x_i) - f(x_{i-1})|; 0 < x_i < \dots < x_1 = 1 \text{ and } \lim_{i \rightarrow \infty} x_i = 0 \right\} < \infty. \quad (8.17)$$

A vector function $\mu : (0, 1] \rightarrow \mathbb{R}^E$ (or equivalently a vector $\mu(c) = \{\mu_i(c)\}_{i \in E}$) is said to be of bounded variation if

$$\sup \left\{ \sum_{i=1}^{\infty} |\mu(x_i) - \mu(x_{i-1})|; 0 < x_i < \dots < x_1 = 1 \text{ and } \lim_{i \rightarrow \infty} x_i = 0 \right\} < \infty. \quad (8.18)$$

◇

Definition 8.3. *Bounded Variation Extension*

The vector function $\bar{\pi} : (0, 1] \rightarrow \mathbb{R}^E$ is called a *bounded variation extension* of $\{\pi(n)\}_{n > 0}$ if there exists a sequence $\{c_n\}_{n \geq 0}$ in $(0, 1]$ decreasing to 0 and such that $\bar{\pi}(c_n) = \pi(n)$ for all $n \geq 0$.

Theorem 8.4. *Bounded Variation and Strong Ergodicity*

Suppose that $\{\mathbf{P}(n)\}_{n \geq 0}$ is weakly ergodic and that for all $n \geq 0$, there exists a probability vector $\pi(n)$ such that $\pi(n)\mathbf{P}(n) = \pi(n)$. If there exists a bounded variation extension $\bar{\pi}(c)$ of $\{\pi(n)\}_{n \geq 0}$, the chain is strongly ergodic.

Proof. We have

$$\sum_{n \geq 0} |\pi(n+1) - \pi(n)| = \sum_{n \geq 0} |\bar{\pi}(c_{n+1}) - \bar{\pi}(c_n)| < \infty,$$

since $\bar{\pi}(c)$ is an extension of $\{\pi(n)\}_{n \geq 0}$ and of bounded variation, and the conclusion follows by Theorem 8.1. □

Let $\bar{\mathbf{P}}(c)$ be an extension of $\{\mathbf{P}(n)\}_{n \geq 0}$; i.e., there exists a sequence $\{c_n\}_{n \geq 0}$ in $(0, 1]$, decreasing to 0 as n goes to infinity and such that for all $n \geq 0$,

$$\bar{\mathbf{P}}(c_n) = \mathbf{P}(n). \quad (8.19)$$

Suppose that for each $c \in (0, 1]$, there exists a probability vector $\bar{\pi}(c)$ such that

$$\bar{\pi}(c)\bar{\mathbf{P}}(c) = \bar{\pi}(c). \quad (8.20)$$

Is it enough for $\bar{\pi}(c)$ to be of bounded variation that $\bar{\mathbf{P}}(c)$ be of bounded variation, i.e., that

$$\sup \left\{ \sum_{i=1}^{\infty} |\bar{\mathbf{P}}(x_{i+1}) - \bar{\mathbf{P}}(x_i)|; 0 < x_i < \dots < x_1 = 1 \text{ and } \lim_{i \in \infty} x_i = 0 \right\} < \infty?$$

A simple counterexample shows that this is not the case.

Example 8.2. Counterexample

Let

$$P(n) = \begin{pmatrix} 1 - e^{-n} & e^{-n} \\ e^{-n} \sin^2\left(\frac{n\pi}{2}\right) & 1 - e^{-n} \sin^2\left(\frac{n\pi}{2}\right) \end{pmatrix},$$

$$\bar{\mathbf{P}}(c) = \begin{pmatrix} 1 - e^{-1/c} & e^{-1/c} \\ e^{-1/c} \sin^2\left(\frac{\pi}{2c}\right) & 1 - e^{-1/c} \sin^2\left(\frac{\pi}{2c}\right) \end{pmatrix}.$$

Clearly, $\bar{\mathbf{P}}(c)$ is a bounded variation extension of $\{\mathbf{P}(n)\}_{n \geq 0}$. If the corresponding stationary probability $\bar{\pi}(c)$ were of bounded variation, then as shown in the proof of Theorem 6.5, $\sum_{n \geq 0} |\pi(n+1) - \pi(n)|$ would be finite. Computations give for the second coordinate of $\pi(n)$

$$\pi(n)_2 = \left(1 + \sin^2\left(\frac{n\pi}{2}\right) \right)^{-1},$$

a quantity that oscillates between 1 and $\frac{1}{2}$. Therefore, $\sum_{n \geq 0} |\pi(n+1) - \pi(n)| = \infty$. \diamond

In order to give conditions on $\bar{\mathbf{P}}(c)$ ensuring that $\bar{\pi}(c)$ is of bounded variation, we shall first give a precise description of $\bar{\pi}(c)$ in terms of the entries of $\bar{\mathbf{P}}(c)$. This can be done in the case where E is finite, identified with $\{1, \dots, N\}$ for simplicity. Indeed $\bar{\pi}(c)$ is a solution of the balance equations

$$\bar{\pi}(c)_i = \sum_{j=1}^N \bar{\mathbf{P}}(c)_{ji} \bar{\pi}(c)_j,$$

for $i \in [1, N - 1]$, together with the normalizing equation

$$\sum_{i=1}^N \bar{\pi}(c)_i = 1.$$

That is, in matrix form,

$$\bar{\pi}(c) \begin{bmatrix} 1 - \bar{\mathbf{P}}(c) & \dots & -\bar{\mathbf{P}}(c)_{N-1,1} & -\bar{\mathbf{P}}(c)_{N,1} \\ -\bar{\mathbf{P}}(c) & \dots & -\bar{\mathbf{P}}(c)_{N-1,2} & -\bar{\mathbf{P}}(c)_{N,2} \\ \vdots & & \vdots & \vdots \\ -\bar{\mathbf{P}}(c)_{1,N-1} & \dots & 1 - \bar{\mathbf{P}}(c)_{N-1,N-1} & -\bar{\mathbf{P}}(c)_1 \\ \vdots & & \vdots & \vdots \\ 1 & & 1 & 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}. \tag{8.21}$$

Cramer's rule gives

$$\bar{\pi}(c)_i = \frac{A_i(c)}{B_i(c)},$$

where $A_i(c)$ and $B_i(c)$ are finite sums and differences of finite products of the entries of $\bar{\mathbf{P}}(c)$.

Example 8.3. *Rational Polynomial–Exponential Transition Matrices*

This case covers most applications. The elements of $\bar{\mathbf{P}}(c)$ are ratios of functions of the type

$$\sum_{j=1}^n Q_j \left(\frac{1}{c} \right) e^{\lambda_j/c}, \quad (8.22)$$

where the Q_j are polynomial functions and the λ_j are real numbers. Then so are the elements of $\bar{\pi}(c)$, as well as their derivatives with respect to c . But ratios of terms of type (8.22) have for sufficiently small $c > 0$ a constant sign. Therefore a given element $\bar{\pi}(c)_i = \psi(c)$ is such that

(α) for some $c^* > 0$, $\psi : (0, c^*] \rightarrow \mathbb{R}$ is monotone and bounded;

(β) $\psi : (0, 1] \rightarrow \mathbb{R}$ is continuously differentiable.

Properties (α) and (β) are sufficient to guarantee that $\psi : (0, 1] \rightarrow \mathbb{R}$ is of bounded variation. \diamond

We have spent some time explaining how the sufficient condition of strong ergodicity (8.10) can be checked. One may wonder whether this is really worthwhile, and whether a weaker and easier to verify condition is available. A natural candidate for a sufficient condition of strong ergodicity, *given weak ergodicity*, is

$$\lim_{n \uparrow \infty} |\pi(n) - \pi| = 0, \quad (8.23)$$

for some probability π .

Example 8.4. *Counterexample*

Define for all $n \geq 1$,

$$\mathbf{P}(2n-1) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \mathbf{P}(2n) = \begin{pmatrix} 0 & 1 \\ 1 - \frac{1}{2n} & \frac{1}{2n} \end{pmatrix}.$$

The sequence $\{\mathbf{P}(n)\}_{n \geq 1}$ is weakly ergodic (Problem 6.8.2). The corresponding stationary distributions are

$$\pi(2n-1) = \left(\frac{1}{2}, \frac{1}{2} \right), \quad \pi(2n) = \left(\frac{2n-1}{4n-1}, \frac{2n}{4n-1} \right),$$

and therefore (8.23) is satisfied with $\pi = \left(\frac{1}{2}, \frac{1}{2} \right)$. On the other hand, if we define for all $k \geq 1$

$$R(k) = \mathbf{P}(2k)\mathbf{P}(2k+1) = \begin{pmatrix} 1 - \frac{1}{2k} & \frac{1}{2k} \\ 0 & 1 \end{pmatrix}$$

and

$$S(k) = \mathbf{P}(2k-1)\mathbf{P}(2k) = \begin{pmatrix} 1 & 0 \\ \frac{1}{2k} & 1 - \frac{1}{2k} \end{pmatrix},$$

then the sequences $\{R(k)\}_{k \geq 1}$ and $\{S(k)\}_{k \geq 1}$ are weakly ergodic (exercise), and their stationary distributions are constant, equal to $(1, 0)$ and $(0, 1)$, respectively. Therefore, by Theorem 3.1, they are strongly ergodic, and in particular,

$$\lim_{k \uparrow} \mathbf{P}(1)\mathbf{P}(2) \cdots \mathbf{P}(2k-1)\mathbf{P}(2k) = \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix}$$

and

$$\lim_{k \uparrow} \mathbf{P}(1)(\mathbf{P}(2)\mathbf{P}(3) \cdots \mathbf{P}(2k)\mathbf{P}(2k+1)) = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix}.$$

Therefore, the sequence $\{\mathbf{P}(n)\}_{n \geq 1}$ is not strongly ergodic. \diamond

Problem 6.8.4 shows what can be added to (8.23) to obtain a sufficient condition of strong ergodicity.

We shall quote without proof the following natural result; see Theorem V.4.5 of (Isaacson and Madsen, 1976).

Theorem 8.5.

Let $\{\mathbf{P}(n)\}_{n \geq 1}$ be a sequence of transition matrices each having at least one stationary distribution, and such that

$$\lim_{n \uparrow \infty} |\mathbf{P}(n) - \mathbf{P}| = 0 \tag{8.24}$$

for some *ergodic* transition matrix \mathbf{P} . This sequence is strongly ergodic.

Note the requirement that \mathbf{P} be ergodic. This condition will be found too strong in the study of the convergence of simulated annealing algorithms, where typically the limit transition matrix is reducible (see Chapter 7).

Problems

6.1.1 For the homogeneous Markov chain with state space $E = \{1, 2, 3\}$ and transition matrix

$$\mathbf{P} = \begin{pmatrix} 1 - \alpha & \alpha & 0 \\ 0 & 1 - \beta & \beta \\ \gamma & 0 & 1 - \gamma \end{pmatrix},$$

where $\alpha, \beta, \gamma \in (0, 1)$, compute $\lim_{n \uparrow \infty} \mathbf{P}^n$ and give the corresponding rate of convergence.

6.1.2 With the notation of the Section 1.2 on quasi-stationary distributions, show that for any transient state j ,

$$\lim_{m \uparrow \infty, n \uparrow \infty} P(X_n = j | v > m + n) = \frac{u_1(j)v_1(j)}{\sum_{i \in T} u_1(i)v_1(i)}.$$

6.1.3 Let \mathbf{P} be a stochastic matrix on the finite state space E , and let T and R be the sets of transient states and recurrent states, respectively. Suppose, moreover, that the recurrent classes consist of only one state; that is, the block decomposition of the transition matrix with respect to the partition $R + T = E$ is

$$\mathbf{P} = \begin{pmatrix} I & 0 \\ B & Q \end{pmatrix},$$

where I is the $|R| \times |R|$ identity matrix. Show that \mathbf{P} is diagonalizable if and only if \mathbf{Q} is diagonalizable.

6.1.4 Show that for a finite-state ergodic HMC with *diagonalizable* transition matrix, the existence of a finite time k such that irrespective of the initial distribution X_k has the stationary distribution π is equivalent to the fact that $\{X_n\}_{n \geq 0}$ is an i.i.d. sequence. Give a simple counter-example when the hypothesis of diagonalizability of the transition matrix is abandoned.

6.3.1 Show that

$$\chi^2(p_i(n); \pi(\cdot)) = \sum_{j=2}^r \lambda_j^{2n} v_j(i)^2,$$

where v_j is the j th righteigenvector associated with the reversible ergodic pair (\mathbf{P}, π) , and λ_j is the corresponding eigenvalue.

6.3.2 Show that λ_2 is equal to the maximum of the correlation coefficient between $f(X_0)$ and $f(X_1)$ among all real-valued functions f such that $E[f(X_0)] = 0$, where $\{X_n\}_{n \geq 0}$ is a stationary HMC corresponding to (\mathbf{P}, π) , and λ_2 is the SLE of \mathbf{P} .

6.4.1 Prove the version of Theorem 4.1 where Poincaré's coefficient κ is replaced by

$$K = \max_e Q(e)^{-1} \sum_{\gamma_j, e \in \gamma_j} |\gamma_{ij}| \pi(i) \pi(j),$$

where $|\gamma|$ is the length of path γ . (This version is due to Sinclair, and Jim Fill pointed out to me that it generally does better than that of Theorem 4.1, when the two versions differ).

6.4.2 Consider a random walk on a graph (see Exaple 6.3 of Chapter 2), where the graph is now a full binary tree of depth L . Show that the second largest eigenvalue λ_2 satisfies

$$\lambda_2 \leq 1 - \frac{1}{9L2^{L-1}}.$$

6.4.3 This is a continuation of the previous problem. Explain why the formula (4.10) does not apply directly. Show that

$$\lambda_2 \geq 1 - 2(2^L - 1) \left(1 - \frac{1}{2^{L+1} - 2} \right),$$

which is equivalent for large L to $1 - 2^{-L-1}$. Hint: Use Rayleigh's characterization with x as follows: $x(i) = 0, 1$, or -1 according to whether i is the root, a vertex on the right of the tree, or one on the left.

6.5.1 Let $M_p(E)$ be a collection of probability distributions on the countable set E . Show that for all $\alpha, \beta \in M_p(E)$,

$$s(\alpha; \beta) = \inf \left\{ s \geq 0; \alpha = (1-s)\beta + s\gamma, \gamma \in M_p(E) \right\},$$

where s denotes the separation pseudo-distance.

6.5.2 Show that condition (α) of Definition 5.3 is equivalent to either one of the following two conditions:

(β) For all $i \in E$ and all $n \geq 0$,

$$P(X_n = i | T = n) = \pi(i).$$

(γ) For all $i \in E$ and all $n \geq 0$,

$$P(X_n = i | T \leq n) = \pi(i).$$

6.6.1 Let π be the stationary distribution of an ergodic Markov chain with finite state space, and denote by T_i the return time to state i . Let S_Z be the time necessary to visit for the first time the random state Z chosen according to the distribution π , independently of the chain. Show that $E_i[S_Z]$ is independent of i , and give its expression in terms of the fundamental matrix.

6.6.2 Show that $\max \left\{ \frac{v(f, \mathbf{P}, \pi)}{\text{Var}_\pi(f)}; f \neq 0 \right\} = \frac{1+\lambda_2}{1-\lambda_2}$. Recall the definition of $v(f, \mathbf{P}, \pi)$: it is the left-hand side of (6.1.18), that is, the asymptotic variance of the ergodic estimate of $E_\pi[f(X_0)]$.

6.7.1. Let \mathbf{P} be a transition matrix on E . Show that $\delta(\mathbf{P}) = 0$ iff the rows of \mathbf{P} are identical, and that $\delta(\mathbf{P}) = 1$ iff there exist two rows of \mathbf{P} that are orthogonal.

6.7.2. Let \mathbf{P} be an ergodic transition matrix on the finite set E . Show that its SLEM is bounded above by $\inf_k \delta(\mathbf{P}^k)^{1/k}$.

6.7.3. Let \mathbf{P} be an ergodic transition matrix on the finite state space E . Suppose, moreover, that \mathbf{P} is diagonalizable (for simplicity, since the result is the same for the general case). Show that

$$\ln(\text{SLEM}(\mathbf{P})) = \lim_{k \rightarrow \infty} \frac{1}{k} \ln(\delta(\mathbf{P}^k)).$$

6.7.4. Prove in detail (7.10).

6.7.5. Let \mathbf{P} and \mathbf{Q} be two ergodic transition matrices on the finite state space E , with respective stationary distributions π and σ . Let $f : E \rightarrow \mathbb{R}$ be such that $\|f\|_\infty = \sup_{i \in E} |f(i)|$. Show that

$$|\langle f \rangle_\pi - \langle f \rangle_\sigma| \leq \|f\|_\infty \frac{\alpha}{1 + \alpha - \beta},$$

where $\alpha = \alpha(\mathbf{P}, \mathbf{Q})$ and $\beta = \beta(\mathbf{P}, \mathbf{Q})$ are defined by (7.8) and (7.9), respectively. Show that

$$d_V(\pi, \sigma) \leq \frac{\alpha}{1 + \alpha - \beta}.$$

6.7.6. Show that for two transition matrices \mathbf{P} and \mathbf{Q} on the same space,

$$|\delta(\mathbf{P}) - \delta(\mathbf{Q})| \leq \|\mathbf{P} - \mathbf{Q}\|.$$

6.8.1 For $n \geq 0$, define

$$\mathbf{P}(2n-1) = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & 0 \end{pmatrix}, \quad \mathbf{P}(2n) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Show that the nonhomogeneous Markov chain with transition matrices $\{\mathbf{P}(n)\}_{n \geq 0}$ is *not* strongly ergodic (Hint: take $\mu = (0, 1)$ as initial distribution, and compute the distributions of the chain at even and odd times). Show that the chain is weakly ergodic.

6.8.2 For $n \geq 0$, define

$$\mathbf{P}(2n-1) = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}, \quad \mathbf{P}(2n) = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{6} & \frac{1}{3} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{6} \end{pmatrix}.$$

Show directly that the nonhomogeneous Markov chain with transition matrices $\{\mathbf{P}(n)\}_{n \geq 0}$ is strongly ergodic.

6.8.3 Show that the sequence $\{\mathbf{P}(n)\}_{n \geq 0}$ defined by

$$\mathbf{P}(n) = \begin{pmatrix} \frac{1}{3} + \frac{1}{n} & \frac{2}{3} - \frac{1}{n} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$$

is strongly ergodic.

6.8.4 (Isaacson and Madsen, 1976) Let $\{\mathbf{P}(n)\}_{n \geq 0}$ be a sequence of transition matrices such that for some $D < \infty$ and for all $k \geq 1$,

$$\sum_{j=0}^{k-1} \delta(\mathbf{P}(j, k)) \leq D.$$

Show that this sequence is weakly ergodic.

Show that it is strongly ergodic if in addition there exists a probability vector π such that (8.23) holds, where $\pi(n)$ is a stationary distribution, assumed to exist but not necessarily unique, of $\mathbf{P}(n)$. (Hint: Mimic the proof of Theorem 6.1, only adapting the argument for the term B thereof.)

Show that

$$\mathbf{P}(n) = \begin{pmatrix} \frac{1}{n} & 1 - \frac{1}{n} \\ 1 & 0 \end{pmatrix}$$

satisfies the conditions of Theorem 8.1, but not those of the theorem just proved.

Gibbs Fields and Monte Carlo Simulation

1 Markov Random Fields

1.1 Neighborhoods and Local Specification

The Markov property of a stochastic sequence $\{X_n\}_{n \geq 0}$ implies that for all $n \geq 1$, X_n is independent of $(X_k, k \notin \{n-1, n, n+1\})$ given (X_{n-1}, X_{n+1}) .

If we call n a site, X_n the value of the process at site n , and the set $\{n-1, n+1\}$ the neighborhood of site n , the above property can be rephrased as: For all $n \geq 1$, the value at site n is independent of the values at sites $k \notin \{n-1, n, n+1\}$ given the values in the neighborhood of site n . This is the starting point for a natural generalization of the Markov property.

Definition 1.1. *Random Field*

Let S be a *finite* set, with elements denoted by s and called *sites*, and let Λ be a finite set called the *phase space*. A *random field* on S with phases in Λ is a collection $X = \{X(s)\}_{s \in S}$ of random variables $X(s)$ with values in Λ .

A random field can be regarded as a random variable taking its values in the *configuration space* Λ^S . A configuration $x \in \Lambda^S$ is of the form $x = (x(s), s \in S)$, where $x(s) \in \Lambda$ for all $s \in S$. For a given configuration x and a given subset $A \subset S$, define

$$x(A) = (x(s), s \in A),$$

the restriction of x to A . If $S \setminus A$ denotes the complement of A in S , one writes $x = (x(A), x(S \setminus A))$. In particular, for fixed $s \in S$, $x = (x(s), x(S \setminus s))$, where $S \setminus s$ is a shorter way of writing $S \setminus \{s\}$.

Of special interest are the Markov fields characterized by *local interactions*. In order to state the Markov property for random fields, we need first to introduce a *topology* on the sites.

Definition 1.2. *Neighborhoods*

A *neighborhood system* on S is a family $N = \{\mathcal{N}_s\}_{s \in S}$ of subsets of S such that for all $s \in S$,

- (i) $s \notin \mathcal{N}_s$,
- (ii) $t \in \mathcal{N}_s \Rightarrow s \in \mathcal{N}_t$.

The subset \mathcal{N}_s is called the *neighborhood* of site s . The couple (S, N) is called a *graph*, or a *topology*. The *boundary* of $A \subset S$ is, by definition, the set $\partial A = (\cup_{s \in A} \mathcal{N}_s) \setminus A$.

In the graph interpretation, S is the set of *vertices* and N defines the *edges*: Sites s and t are linked by an edge if and only if they are neighbors, i.e., $t \in \mathcal{N}_s$.

In the following, $\tilde{\mathcal{N}}_s$ denotes the set $\mathcal{N}_s \cup \{s\}$.

Definition 1.3. *Markov Random Field*

The random field X is called a *Markov random field* (MRF) with respect to the neighborhood system N if for all sites $s \in S$ the random variables $X(s)$ and $X(S \setminus \tilde{\mathcal{N}}_s)$ are independent given $X(\mathcal{N}_s)$.

In mathematical symbols:

$$P(X(s) = x(s) \mid X(S \setminus s) = x(S \setminus s)) = P(X(s) = x(s) \mid X(\mathcal{N}_s) = x(\mathcal{N}_s)) \quad (1.1)$$

for all $s \in S$, $x \in \Lambda^S$. Property (1.1) is clearly of the Markov type: the distribution of the phase at a site is directly influenced only by the phases of the neighboring sites.

Note that any random field is Markovian with respect to the trivial topology, where the neighborhood of any site is the whole set S . However, as will become evident, the interesting Markov fields (from the point of view of modeling, simulation, and optimization) are those with relatively small neighborhoods.

Definition 1.4. *Local Specification*

The *local characteristic* of the MRF at site s is the function $\pi^s : \Lambda^S \rightarrow [0, 1]$ defined by

$$\pi^s(x) = P(X(s) = x(s) \mid X(\mathcal{N}_s) = x(\mathcal{N}_s)). \quad (1.2)$$

The family $\{\pi^s\}_{s \in S}$ is called the *local specification* of the MRF.

One sometimes writes

$$\pi^s(x) = \pi(x(s) | x(\mathcal{N}'_s))$$

in order to stress the role of the neighborhood system.

Finiteness of the phase space Λ is so far not essential. An immediate extension not requiring a change in the notation is to a denumerable phase space. Arbitrary phase spaces do not present a difficulty except for notation, since $\pi^s(x) = \pi(x(s) | x(\mathcal{N}'_s))$, considered as a function of $x(s)$ for a fixed $x(\mathcal{N}'_s)$, is now interpreted as a density function with respect to some measure. A simple case that contains all that is needed in applications will be briefly outlined.

Here the phase space is $\Lambda = F \times \mathbb{R}^k$, where F is a denumerable set. Thus $x = (u, v)$, where $u = (u(s), s \in S)$, $v = (v(s), s \in S)$, $u(s) \in F$, $v(s) \in \mathbb{R}^k$. In particular, $X = (U, V) = \{(U(s), V(s))\}_{s \in S}$. The local characteristic π^s is now interpreted according to

$$P(U(s) = \alpha, V(s) \leq \beta | U(\mathcal{N}'_s) = u(\mathcal{N}'_s), V(\mathcal{N}'_s) = v(\mathcal{N}'_s)) = \int_{-\infty}^{\beta} \pi(\alpha, \gamma | u(\mathcal{N}'_s), v(\mathcal{N}'_s)) d\gamma,$$

where the integral is, of course, a multiple integral of order k . Thus $\pi(\alpha, \gamma | u(\mathcal{N}'_s), v(\mathcal{N}'_s))$ is a probability density with respect to γ for fixed α . The Markov property now reads

$$P(U(s) = \alpha, V(s) \leq \beta | U(S \setminus s), V(S \setminus s)) = P(U(s) = \alpha, V(s) \leq \beta | U(\mathcal{N}'_s), V(\mathcal{N}'_s))$$

for all $\alpha \in F$, $\beta \in \mathbb{R}^k$, $s \in S$.

Definition 1.5. *Positivity Condition*

The probability distribution π on the finite configuration space Λ^S , where $S = \{1, 2, \dots, K\}$, is said to satisfy the *positivity condition* if for all $j \in S$, $x_j \in \Lambda$,

$$(\pi_j(x_j) = 0) \Rightarrow (\pi(y_1, \dots, y_{j-1}, x_j, y_{j+1}, \dots, y_K) = 0) \tag{1.3}$$

for all $y_1, \dots, y_{j-1}, y_{j+1}, \dots, y_K \in \Lambda$, where π_j is the marginal distribution on site j .

Theorem 1.1. *At Most One Distribution for a Local Specification*

Two distributions of an MRF with a *finite* configuration space Λ^S that satisfy the positivity condition and have the same local specification are identical.

Proof. Enumerate S as in Definition 1.5. The result follows from the identity

$$\pi(x) = \prod_{i=1}^K \frac{\pi(x_i | x_1, \dots, x_{i-1}, y_{i+1}, \dots, y_K)}{\pi(y_i | x_1, \dots, x_{i-1}, y_{i+1}, \dots, y_K)} \pi(y) \tag{1.4}$$

for any $x, y \in \Lambda^S$ with strictly positive probability. Indeed, if the local specification is specified and y is fixed, then π is determined up to the multiplicative factor $\pi(y)$. But the latter is in turn determined by normalization.

For the proof of the above identity, one checks that

$$\pi(x) = \frac{\pi(x_K | x_1, \dots, x_{K-1})}{\pi(y_K | x_1, \dots, x_{K-1})} \pi(x_1, \dots, x_{K-1}, y_K)$$

by simply rewriting the conditional probabilities thereof using Bayes's definition. Similarly,

$$\pi(x_1, \dots, x_{K-1}, y_K) = \frac{\pi(x_{K-1} | x_1, \dots, x_{K-2}, y_K)}{\pi(y_{K-1} | x_1, \dots, x_{K-2}, y_K)} \pi(x_1, \dots, x_{K-2}, y_{K-1}, y_K)$$

and so forth. The above calculations make sense because the positivity condition and the strict positivity of $\pi(x)$ and $\pi(y)$ imply that for all $j \in [1, K]$, $\pi(x_1, \dots, x_j, y_{j+1}, \dots, y_K) > 0$. \square

Note that Theorem 1.1 is not true in general for an infinite number of sites, and this is why the basic assumption on the finiteness of S was recalled. In the case where S is infinite, there may be several distributions corresponding to a given local specification: One then speaks of *phase transition* (Section 5 gives a brief introduction to this topic).

The positivity condition is also important (see Problem 7.1.3).

1.2 Cliques, Potential, and Gibbs Distributions

The basic notion in this subsection comes from physics, where it was introduced by Gibbs (1902). Consider the probability distribution

$$\pi_T(x) = \frac{1}{Z_T} e^{-\frac{1}{T} \mathcal{E}(x)} \quad (1.5)$$

on the configuration space Λ^S , where $T > 0$ is the *temperature*, $\mathcal{E}(x)$ is the *energy* of configuration x , and Z_T is the normalizing constant, called the *partition function*. Since $\pi_T(x)$ takes its values in $[0, 1]$, necessarily $-\infty < \mathcal{E}(x) \leq +\infty$. Note that $\mathcal{E}(x) < \infty$ if and only if $\pi_T(x) > 0$.

Such distributions are interesting for physicists when the energy is expressed in terms of a potential function describing the local interactions. The notion of *clique* then plays a central role. Remember that in ordinary parlance, a clique is a group of people who know (and favor) one another.

Definition 1.6. Cliques

Any singleton $\{s\}$ is a clique. A subset $C \subset S$ with more than one element is called a clique of the graph (S, N) if and only if any two distinct sites of C are mutual neighbors. A clique C is called *maximal* if for any site s , $C \cup \{s\}$ is *not* a clique.

Example 1.1. Ferromagnets

In statistical physics, the following model is regarded as a qualitatively correct idealization of a piece of ferromagnetic material. Here $S = \mathbb{Z}_m^2 = \{(i, j) \in \mathbb{Z}^2, i, j \in [1, m]\}$ and

$\Lambda = \{+1, -1\}$, where ± 1 is the orientation of the magnetic spin at a given site. Fig. 7.1.1 depicts two particular neighborhood systems, their respective cliques, and the boundary of a 2×2 square for both cases. \diamond




Figure 7.1.1. Two examples of neighborhoods, cliques, and boundaries

Definition 1.7. *Gibbs Potential and Gibbs Distribution*

A Gibbs potential on Λ^S relative to the neighborhood system N is a collection $\{V_C\}_{C \in \mathcal{C}_S}$ of functions $V_C : \Lambda^S \rightarrow \mathbb{R} \cup \{+\infty\}$ such that

- (i) $V_C \equiv 0$ if C is not a clique,
(ii) for all $x, x' \in \Lambda^S$ and all $C \subset S$,

$$(x(C) = x'(C)) \Rightarrow (V_C(x) = V_C(x')). \quad (1.6)$$

The energy function $\mathcal{E} : \Lambda^S \rightarrow \mathbb{R} \cup \{+\infty\}$ is said to *derive from the potential* $\{V_C\}_{C \subset S}$ if

$$\mathcal{E}(x) = \sum_C V_C(x). \quad (1.7)$$

The function V_C depends only on the phases at the sites inside subset C . One could write more explicitly $V_C(x(C))$ instead of $V_C(x)$, but this notation will not be used.

In this context, the distribution in (1.5) is called a *Gibbs distribution*.

Example 1.2. Ising Model

This model was introduced by Ising (1925) for understanding qualitatively the phenomenon of phase transition in ferromagnetic materials. In Ising's *finite* model, $S = \mathbb{Z}_m^2$, $\Lambda = \{+1, -1\}$, and the neighborhood system is as in (α) of Fig. 7.1.1. The Gibbs potential is

$$V_{\{s\}}(x) = -\frac{H}{k}x(s),$$

$$V_{\{s,t\}}(x) = -\frac{J}{k}x(t)x(s),$$

where $\langle s, t \rangle$ is the 2-element clique $\{s, t\}$, where $t \in \mathcal{N}_s$. Here, k is the Boltzmann constant, H is the external magnetic field, and J is the internal energy of an elementary magnetic dipole. The energy function corresponding to this potential is therefore

$$\mathcal{E}(x) = -\frac{J}{k} \sum_{\langle s,t \rangle} x(s)x(t) - \frac{H}{k} \sum_{s \in S} x(s). \quad \diamond$$

Example 1.3. Neural Network

In this example, the graph structure (S, N) is as in the Ising model, but now the phase space is $\Lambda = \{0, 1\}$. A site s is interpreted as being a *neuron* that is *excited* if $x(s) = 1$ and *inhibited* if $x(s) = 0$. If $t \in \mathcal{N}_s$, one says that there is a *synapse* from s to t , and such a synapse has a *strength* w_{st} . If $w_{st} > 0$, one says that the synapse is *excitatory*; otherwise it is called *inhibitory*. The energy function is

$$\mathcal{E}(x) = \sum_{s \in S} \sum_{t \in \mathcal{N}_s} w_{ts} x(t)x(s) - \sum_{s \in S} h_s x(s),$$

where h_s is called the *threshold* of neuron s (we shall understand why later). From the expression

$$\mathcal{E}(x) = \sum_{\langle s,t \rangle} (w_{ts} + w_{st})x(t)x(s) - \sum_{s \in S} h_s x(s),$$

we see that the energy derives from the Gibbs potential

$$V_{\{s\}}(x) = -h_s x(s),$$

$$V_{\{s,t\}}(x) = (w_{ts} + w_{st})x(t)x(s).$$

From a biological point of view, the above model is not interesting because the coefficient of $x(t)x(s)$ is symmetric in t, s , and this implies a symmetry between two interacting neurons that is not found in nature. However, this model is of interest to artificial intelligence and is then called an artificial neural network. \diamond

One of the challenges associated with Gibbs models is obtaining explicit formulas for averages, considering that it is generally hard to compute the partition function. However, physicists have gained some expertise in this domain with relatively simple models; see, for instance, (Baxter, 1982). Here is the simplest one, which Ising solved in 1925.

Example 1.4. *Ising's Toric Model*

This is the classical Ising model of Example 1.2, except that the site space $S = \{1, 2, \dots, N\}$ consists of N points arranged in this order on a circle. In particular, sites 1 and N are neighbors. The neighbors of site i are $i + 1$ and $i - 1$, with the convention that site $N + 1$ is site 1. The phase space is $\Lambda = \{+1, -1\}$. A configuration $x \in E = \Lambda^S$ is denoted by (x_1, x_2, \dots, x_N) . Therefore, with $K = \frac{J}{kT}$ and $h = \frac{H}{kT}$,

$$\pi_T(x) = \frac{1}{Z_N} e^{K \sum_{i=1}^N x_i x_{i+1} + h \sum_{i=1}^N x_i},$$

where the partition function is

$$Z_N = \sum_{x \in E} e^{K \sum_{i=1}^N x_i x_{i-1} + h \sum_{i=1}^N x_i}.$$

It takes the convenient form

$$Z_N = \sum_{x \in E} R(x_1, x_2)R(x_2, x_3) \dots R(x_N, x_1), \tag{1.8}$$

where

$$R(a, b) = e^{Kab + \frac{1}{2}h(a+b)}.$$

Defining the *transfer matrix*

$$R = \begin{pmatrix} R(+1, +1) & R(+1, -1) \\ R(-1, +1) & R(-1, -1) \end{pmatrix} = \begin{pmatrix} e^{K+h} & e^{-K} \\ e^{-K} & e^{K-h} \end{pmatrix},$$

we see that the general term of the sum in (1.8) is equal to

$$\sum_{x_1 \in \Lambda} \sum_{x_1, \dots, x_N \in \Lambda} R(x_1, x_2)R(x_2, x_3) \dots R(x_N, x_1) = \sum_{x_1 \in \Lambda} R^N(x_1, x_1).$$

Therefore, Z_N is the trace of R^N . In particular, if we denote by λ_+ and λ_- the eigenvalues of R , then

$$Z_N = (\lambda_+)^N + (\lambda_-)^N. \quad (1.9)$$

The reader is directed to Problem 7.5.3 for further study of this model. The *transfer matrix method* can be applied in principle to other translationally invariant toric models. For further examples of exactly solved models, which are all kin to the basic Ising model, see (Baxter, 1982). \diamond

2 Gibbs–Markov Equivalence

2.1 From the Potential to the Local Specification

Gibbs distributions with an energy deriving from a Gibbs potential relative to a neighborhood system are distributions of Markov fields relative to the same neighborhood system. This result admits a (partial) converse, given in the next subsection.

Theorem 2.1. *Gibbs Fields Are Markov Fields*

If X is a random field with a distribution π , of the form (1.5), where the energy $\mathcal{E}(x)$ derives from a Gibbs potential $\{V_C\}_{C \in \mathcal{S}}$ relative to the neighborhood system N , then X is Markovian relative to the same neighborhood system N . Moreover, its local specification is given by the formula

$$\pi^s(x) = \frac{e^{-\sum_{C \ni s} V_C(x)}}{\sum_{\lambda \in \Lambda} e^{-\sum_{C \ni s} V_C(\lambda, x(S \setminus s))}}, \quad (2.1)$$

where the notation $\sum_{C \ni s}$ means that the sum extends over the sets C that contain the site s .

Proof. First observe that the right-hand side of (2.1) depends on x only through $x(s)$ and $x(\mathcal{N}_s)$. Indeed, $V_C(x)$ depends only on $(x(t), t \in C)$, and for a clique C , if $t \in C$ and $s \in C$, then either $t = s$ or $t \in \mathcal{N}_s$.

Therefore, if one can show that $P(X(s) = x(s) | X(S \setminus s) = x(S \setminus s))$ equals the right-hand side of (2.1), and in particular is a function of $x(s)$ and $x(\mathcal{N}_s)$ only, then by Theorem 2.3 of Chapter 1, the Markov property (1.1) and equality (2.1) will be proven. By definition of conditional probability,

$$P(X(s) = x(s) | X(S \setminus s) = x(S \setminus s)) = \frac{\pi(x)}{\sum_{\lambda \in \Lambda} \pi(\lambda, x(S \setminus s))}. \quad (2.2)$$

But

$$\pi(x) = \frac{1}{Z} e^{-\sum_{C \ni s} V_C(x) + \sum_{C \ni s} V_C(x)},$$

and similarly,

$$\pi(\lambda, x(S \setminus s)) = \frac{1}{Z} e^{-\sum_{C \ni s} V_C(\lambda, x(S \setminus s)) - \sum_{C \ni s} V_C(\lambda, x(S \setminus s))}.$$

If C is a clique and s is not in C , then $V_C(\lambda, x(S \setminus s)) = V_C(x)$ and is therefore independent of $\lambda \in \Lambda$. Therefore, the righthand side of (2.2) is found, after factoring out $\exp \left\{ - \sum_{C \ni s} V_C(x) \right\}$, to be equal to the righthand side of (2.1). \square

The local energy at site s of configuration x is

$$\mathcal{E}_s(x) = \sum_{C \ni s} V_C(x).$$

With this notation, (2.1) becomes

$$\pi^s(x) = \frac{e^{-\mathcal{E}_s(x)}}{\sum_{\lambda \in \Lambda} e^{-\mathcal{E}_s(\lambda, x(S \setminus s))}}.$$

Example 2.1. Neural Network

This is a continuation of Example 1.3. One finds by application of (2.1), and recalling that $\Lambda = \{0, 1\}$, that

$$\pi_T^s(x) = \frac{e^{-\frac{1}{T} \{ \sum_{t \in N_s^+} (w_{ts} + w_{st}) x(t) - h_s \} x(s)}}{1 + e^{-\frac{1}{T} \{ \sum_{t \in N_s^+} (w_{ts} + w_{st}) x(t) - h_s \}}}. \tag{2.3}$$

\diamond

Example 2.2. Ising Model

This is a continuation of Example 1.2. The local characteristics in the Ising model are

$$\pi_T^s(x) = \frac{e^{\frac{1}{kT} \{ J \sum_{t:|t-s|=1} x(t) + H \} x(s)}}{e^{\frac{1}{kT} \{ J \sum_{t:|t-s|=1} x(t) + H \}} + e^{-\frac{1}{kT} \{ J \sum_{t:|t-s|=1} x(t) + H \}}}. \tag{2.4}$$

The local energy at s is

$$\mathcal{E}_s(x) = \frac{1}{k} \left\{ J \sum_{t:|t-s|=1} x(t) + H \right\} x(s). \tag{2.5}$$

\diamond

2.2 From the Local Specification to the Potential

Theorem 2.1 above is the direct part of the *Gibbs–Markov equivalence* theorem: A Gibbs distribution relative to a neighborhood system is the distribution of a Markov field with respect to the same neighborhood system. The converse part, due to Hammersley and Clifford (1968), is important from a theoretical point of view, since together with the direct part it concludes that Gibbs distributions and MRFs are essentially the same objects, with a provision explained in the statement below.

Theorem 2.2. *Hammersley–Clifford Theorem*

Let π be the distribution of a Markov random field with respect to a graph (S, N) satisfying the positivity condition. Then

$$\pi(x) = \frac{1}{Z} e^{-\mathcal{E}(x)}$$

for some energy function $\mathcal{E}(x)$ deriving from a Gibbs potential $\{V_C\}_{C \subset S}$ associated with the topology (S, N) .

Proof. The following proof, due to Grimmett (1973), is based on the Möbius formula.

Lemma 2.1. *Möbius Formula*

Let Φ and Ψ be two set functions defined on $P(E)$, the collection of subsets of a finite set E . The two statements below are equivalent:

$$\Phi(A) = \sum_{B \subset A} (-1)^{|A-B|} \Psi(B), \text{ for all } A \subset E, \quad (2.5)$$

$$\Psi(A) = \sum_{B \subset A} \Phi(B), \text{ for all } A \subset E, \quad (2.6)$$

where $|C|$ is the number of elements of the set C .

Proof. We first show that (2.5) implies (2.6). Write the right-hand side of (2.6) using (2.5):

$$\begin{aligned} \sum_{B \subset A} \Phi(B) &= \sum_{B \subset A} \sum_{D \subset B} (-1)^{|B-D|} \Psi(D) \\ &= \sum_{D \subset A} \sum_{C \subset A-D} (-1)^{|C|} \Psi(D) = \sum_{D \subset A} \Psi(D) \sum_{C \subset A-D} (-1)^{|C|}. \end{aligned}$$

But if $A - D = \emptyset$,

$$\sum_{C \subset A-D} (-1)^{|C|} = (-1)^{|\emptyset|} = (-1)^0 = 1,$$

whereas if $A - D \neq \emptyset$,

$$\begin{aligned} \sum_{C \subset A-D} (-1)^{|C|} &= \sum_{k=0}^{|A-D|} (-1)^k \text{card} \{C; |C| = k, C \subset A - D\} \\ &= \sum_{k=0}^{|A-D|} (-1)^k \binom{|A-D|}{k} = (1-1)^{|A-D|} = 0, \end{aligned}$$

and therefore

$$\sum_{D \subset A} \Psi(D) \sum_{C \subset A-D} (-1)^{|C|} = \Psi(A).$$

We now show that (2.6) implies (2.5). Write the right-hand side of (2.5) using (2.6):

$$\sum_{B \subset A} (-1)^{|A-B|} \Psi(B) = \sum_{D \subset B \subset A} (-1)^{|A-B|} \Phi(D) = \sum_{D \subset A} \Phi(D) \sum_{C \subset A-D} (-1)^{|C|}$$

By the same argument as above, the last quantity equals $\Phi(A)$. □

We now prove Theorem 2.2. Let 0 be a fixed element of the phase space Λ . Also, let 0 denote the configuration with all phases equal to 0 . The context will prevent the confusion between $0 \in \Lambda$ and $0 \in \Lambda^S$. We may suppose in view of the positivity condition that $\pi(0) > 0$. Let x be a configuration, and let A be a subset of S . The notation x^A represents a configuration of Λ^S coinciding with x on A , and with phase 0 outside A .

Define for $A \subset S, x \in \Lambda^S$,

$$V_A(x) = \sum_{B \subset A} (-1)^{|A-B|} \ln \frac{\pi(0)}{\pi(x^B)}. \tag{2.7}$$

From the Möbius formula,

$$\ln \frac{\pi(0)}{\pi(x^A)} = \sum_{B \subset A} V_B(x).$$

Taking $A = S$ gives $\ln \frac{\pi(0)}{\pi(x)} = \sum_{B \subset S} V_B(x)$; that is,

$$\pi(x) = \pi(0)e^{-\sum_{A \subset S} V_A(x)}.$$

It remains to show that V_A depends only on the phases on A , and that $V_A \equiv 0$ if A is not a clique of (S, N) .

If $x, y \in \Lambda^S$ are such that $x(A) = y(A)$, then for any $B \subset A, x^B = y^B$, and therefore, by (2.7), $V_A(x) = V_A(y)$.

Using an arbitrary site $t \in A$, write (2.7) as follows:

$$\begin{aligned} V_A(x) &= \left\{ \sum_{B \subset A, B \not\ni t} + \sum_{B \subset A, B \ni t} \right\} (-1)^{|A-B|} \ln \frac{\pi(0)}{\pi(x^B)} \\ &= \sum_{B \subset A-t} (-1)^{|A-B|} \left\{ \ln \frac{\pi(0)}{\pi(x^B)} - \ln \frac{\pi(0)}{\pi(x^{B+t})} \right\}. \end{aligned}$$

That is,

$$V_A(x) = \sum_{B \subset A-t} (-1)^{|A-B|} \ln \frac{\pi(x^{B+t})}{\pi(x^B)}. \tag{2.8}$$

Now, if t is not in $B \subset A$, then

$$\frac{\pi(x^{B+t})}{\pi(x^B)} = \frac{\pi^t(x^{B+t})}{\pi^t(x^B)},$$

and therefore

$$V_A(x) = \sum_{B \subset A-t} (-1)^{|A-B|} \ln \frac{\pi^t(x^{B+t})}{\pi^t(x^B)},$$

and, by the same calculations that led to (2.8),

$$V_A(x) = - \sum_{B \subset A} (-1)^{|A-B|} \ln \pi^t(x^B). \tag{2.9}$$

Recall that $t \in A$, and therefore, if A is not a clique, one can find $s \in A$ such that s is not a neighbor of t . Fix such an s , and split the sum in (2.9) as follows:

$$\begin{aligned} V_A(x) &= - \sum_{B \subset A-s-t} (-1)^{|A-B|} \ln \pi^t(x^B) - \sum_{B \subset A-s-t} (-1)^{|A-(B+t)|} \ln \pi^t(x^{B+t}) \\ &\quad - \sum_{B \subset A-s-t} (-1)^{|A-(B+s)|} \ln \pi^t(x^{B+s}) - \sum_{B \subset A-s-t} (-1)^{|A-(B+s+t)|} \ln \pi^t(x^{B+s+t}) \\ &= - \sum_{B \subset A-s-t} (-1)^{|A-B|} \ln \frac{\pi^t(x^B)\pi^t(x^{B+t+s})}{\pi^t(x^{B+s})\pi^t(x^{B+t})}. \end{aligned}$$

But since $s \neq t$ and $s \notin \mathcal{N}_t$, we have $\pi^t(x^B) = \pi^t(x^{B+s})$ and $\pi^t(x^{B+t}) = \pi^t(x^{B+t+s})$, and therefore $V_A(x) = 0$. \square

It must be noted that the energy function \mathcal{E} and the partition function are not unique, since one can add a constant to the energy function and then multiply the normalizing factor by an appropriate constant. Likewise, and more importantly, the Gibbs potential associated with π is not unique. However, uniqueness can be forced into the result if a certain property is imposed on the potential, namely normalization with respect to a fixed phase value: See Theorem 2.3 below.

In many situations, the potential as well as the topology of S can be obtained directly from the expression of the energy.

Definition 2.1. *Normalized Potential*

A Gibbs potential $\{V_C\}_{C \subset S}$ is said to be normalized with respect to a given phase in Λ , conventionally denoted by θ , if $V_C(x) = 0$ whenever there exists $t \in C$ such that $x(t) = \theta$.

Theorem 2.3. *Uniqueness of Normalized Potential*

There exists one and only one potential normalized with respect to a given phase $\theta \in \Lambda$ corresponding to a Gibbs distribution.

Proof. Formula (2.8) produces for potential (2.7) the expression

$$V_C(x) = \sum_{B \subset C-t} (-1)^{|A-B|} \ln \frac{\pi(x^{B+t})}{\pi(x^B)}$$

independent of t in the clique C . In particular, choosing any $t \in C$ such that $x(t) = \theta$, $x^{B+t} = x^B$ for all $B \subset C - t$, and therefore $V_C(x) = 0$. This shows that (2.7) gives a potential normalized with respect to $\theta \in \Lambda$. To prove uniqueness, suppose that

$$\pi(x) = \frac{1}{Z_U} e^{-\mathcal{E}_U(x)} = \frac{1}{Z_W} e^{-\mathcal{E}_W(x)}$$

for two energy functions \mathcal{E}_U and \mathcal{E}_W deriving from potentials U and W , respectively, both normalized with respect to $0 \in \Lambda$. Since $\mathcal{E}_U(0) = \sum_{C \in \mathcal{C}} U_C(0) = 0$, and similarly $\mathcal{E}_W(0) = 0$, it follows that $Z_U = Z_W = \pi(0)^{-1}$, and therefore $\mathcal{E}_U \equiv \mathcal{E}_W$. Suppose that $U_A = W_A$ for all $A \in \mathcal{C}$ such that $|A| \leq k$ (property \mathcal{P}_k). It remains to show, in view of a proof by induction, that $\mathcal{P}_k \Rightarrow \mathcal{P}_{k+1}$ and that \mathcal{P}_1 is true.

To prove $\mathcal{P}_k \Rightarrow \mathcal{P}_{k+1}$, fix $A \subset S$ with $|A| = k + 1$. To prove that $U_A \equiv W_A$ it suffices to show that $U_A(x) = W_A(x)$ for all $x \in \Lambda^S$ such that $x = x^A$. Fix such an x . Then

$$\mathcal{E}_U(x) = \sum_C U_C(x) = \sum_{C \subset A} U_C(x),$$

since x has phase 0 outside A and U is normalized with respect to 0. Also,

$$\mathcal{E}_U(x) = \sum_{C \subset A} U_C(x) = U_A(x) + \sum_{C \subset A, |C| \leq k} U_C(x), \tag{2.10}$$

with a similar equality for $\mathcal{E}_W(x)$. Therefore, since $\mathcal{E}_U(x) = \mathcal{E}_W(x)$, we obtain $U_A(x) = W_A(x)$ in view of the induction hypothesis. The root \mathcal{P}_1 of the induction hypothesis is true, since when $|A| = 1$, (2.10) becomes $\mathcal{E}_U(x) = U_A(x)$, and similarly, $\mathcal{E}_W(x) = W_A(x)$, so that $U_A(x) = W_A(x)$ is a consequence of $\mathcal{E}_U(x) = \mathcal{E}_W(x)$. \square

Example 2.3. *Markov Chains as Markov Fields.*

Let $S = \{0, 1, \dots, N\}$ and $\Lambda = E$, a finite space. A random field X on S with phase space Λ is therefore a vector X with values in E^{N+1} . Suppose that X_0, \dots, X_N is a homogeneous Markov chain with transition matrix $\mathbf{P} = \{p_{ij}\}_{i,j \in E}$ and initial distribution $\nu = \{\nu_i\}_{i \in E}$. In particular, with $x = (x_0, \dots, x_N)$,

$$\pi(x) = \nu_{x_0} p_{x_0 x_1} \cdots p_{x_{N-1} x_N},$$

that is,

$$\pi(x) = e^{-\mathcal{E}(x)},$$

where

$$\mathcal{E}(x) = -\ln \nu_{x_0} - \sum_{n=0}^{N-1} (\ln p_{x_n x_{n+1}}).$$

Clearly, this energy derives from a Gibbs potential associated with the nearest-neighbor topology for which the cliques are, besides the singletons, the pairs of adjacent sites. The nonnull potential functions can be chosen as follows:

$$V_{\{0\}}(x) = -\ln \nu_{x_0}, \quad V_{\{n, n+1\}}(x) = -\ln p_{x_n x_{n+1}}.$$

The local characteristic at site $s = n \in [2, N - 1]$ is given by Formula (2.1),

$$\pi^n(x) = \frac{\exp(\ln p_{x_{n-1} x_n} + \ln p_{x_n x_{n+1}})}{\sum_{y \in E} \exp(\ln p_{x_{n-1} y} + \ln p_{y x_{n+1}})};$$

that is,

$$\pi^n(x) = \frac{P_{x_{n-1}, x_n} P_{x_n, x_{n+1}}}{P_{x_{n-1}, x_{n+1}}^{(2)}},$$

where $p_{ij}^{(2)}$ is the general term of the two-step transition matrix \mathbf{P}^2 . Similar computations give $\pi^0(x)$ and $\pi^N(x)$.

In view of the neighborhood structure, for $n \in [2, N - 1]$, X_n is independent of $X_0, \dots, X_{n-2}, X_{n+2}, \dots, X_N$ given X_{n-1} and X_{n+1} . \diamond

Example 2.4. *Two-State Markov Chain*

Example 2.3 will be particularized to a Markov chain with state space $E = \{-1, 1\}$ and transition matrix

$$\mathbf{P} = \begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix},$$

where $\alpha, \beta \in (0, 1)$, and with the stationary initial distribution

$$(v_0, v_1) = \frac{1}{\alpha + \beta}(\beta, \alpha).$$

Observing that if $x \in \{-1, 1\}$, then $\frac{1-x}{2} = 1_{x=-1}$ and $\frac{1+x}{2} = 1_{x=+1}$, we have

$$4 \ln p_{xy} = (1-x)(1+y) \ln(\alpha) + (1-x)(1-y) \ln(1-\alpha) \\ + (1+x)(1-y) \ln(\beta) + (1+x)(1+y) \ln(1-\beta)$$

and

$$2 \ln v(x) = (1-x) \ln \left(\frac{\beta}{\alpha + \beta} \right) + (1+x) \ln \left(\frac{\alpha}{\alpha + \beta} \right).$$

Therefore, particularizing the expression for the energy in Example 2.3, we obtain

$$-\mathcal{E}(x) = \frac{1}{4} N \ln(\beta\alpha(1-\beta)(1-\alpha)) + \frac{1}{2} \ln \left(\frac{\beta\alpha}{(\alpha + \beta)^2} \right) \\ + \frac{1}{2} \ln \left(\frac{1-\beta}{1-\alpha} \right) \sum_{i=0}^N x_i + \frac{1}{4} (x_0 + x_N) \ln \left(\frac{\alpha(1-\beta)}{\beta(1-\alpha)} \right) \\ + \frac{1}{4} \left(\sum_{n=0}^{N-1} x_n x_{n+1} \right) \ln \left(\frac{(1-\alpha)(1-\beta)}{\alpha\beta} \right).$$

Here the normalized potential with respect to 0 appears naturally in the nonconstant terms. \diamond

Example 2.5. *Unilateral MRF*

Example 2.3 shows that the Markovian character of a Markov chain may be expressed in two ways, the choice being between a unilateral description in terms of the transition matrix

and a bilateral description in terms of the local characteristics. The notion of unilateral Markov random field generalizes this situation.

A unilateral Markov random field receives a description analogous to that of a Markov chain, except that the time sequence $1, 2, 3, \dots$ is now a sequence of sites s_1, s_2, \dots, s_N exhausting all sites of S without repetition. Time is indirectly involved in that at least for some applications in image processing, the sites of S are scanned sequentially in the above order. For this reason, such enumeration is called a *scanning* of S .

From Bayes’s sequential rule,

$$\pi(x) = P(X(s_1) = x(s_1)) \prod_{j=2}^N P(X(s_j) = x(s_j) \mid X(s_1) = x(s_1), \dots, X(s_{j-1}) = x(s_{j-1})).$$

Let X be a random field on S with phase space Λ , and let s_1, \dots, s_N be a scanning of S . Let $\mathcal{P}_1 = \mathcal{P}(s_1) = \emptyset, \mathcal{P}_2 = \mathcal{P}(s_2), \dots, \mathcal{P}_N = \mathcal{P}(s_N)$ be a sequence of subsets of S such that for all $j \geq 1$,

$$\mathcal{P}(s_j) \subset \{s_1, \dots, s_{j-1}\} = S_{j-1}.$$

Suppose that for all $j \geq 1$,

$$P(X(s_j) = x(s_j) \mid X(S_{j-1}) = x(S_{j-1})) = P(X(s_j) = x(s_j) \mid X(\mathcal{P}(s_j)) = x(\mathcal{P}(s_j))). \tag{2.11}$$

Then $\{\mathcal{P}(s)\}_{s \in S}$ is called a *past relevance structure* with respect to the scanning s_1, \dots, s_N ; $\mathcal{P}(s)$ is called the *relevant past* at site s ; and X is called a *unilateral MRF* with respect to $\{\mathcal{P}(s)\}_{s \in S}$ and the scanning s_1, \dots, s_N .

The *transition function* of X at site $s \in S$ is the following function of $x(s)$ and $x(\mathcal{P}(s))$:

$$p(x(s) \mid x(\mathcal{P}(s))) = P(X(s) = x(s) \mid X(\mathcal{P}(s)) = x(\mathcal{P}(s))), \tag{2.12}$$

and therefore, for any configuration x ,

$$\pi(x) = \prod_{j=1}^N p(x(s_j) \mid x(\mathcal{P}(s_j))) \tag{2.13}$$

The random field of Definition 2.2 is an MRF with respect to the neighborhood system given by

$$\mathcal{N}_s = \left(\bigcup_{t: \mathcal{P}(t) \ni s} \{\mathcal{P}(t) \cup t\} \setminus s \right) \cup \mathcal{P}(s). \tag{2.14}$$

To prove this, it suffices, by Lemma 7.1 of Chapter 1, to show that $P(X(s) = x(s) \mid X(S \setminus s) = x(S \setminus s))$ depends only upon $x(\mathcal{N}_s \cup \{s\})$, where \mathcal{N}_s is given by (2.14). Bayes’s rule gives

$$\pi^s(x) = \frac{\pi(x(s), x(S \setminus s))}{\sum_{\lambda \in \Lambda} \pi(\lambda, x(S \setminus s))}. \tag{2.15}$$

From (2.13) we have $\pi(\lambda, x(S \setminus s)) = A \times B$, where A is independent of λ and therefore factors out in the last ratio, and

$$B = p(\lambda \mid x(\mathcal{P}(s))) \times \prod_{t:s \in \mathcal{P}(t)} p(x(t) \mid \lambda, x(\mathcal{P}(t) \setminus s)).$$

Therefore,

$$\pi^*(x) = \frac{p(x(s) \mid x(\mathcal{P}(s))) \times \prod_{t:s \in \mathcal{P}(t)} p(x(t) \mid x(\mathcal{P}(t)))}{\sum_{\lambda \in \Lambda} p(\lambda \mid x(\mathcal{P}(s))) \times \prod_{t:s \in \mathcal{P}(t)} p(x(t) \mid \lambda, x(\mathcal{P}(t) \setminus s))}, \quad (2.16)$$

a quantity depending only upon $x(\mathcal{N}_s \cup \{s\})$ as announced. \diamond

3 Image Models

3.1 Textures

For the purpose of image synthesis, one seeks Gibbs distributions describing pictures featuring various textures, lines separating patches with different textures (boundaries), lines per se (roads, rail tracks), randomly located objects (moon craters), etc. The corresponding model is then checked by simulation (see Section 6): Images are drawn from the proposed Gibbs distribution, and some tuning of the parameters is done, until the images correspond to what is desired. Image synthesis is an art based on trial and error, and fortunately guided by some general principles. But these principles are difficult to formalize, and we shall mainly resort to simple examples with a pedagogical value. Note, however, that there is a domain of application where the model need not be very accurate, namely Bayesian estimation. As a matter of fact, the models proposed in this section have been devised in view of applications to Bayesian restoration of degraded pictures (see Section 4). We shall begin with a famous all-purpose texture model.

Example 3.1. Autobinomial Model

This model was introduced by Besag (1974) to describe the texture of various materials. The set of sites is $S = \mathbb{Z}_m^2$, and the phase space is $\Lambda = \{0, 1, \dots, L\}$. In the context of image processing, a site s is a pixel (PICTure ELEment), and a phase $\lambda \in \Lambda$ is a shade of grey, or a color. The neighborhood system is

$$\mathcal{N}_s = \{t \in S; t \neq s \text{ and } \|t - s\|^2 \leq d\}, \quad (3.1)$$

where d is a fixed positive integer and where $\|t - s\|$ is the euclidean distance between s and t . In this model the only cliques participating in the energy function are singletons and pairs of mutual neighbors. The set of cliques appearing in the energy function is a disjoint sum of collections of cliques

$$\mathcal{C} = \sum_{j=1}^{m(d)} \mathcal{C}_j,$$

where \mathcal{C}_1 is the collection of singletons, and all pairs $\{s, t\}$ in \mathcal{C}_j , $2 \leq j \leq m(d)$, have the same distance $\|t - s\|$ and the same direction, as shown in Fig. 7.3.1. The potential is given by

$$V_C(x) = \begin{cases} -\ln \binom{L}{x(s)} + \alpha_1 x(s) & \text{if } C = \{s\} \in \mathcal{C}_1, \\ \alpha_j x(s)x(t) & \text{if } C = \{s, t\} \in \mathcal{C}_j, \end{cases}$$

where $\alpha_j \in \mathbb{R}$. For any clique C not of type \mathcal{C}_j , $V_C \equiv 0$.

The terminology (“autobinomial”) is motivated by the fact that the local system has the form

$$\pi^s(x) = \binom{L}{x(s)} \tau^{x(s)} (1 - \tau)^{L-x(s)}, \tag{3.2}$$

where τ is a parameter depending on $x(\mathcal{N}_s)$ as follows:

$$\tau = \tau(\mathcal{N}_s) = \frac{e^{-(\alpha, b)}}{1 + e^{-(\alpha, b)}}. \tag{3.3}$$

Here (α, b) is the scalar product of

$$\alpha = (\alpha_1, \dots, \alpha_{m(d)}) \text{ and } b = (b_1, \dots, b_{m(d)}),$$

where $b_1 = 1$, and for $j \in [2, m(d)]$,

$$b_j = b_j(x(\mathcal{N}_s)) = x(r) + x(t),$$

where $\{r, s\}$ and $\{t, s\}$ are the two pairs in \mathcal{C}_j containing s . To prove (3.2), one uses the explicit formula (2.1) giving the local characteristic at site s to obtain

$$\pi^s(x) = \frac{\exp \left\{ \ln \binom{L}{x(s)} - \alpha_1 x(s) - \left[\sum_{j=2}^{m(d)} \alpha_j \sum_{t:\{s,t\} \in \mathcal{C}_j} x(t) \right] x(s) \right\}}{\sum_{\lambda \in \Lambda} \exp \left\{ \ln \binom{L}{\lambda} - \alpha_1 \lambda - \left[\sum_{j=2}^{m(d)} \alpha_j \sum_{t:\{s,t\} \in \mathcal{C}_j} x(t) \right] \lambda \right\}}.$$

The numerator of the fraction equals

$$\binom{L}{x(s)} e^{-(\alpha, b)x(s)},$$

and the denominator is

$$\sum_{\lambda \in \Lambda} \binom{L}{\lambda} e^{-(\alpha, b)\lambda} = \sum_{\ell=0}^L \binom{L}{\ell} (e^{-(\alpha, b)})^\ell = (1 + e^{-(\alpha, b)})^L.$$

Equality (3.2) then follows.

Expression (3.2) shows that τ is the average level of grey at site s , given $x(\mathcal{N}_s)$, and expression (3.3) shows that τ is a function of (α, b) . The parameter α_j controls the bond in the direction and at the distance that characterize \mathcal{C}_j . \diamond

Generalizations of the autobinomial model are easy to conceive: Replace $\binom{L}{x(s)}$ in the expression of the potential by a function of $x(s)$. See Problem 7.3.3 featuring the *auto-Poisson model*.




Figure 7.3.1. Neighborhoods and cliques of three autobinomial models

3.2 Lines and Points

Let $X = \{X(s)\}_{s \in S}$ be a random field on S with phase space Λ , with the following structure:

$$S = S_1 + S_2, \quad \Lambda = \Lambda_1 \cup \Lambda_2,$$

and

$$\begin{aligned} X(s) &= Y(s_1) \in \Lambda_1 \text{ if } s = s_1 \in S_1 \\ &= Z(s_2) \in \Lambda_2 \text{ if } s = s_2 \in S_2. \end{aligned}$$

Here S_1 and S_2 are two disjoint collections of sites that can be of a different nature, or have different functions, and Λ_1 and Λ_2 are phase spaces that need not be disjoint. Define

$$Y = \{Y(s_1)\}_{s_1 \in S_1}, \quad Z = \{Z(s_2)\}_{s_2 \in S_2},$$

so that

$$X = (Y, Z).$$

As such, the field X is just the juxtaposition of Y and Z . The statistical interactions between Y and Z are described by the distribution of X derived from an energy function, or equivalently from a potential, with a neighborhood structure N . Efficient models have neighborhoods where both types of sites are present, and their neighborhoods are relatively small.

In some situations, Y is the *observed* field, and Z is the *hidden* field. The introduction of a hidden field is in principle motivated by physical considerations. From the computational point of view, it is justified by the fact that the field Y alone usually has a Gibbsian description with large neighborhoods, whereas $X = (Y, Z)$ hopefully has small neighborhoods.

Example 3.2. The Pixel–Edge Model

The philosophy supporting the model introduced by D. and S. Geman (1984) is the following. A digitized image can be viewed as a realization of a random field on $S^P = \mathbb{Z}_m^2$. A site could be, for instance, a pixel on a digital television screen, and therefore S^P will be called the set of *pixel* sites. For an observer, there is, in general, more in an image than just the colors at each pixel. For instance, an image can be perceived as a juxtaposition of zones with various textures separated by lines. However, these lines, or contours, are not seen directly on the pixels, they are inferred from them by some processing in the brain. On the other hand, if one is to sketch the picture observed on the screen, one would most likely start by drawing the lines. In any case, textures and contours are very much linked, and one should seek a description featuring the interaction between them. But as was mentioned, contours do not exist on the digital screen, they are hidden, or more accurately, they are virtual.

In the model of D. and S. Geman, there is a set S^E of *edge* sites, one between each pair of adjacent pixel sites, as indicated in Fig. 7.3.2a. The possible values of the phase on an edge site are blank or bar: horizontal (resp. vertical) for an edge site between two pixel sites forming a vertical (resp. horizontal) segment, as shown in Fig. 7.3.2b. In Fig. 7.3.2b not all edge sites between two pixels with a different color have a bar, because a good model should not systematically react to what may be accidents in the global structure.

Let (i, j) denote a pixel site and (α, β) an edge site (these are the coordinates of the sites in two distinct orthogonal frames). The random field on the pixels is denoted by $X^P = \{X_{ij}^P\}_{(i,j) \in S^P}$, and that on the edge sites is $X^E = \{X_{\alpha\beta}^E\}_{(\alpha,\beta) \in S^E}$; X^P is the observed

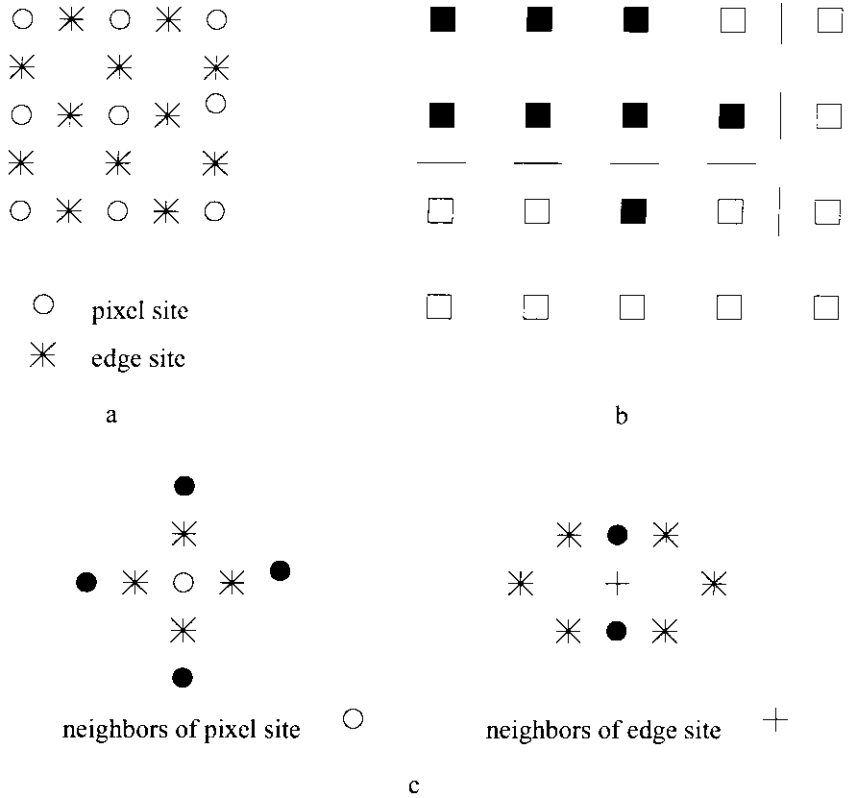


Figure 7.3.2. Example of a neighborhood structure for the pixel-edge model

image, and X^L is the hidden, or virtual, line field. The distribution of the field $X = (X^P, X^E)$ is described by an energy function $\mathcal{E}(x^P, x^E)$:

$$\pi(x^P, x^E) = \frac{1}{Z} e^{-\mathcal{E}(x^P, x^E)}, \tag{3.4}$$

where $x^P = \{x_{ij}^P\}_{(i,j) \in S^P}$, $x^E = \{x_{\alpha\beta}^E\}_{(\alpha,\beta) \in S^E}$. The energy function derives from a potential relative to some neighborhood system, a particular choice of which is pictured in Fig. 7.3.2c. The energy function separates into two parts

$$\mathcal{E}(x^P, x^E) = \mathcal{E}_1(x^P, x^E) + \mathcal{E}_2(x^E), \tag{3.5}$$

where \mathcal{E}_1 features only cliques formed by a pair of neighboring pixels and the edge pixel in between, whereas \mathcal{E}_2 features only the diamond cliques shown in Fig. 7.3.3. The energy $\mathcal{E}_1(x^P, x^E)$ governs the relation between an edge and its two adjacent pixels. For instance,

for some real constant $\alpha > 0$ and some function φ ,

$$\mathcal{E}_1(x^P, x^E) = -\alpha \sum_{(1,2)} \varphi(x_1^P - x_2^P)x_{(1,2)}^E,$$

where $(1, 2)$ represents a pair of adjacent pixel sites and $x_{(1,2)}^E$ is the value of the phase on the edge site between the pixel sites 1 and 2, say 0 for a blank and 1 for a bar. A possible choice of φ is

$$\varphi(x) = \begin{cases} +1 & \text{if } x \neq 0, \\ -1 & \text{if } x = 0. \end{cases}$$

Since the most probable configurations are those with low energy, this model favors bars between two adjacent pixel sites with different colors, as is natural. More sophisticated choices of φ with a similar effect are possible. The organization of the contours is controlled by the energy

$$\mathcal{E}_2(x^E) = \beta \sum_D w_D(x^E),$$

where $\beta > 0$ and the sum extends to all diamond cliques, and w_D is a function depending only on the phases of the four edge sites of the diamond clique D . Up to rotations of $\frac{\pi}{2}$, there are six possible values for the four-vector of phases on a given diamond clique D , as shown in Fig. 7.3.3. If the modeler believes that for the type of images he is interested in the likelihood of the configurations shown in Fig. 7.3.3 decreases from left to right, then the values of $w_D(x^E)$ attributed to these configurations will increase from left to right. This is generally the case because four-country border points are rare, broken lines also, and the same is true to a lesser extent for three-country border points. Also, when the picture is not a clutter of lines, the no-line configuration is the most likely.




Figure 7.3.3. Six configurations of the diamond clique

Although this example is representative, it admits many variations. However, too many sophisticated features could ruin the model. The purpose of the model of D. and S. Geman is not so much to do image synthesis as to have a reasonable a priori model for the image F in view of Bayesian restoration of this image from a noisy version of it. This topic will be taken up later on in Example 4.2, where the model will be further discussed.

Example 3.3. Points

A (Markov) random field X on S with phase space

$$\Lambda = \{0, 1\}$$

is also called a (Markov) *point process* on S . If $X(s) = 1$, one says that there is a *point* at s . When taking the point process point of view, a configuration x is identified with its support:

$$\text{supp}(x) = \{s \in S; x(s) = 1\}.$$

The notation $x \cup s$ represents the configuration identical to x but with phase 1 at site s . If $s \in x$, then $x = x \cup s$. The configuration $x \setminus s$ is x with phase 0 at s , and it is identical to x if $s \notin x$.

Let π be a probability distribution on the configuration space $\{0, 1\}^S$. The positivity condition implies that

$$(y \subset x) \implies (\pi(x) > 0 \implies \pi(y) > 0).$$

The following result is a particular case of a result due to Kelly and Ripley (1977) in the context of Gibbs point processes on a general space.

Let π be a probability distribution on $\{0, 1\}^S$ satisfying the positivity condition. There exists a function $\varphi : \{0, 1\}^S \rightarrow (0, 1]$ such that

$$\pi(x) = \prod_{y \subset x} \varphi(y) \tag{3.6}$$

and

$$(\varphi(y) < 1) \implies ((s, t \in x \implies s \sim t) \text{ or } y = \emptyset),$$

where $s \sim t$ means that s and t are neighbors.

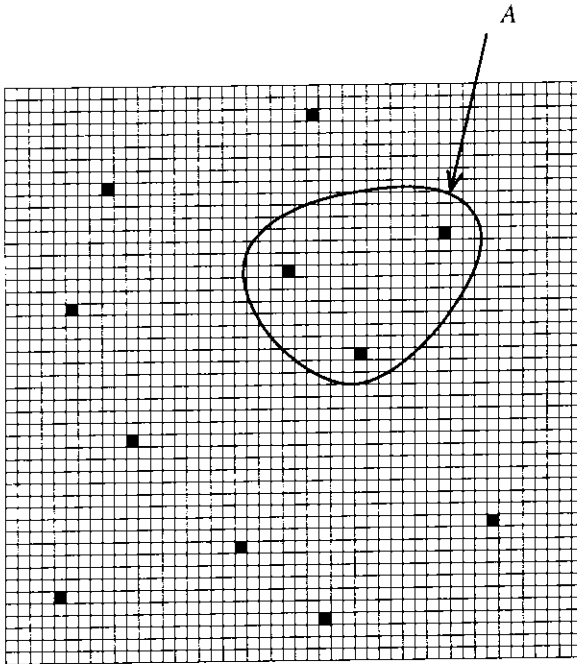
This is a rephrasing of the Hammersley–Clifford theorem, according to which there exists a potential $\{V_C\}_{C \subset S}$ normalized with respect to the null phase, and such that

$$\pi(x) = \frac{1}{Z} e^{-\sum_C V_C(x)}.$$

Now, $V_C(x) > 0$ if and only if C is a clique and $C \subset x$ (normalized potential). Identifying C with a nonempty configuration y , and setting

$$\varphi(y) = e^{-V_y(y)},$$

we obtain the announced result, having observed that if $y \subset x$ then $V_y(x) = V_y(y)$, and that $\varphi(\emptyset) = \frac{1}{Z} = \pi(\emptyset)$. \diamond



$$\Phi(A) = 3$$

Figure 7.3.4. A point process

4 Bayesian Restoration of Images

4.1 MAP Likelihood Estimation

A number of estimation problems arising in various areas of statistics and engineering, and particularly in image processing, are solved by a method of statistical analysis known as the *maximum a posteriori* (MAP) likelihood method. The theory will not be presented, but the examples below will show its substance. These examples compute the a posteriori probability, or conditional probability, of a given MRF X with respect to another MRF Y , the observed field, say,

$$\pi(x | y) = P(X = x | Y = y),$$

and the MAP method estimates the nonobservable field X given the observed value y of Y , by $\hat{x} = \hat{x}(y)$, the value of x that maximizes $\pi(x | y)$:

$$\hat{x}(y) = \arg \max_x \pi(x | y).$$

Usually, this maximization problem is doomed by combinatorial explosion and by the complexity of standard methods of optimization. However, if

$$\pi(x | y) \propto e^{-\mathcal{E}(x | y)} \quad (4.1)$$

(the proportionality factor depends only on y and is therefore irrelevant to maximization with respect to x) with an energy $\mathcal{E}(x | y)$ that as a function of x corresponds to a topology N with small neighborhoods, then one can use a simulated annealing algorithm or a related algorithm (see Section 8).

Example 4.1. *Random Flips or Multiplicative Noise*

Let X, Z be random fields on $S = \mathbb{Z}_m^2$ with phase space $\Lambda = \{-1, +1\}$, and define the field $Y = XZ$ by $y(s) = x(s)z(s)$, $s \in S$. The field Z will be interpreted as multiplicative noise, and one can call it a random flip field, because what it does to X is to flip the phase at site s if $z(s) = -1$.

The computation below uses the fact that if $a, b, c \in \Lambda^S$, where $\Lambda = \{-1, +1\}$, then $ab = c \Leftrightarrow b = ac$:

$$\begin{aligned} P(Y = y)P(X = x | Y = y) &= P(X = x, Y = y) = P(X = x, ZX = y) \\ &= P(X = x, Zx = y) = P(X = x, Z = yx). \end{aligned}$$

In particular, if the noise field Z is independent of the original field X , then

$$\pi(x | y) \propto P(X = x)P(Z = yx).$$

The random field X has the distribution

$$P(X = x) \propto e^{-\mathcal{E}(x)}.$$

Suppose that $(Z(s), s \in S)$ is a family of i.i.d. random variables, with $P(Z(s) = -1) = p$, $P(Z(s) = +1) = q = 1 - p$. Therefore (see Problem 7.2.5),

$$P(Z = z) = \prod_{s \in S} P(Z(s) = z(s)) \propto e^{\gamma \sum_{s \in S} z(s)}, \quad (4.2)$$

where $\gamma = \frac{1}{2} \ln \left(\frac{1-p}{p} \right)$. Finally,

$$\pi(x | y) \propto e^{-\mathcal{E}(x) + \gamma \sum_{s \in S} y(s)x(s)}. \quad (4.3)$$

Note that if $p > \frac{1}{2}$, then $\gamma > 0$. ◇

Example 4.2. *Image Restoration*

This example continues Example 3.2. Recall that we have a Gibbs field $X = (X^P, X^E)$ corresponding to some energy function $\mathcal{E}(x^P, x^E)$, which need not be made precise here (see, however, Example 3.2).

The image X^P is degraded into a noisy image Y , and it is this corrupted image that is observed. Degradation combines two effects: blurring, and a possibly nonlinear interaction of the blurred image and the noise. Specifically,

$$Y_{ij} = \varphi(H(X^P)_{ij}, N_{ij}), \quad (4.4)$$

where (i, j) is a pixel site and φ , H , and N are defined as follows.

First $N = \{N_{ij}\}_{(i,j) \in P_m}$ is, for instance, a family of independent centered Gaussian random variables with common variance σ^2 , and is independent of (X^P, X^E) . As for $H(X^P)$, it is the random field obtained by blurring X^P , that is,

$$H(X^P)_{ij} = \sum_{k,\ell} H_{k\ell} X^P_{i-k,j-\ell}, \quad (4.5)$$

where $H = \{H_{k\ell}\}_{-N \leq k, \ell \leq N}$ is the *blurring matrix*. In (4.5), $X^P_{i-k,j-\ell} = 0$ if $(i-k, j-\ell) \notin S^P$. A typical blurring matrix is

$$H = \begin{pmatrix} 1/80 & 1/80 & 1/80 \\ 1/80 & 9/10 & 1/80 \\ 1/80 & 1/80 & 1/80 \end{pmatrix}.$$

for which $N = 1$. In this case

$$H(X^P)_{ij} = \frac{9}{10} X^P_{ij} + \frac{1}{80} \left(\sum_{\ell} X^P_{k,\ell} \right),$$

where the sum extends to the pixel sites adjacent to (i, j) . As for φ , it is a function such that for fixed a , the function $b \rightarrow \varphi(a, b)$ is invertible. The inverse of this function, for fixed a , is then denoted by $b \rightarrow \psi(a, b)$. A typical example for φ is the additive noise model

$$Y_{ij} = H(X^P)_{ij} + N_{ij}. \quad (4.6)$$

To estimate X given Y , one is led to compute the a posteriori probability of image x given that the noisy image is y :

$$\pi(x^P, x^E | y) = P(X^P = x^P, X^E = x^E | Y = y).$$

Writing $\pi(y) = P(Y = y)$, we have

$$\begin{aligned} \pi(x^P, x^E | y) &= \pi(y) P(X^P = x^P, X^E = x^E, Y = y) \\ &= \pi(y) P(X^P = x^P, X^E = x^E, \varphi(H(X^P), N) = y) \\ &= \pi(y) P(X^P = x^P, X^E = x^E, N = \psi(H(x^P), y)) \\ &= \pi(y) P(X^P = x^P, X^E = x^E) P(N = \psi(H(x^P), y)). \end{aligned}$$

The reader will have noticed the abuse of notation by which the continuous character of N was ignored: The second to fourth terms are actually probability densities, and similarly for

$$P(N = \psi(H(x^P), y)) \propto e^{-\frac{1}{2\sigma^2} \|\psi(H(x^P), y)\|^2}.$$

Using the expression of the distribution of the pixel+line image in terms of the energy function, one finds

$$\pi(x^P, x^E | y) \propto e^{-\mathcal{E}(x^P, x^E) - \frac{1}{2\sigma^2} \|\psi(H(x^P), y)\|^2}. \quad (4.7)$$

Therefore the *a posteriori* distribution of (X^P, X^E) given $Y = g$ is a Gibbs distribution corresponding to the energy function

$$U(x^P, x^E) = \mathcal{E}(x^P, x^E) + \frac{1}{2\sigma^2} \|\psi(H(x^P), y)\|^2. \quad (4.8)$$

For instance, if the noise is additive, as in (4.6), then

$$U(x^P, x^E) = \mathcal{E}(x^P, x^E) + \frac{1}{2\sigma^2} \|y - H(x^P)\|^2. \quad (4.9)$$

◇

Example 4.3. Bernoulli–Gaussian Model

Let $\{Y_n\}_{1 \leq n \leq N}$ be a real-valued stochastic process of the form

$$Y_n = \sum_{k=1}^N X_k h_{n-k} + Z_n.$$

where $\{Z_n\}_{1 \leq n \leq N}$ is a sequence of independent centered Gaussian random variables of variance σ^2 , $\{X_n\}_{1 \leq n \leq N}$ is an i.i.d. sequence of $\{0, 1\}$ -valued random variables with $P(X_n = 1) = p$, and $\{h_k\}_{k \in \mathbb{Z}}$ is a deterministic function.

This is a particular case of the one-dimensional version of the model in Example 4.3. Here $S = \{1, \dots, N\}$, a configuration $x \in \{0, 1\}^N$ is of the form $x = (x_1, x_2, \dots, x_N)$, X is the original image, Y is the degraded image, Z is the additive noise, and h corresponds to the blurring matrix. For this particular model, the energy of the i.i.d. field X is of the form $\gamma \sum_{i=1}^N x_i$ (see Problem 7.2.6), and the energy of the conditional field $x|y$ is

$$\gamma \sum_{i=1}^N x_i + \frac{1}{2\sigma^2} \sum_{i=1}^N \left| y_i - \sum_{j:1 \leq i-j \leq N} h_j x_{i-j} \right|^2.$$

This model is often used in problems of detection of reflectors. One says that there is a reflector at position i if $X_i = 1$. The function h is a probe signal (radar, sonar), and $\{h_{k-i}\}_{k \in \mathbb{Z}}$ is the signal reflected by the reflector at position i , if any, so that $Y_n = \sum_{k=1}^N X_k h_{n-k}$ is the reflected signal from which the map of reflectors X is to be recovered. The process Z is the usual additive noise of signal processing.

Of course, this model can be considerably enriched by introducing random reflection coefficients or by using a more elaborate a priori model for X , say, a Markov chain model. ◇

4.2 Penalty Methods

Consider the simple model where the image X is additively corrupted by white Gaussian noise N of variance σ^2 , and let Y be the resulting image. Calling $\mathcal{E}(x)$ the energy function of the a priori model, the MAP estimate is

$$\hat{x} = \arg \min_x \left\{ \mathcal{E}(x) + \frac{1}{\sigma^2} \|y - x\|^2 \right\}.$$

If we take an a priori model where all images are equiprobable, that is, the corresponding energy is null, then the above minimization is trivial, leading one to accept the noisy image as if it were the original image. A nontrivial a priori model introduces a penalty term $\mathcal{E}(x)$ and forces a balance between our belief in the observed image, corresponding to a small value of $\|y - x\|^2$, and our a priori expectation as to what we should obtain, corresponding to a small value of $\mathcal{E}(x)$. The compromise between the credibility of the observed image and the credibility of the estimate with respect to the prior distribution is embodied in the criterion $\mathcal{E}(x) + \frac{1}{\sigma^2} \|y - x\|^2$. A non-bayesian mind will, somehow rightly, argue that one cannot even dream of thinking that a correct a priori model is available, and that Gaussian additive white noise is at best an intellectual construction. All that he will retain from the above is the criterion

$$\lambda \mathcal{E}(x) + \|y - x\|^2,$$

with the interpretation that the penalty term $\lambda \mathcal{E}(x)$ corrects undesirable features of the observed image y . One of these is the usually chaotic aspect, at the fine scale. However, he does not attempt to interpret this as due to white noise. In order to correct this effect he introduces a *smoothing* penalty term $\mathcal{E}(x)$, which is small when x is smooth, for instance

$$\mathcal{E}(x) = \sum_{(s,t)} (x(s) - x(t))^2,$$

where the summation extends over pairs of adjacent pixels. One disadvantage of this smoothing method is that it will tend to blur the boundary between two highly contrasted regions. One must choose an edge-preserving smoothing penalty function, for instance

$$\mathcal{E}(x) = \sum_{(s,t)} \Psi(x(s) - x(t)),$$

where

$$\Psi(u) = - \left(1 + \left(\frac{u}{\delta} \right)^2 \right)^{-1},$$

with $\delta > 0$. This energy function favors large contrasts and therefore prevents to some extent blurring of the edges. A more sophisticated penalty function would introduce edges, as in Example 4.2, with a penalty function of the form

$$\mathcal{E}(x^P, x^E) = \mathcal{E}_1(x^P, x^E) + \mathcal{E}_2(x^E),$$

where the term $\mathcal{E}_1(x^P, x^E)$ creates the edges from the pixels, and $\mathcal{E}_2(x^E)$ organizes the edges.

Note that the estimated image now consists of two parts: \hat{x}^P solves the smoothing problem, whereas \hat{x}^E extracts the boundaries. If one is really interested in boundary extraction, a sophisticated line-pixel model is desirable; see (Winkler, 1995) for examples. If one is only interested in cleaning the picture, rough models may suffice.

The import of the Gibbs-Bayes approach with respect to the purely deterministic penalty function approach to image restoration lies in the theoretical possibility of the former to tune the penalty function by means of simulation. Indeed, if one is able to produce a typical image corresponding to the energy-penalty function, one will be able to check with the naked eye whether this penalty function respects the constraints one has in mind, and if necessary to adjust the parameters in it. Simulation is the topic of Sections 6 and 7. Another theoretical advantage of the Gibbs-Bayes approach is the availability of the simulated annealing algorithm to solve the minimization problem arising in MAP likelihood method or in the traditional penalty method. This algorithm is essentially stochastic and is based on the notion of Gibbs measure and on the associated sampling methods (see Section 8).

5 Phase Transitions

5.1 Spontaneous Magnetization

Consider the slightly generalized Ising model of a piece of ferromagnetic material, with spins distributed according to

$$\pi_T(x) = \frac{1}{Z_T} e^{-\frac{\mathcal{E}(x)}{T}}. \quad (5.1)$$

The finite site space is enumerated as $S = \{1, 2, \dots, N\}$, and therefore a configuration x is denoted by $(x(1), x(2), \dots, x(N))$. The energy function has two terms,

$$\mathcal{E}(x) = \mathcal{E}_0(x) - \frac{H}{k} \sum_{i=1}^N x(i), \quad (5.2)$$

where the term $\mathcal{E}_0(x)$ is assumed symmetric, in the sense that for any configuration x ,

$$\mathcal{E}_0(x) = \mathcal{E}_0(-x). \quad (5.3)$$

The constant H is the external magnetic field. The *magnetic moment* of configuration x is

$$m(x) = \sum_{i=1}^N x(i), \quad (5.4)$$

and the *magnetization* is the average *magnetic moment* per site

$$M(H, T) = \frac{1}{N} \sum_{x \in E} \pi_T(x) m(x). \quad (5.5)$$

In Problem 7.5.2, the reader is invited to check that

$$\frac{\partial M(H, T)}{\partial H} \geq 0. \quad (5.6)$$

Also, it is clear that

$$M(-H, T) = -M(H, T) \quad (5.7)$$

and

$$-1 \leq M(H, T) \leq +1. \quad (5.8)$$

In summary, at fixed temperature T , the magnetization $M(H, T)$ is a nondecreasing odd function of H with values in $[-1, +1]$. Also,

$$M(0, T) = 0, \quad (5.9)$$

since for any configuration x , $m(-x) = -m(x)$, and therefore $\pi_T(-x) = \pi_T(x)$ when $H = 0$. Moreover, the magnetization is an analytic function of H .

However, the experimental results seem to contradict the last two assertions. Indeed, if an iron bar is placed in a strong magnetic field H parallel to the axis, it is completely magnetized with magnetization $M(H, T) = +1$, and if the magnetic field is slowly decreased to 0, the magnetization decreases, but tends to a limit $M(0, T) = M_0 > 0$, in disagreement with (5.9). By symmetry, we therefore have a discontinuity of the magnetization at $H = 0$ (see Fig. 7.5.1a), in contradiction to the theoretical analyticity of the magnetization as a function of H .

This discontinuity is called a *phase transition* by physicists, by analogy with the discontinuity in density at a liquid–gas phase transition. It occurs at room temperature, and if the temperature is increased, the residual, or *spontaneous*, magnetization M_0 decreases until it reaches the value 0 at a certain temperature T_c , called the *critical temperature*. Then for $T > T_c$, the discontinuity at 0 disappears, and the magnetization curve is smooth (see Fig. 7.5.1c). At $T = T_c$, the slope at $H = 0$ is infinite, i.e., the *magnetic susceptibility* is infinite (see Fig. 7.5.1b).

The discrepancy with experience and theory observed below the critical temperature is due to the fact that the experimental results describe a situation at the *thermodynamical limit* $N = \infty$. For fixed but large N the theoretical magnetization curve is analytic, but it presents for all practical purposes the same aspect as in Fig. 7.5.1a.

To summarize the experimental results, it seems that below the critical temperature, the spontaneous magnetization has, when no external magnetic field is applied, two “choices.” We shall now explain this phenomenon within the classical Ising model.

5.2 Peierls’s Argument

Consider the Ising model in the absence of an external field ($H = 0$). The energy of a configuration x is of the form

$$\mathcal{E}(x) = -J \sum_{(s,t)} x(s)x(t), \quad (5.10)$$




Figure 7.5.1. Magnetization and critical temperature

where $\langle s, t \rangle$ represents an unordered pair of neighbors. When the cardinal of the site space S is infinite, the sum (5.10) is not defined for all configurations, and therefore one cannot define the Gibbs distribution π_T on Λ^S by formula (5.1). However, the local specification

$$\pi_T^s(x) = \frac{e^{\beta \sum_{\langle s, t \rangle} x(s)x(t)}}{e^{\beta \sum_{\langle s, t \rangle} x(t)} + e^{-\beta \sum_{\langle s, t \rangle} x(t)}}, \quad (5.11)$$

where β is, up to a factor, the inverse temperature, is well-defined for all configurations and all sites.

In the sequel, we shall repeatedly use abbreviated notation. For instance, if π is the distribution of a random field X under probability P , then $\pi((X(A)))$ denotes $P(X(A) = x(A))$, $\pi(x(0) = +1)$ denotes $P(X(0) = +1)$, etc.

A probability distribution π_T on Λ^S is called a solution of the DLR problem if it admits the local specification (5.11). Here DLR stands for Dobrushin (1965), and Lanford and Ruelle (1969).

When $S = K_N = \mathbb{Z}^2 \cap [-N, +N]^2$, we know that there exists a unique solution, given by (5.1). When $S = \mathbb{Z}^2$, Dobrushin has proven the existence of at least one solution of the DLR problem. One way of constructing a solution is to select an arbitrary configuration z , to construct the unique probability distribution $\pi_T^{(N)}$ on Λ^S such that

$$\pi_T^{(N)}(z(S \setminus K_{N-1})) = 1 \quad (5.12)$$

(the field is frozen at the configuration z outside K_{N-1}) and such that the restriction of $\pi_T^{(N)}$ to K_{N-1} has the required local characteristics, given by (5.11), and then let N tend to infinity. A solution π_T of the DLR problem corresponding to the configuration z selected is obtained as follows. For all configurations x and all *finite* subsets $A \subset S$, the following limit exists:

$$\pi_T(x(A)) = \lim_{N \uparrow \infty} \pi_T^{(N)}(x(A)), \quad (5.13)$$

and moreover, there exists a unique random field X with the local specification (5.11) and such that, for all configurations x and all *finite* subsets $A \subset S$,

$$P(X(A) = x(A)) = \pi_T(x(A)).$$

Note that $\pi_T^{(N)}$ depends on the configuration z only through the restriction of z to the boundary $K_N \setminus K_{N-1}$.

If the DLR problem has more than one solution, we say that a *phase transition* occurs. The method given by Dobrushin (1965) to construct a solution suggests a way of proving phase transition when it occurs. It suffices to select two configurations z_1 and z_2 , and to show that for a given finite subset $A \subset S$, the right-hand side of (5.13) is different for $z = z_1$ and $z = z_2$. It has been proven by Peierls (1936) that for sufficiently small values of the temperature, phase transition occurs. Peierls applied the above program with z_1 being the configuration with all spins positive and z_2 the all negative configuration, and with $A = \{0\}$, where 0 denotes the central site of \mathbb{Z}^2 .

Denote then by $\pi_+^{(N)}$ (resp., $\pi_-^{(N)}$) the restriction to K_N of $\pi_T^{(N)}$ when $z = z_1$ (resp., $z = z_2$). We shall prove that if T is large enough, then $\pi_+^{(N)}(x(0) = -1) < \frac{1}{3}$ for all N . By symmetry, $\pi_-^{(N)}(x(0) = +1) < \frac{1}{3}$, and therefore $\pi_-^{(N)}(x(0) = -1) > \frac{2}{3}$. Passing to the limit $N \uparrow \infty$, we see that $\pi_+(x(0) = -1) < \frac{1}{3}$ and $\pi_-(x(0) = -1) > \frac{2}{3}$, and therefore, the limiting distributions are not identical.

We now proceed to the execution of the above program. For all $x \in \Lambda^{K_N}$,

$$\pi_+^{(N)}(x) = \frac{e^{-2\beta n_o(x)}}{Z_+^{(N)}}, \tag{5.14}$$

where $n_o(x)$ is the number of *odd bounds* in configuration x , that is, the number of cliques (s, t) such that $x(s) \neq x(t)$, and where $Z_+^{(N)}$ is the normalization factor. To obtain (5.14), it suffices to observe that

$$-\sum_{(s,t)} x(s)x(t) = n_o(x) - n_e(x),$$

where $n_e(x)$ is the number of even bounds, and that $n_e(x) = M - n_o(x)$, where M is the total number of pair cliques. Therefore,

$$\mathcal{E}(x) = 2\beta n_o(x) - M,$$

from which (5.14) follows.

Before proceeding to the proof of the announced upper bound for $\pi_+^{(N)}(x(0) = -1)$, a few definitions are needed. Actually, no formal definition will be proposed; instead, the reader is referred to pictures. For instance, Fig. 7.5.2 features *circuits* C of various lengths.

For a given configuration x , $C(x; 0)$ denotes the circuit which is the boundary of the largest connected batch of sites with negative phases, containing site 0. It is a *circuit around* 0. If the phase at the central site is positive, then $C(x; 0)$ is the empty set.




Figure 7.5.2. Circuits in the Ising model

For a given configuration x , denote by \tilde{x} the configuration obtained by reversing all the phases inside circuit $C(x; 0)$. For a given circuit C around 0,

$$\pi_+^{(N)}(C(x; 0) = C) = \frac{\sum_{x: C(x; 0) = C} e^{-2\beta n_0(x)}}{\sum_y e^{-2\beta n_0(y)}}.$$

But

$$\sum_z e^{-2\beta n_0(z)} \geq \sum_{y: C(y; 0) = C} e^{-2\beta n_0(\tilde{y})}$$

(one can always associate to a configuration y such that $C(y; 0) = C$ the configuration $z = \tilde{y}$, and therefore the sum on the right-hand side is a subsum of the left-hand side). Therefore,

$$\pi_+^{(N)}(C(x; 0) = C) \leq \frac{\sum_{x: C(x; 0) = C} e^{-2\beta n_0(x)}}{\sum_{x: C(x; 0) = C} e^{-2\beta n_0(\tilde{x})}}.$$

If x is such that $C(x; 0) = C$, then $n_0(\tilde{x}) = n_0(x) - L$, where L is the length of C , and therefore

$$\pi_+^{(N)}(C(x; 0) = C) \leq e^{-2\beta L}.$$

In particular,

$$\pi_+^{(N)}(x(0) = -1) \leq \sum r(L) e^{-2\beta L},$$

where the latter summation is over all lengths L of circuits around 0, and $r(L)$ is the number of nonempty circuits around 0 of length L . The possible lengths are 4, 6, ..., $2f(N)$, where $f(N) \uparrow \infty$ as $N \uparrow \infty$. In order to bound $r(L)$ from above, observe that a circuit around 0 of length L must have at least one point at a distance smaller than or equal to $\frac{L}{2}$ of the central site 0. There are L^2 ways of selecting such a point, and then at most 4 ways of selecting the segment of C starting from this point, and then at most 3 ways of selecting the next connected segment, and so on, so that

$$r(L) \leq 4L^2 3^L.$$

Therefore,

$$\pi_+^{(N)}(x(0) = -1) \leq \sum_{L=4,6,\dots} 4L^2(3e^{-2\beta})^L.$$

Now, the series $\sum_{L=4,6,\dots} L^2 x^L$ has a radius of convergence not less than 1, and therefore, if $3e^{-\beta}$ is small enough, or equivalently if T is large enough, $\pi_+^{(N)}(x(0) = -1) < \frac{1}{3}$ for all N . \square

We shall not pursue the study of phase transition, mainly because it is an advanced topic requiring additional background in probability. It interests mainly statistical physics. This does not mean that engineering applications are immune from phase transitions. However, the situations where phase transitions show up are usually unsafe and are in general to be avoided, whereas in statistical physics they impose themselves, and the physicist cannot get around them.

An elementary introduction to infinite Gibbs fields in statistical physics is the monograph of (Kinderman and Snell, 1980).

6 Gibbs Sampler

6.1 Simulation of Random Fields

Consider a random field that changes randomly with time. We then have a stochastic process $\{X_n\}_{n \geq 0}$, where

$$X_n = (X_n(s), s \in S)$$

and $X_n(s) \in \Lambda$. The state at time n of this process is a random field on S with phases in Λ , or equivalently, a random variable with values in the state space $E = \Lambda^S$, which for simplicity we assume finite. The stochastic process $\{X_n\}_{n \geq 0}$ will be called a *dynamical random field*.

The purpose of the current section is to show how a given random field with probability distribution

$$\pi(x) = \frac{1}{Z} e^{-\mathcal{E}(x)} \quad (6.1)$$

can arise as the stationary distribution of a field-valued Markov chain. This problem is interesting in the following simulation context.

Suppose there exists an *irreducible aperiodic* Markov chain $\{X_n\}_{n \geq 0}$ with state space $E = \Lambda^S$ and stationary distribution (6.1). By the fundamental result of convergence to steady state (Theorem 2.1, Chapter 4), for any initial distribution,

$$\lim_{n \uparrow \infty} d_V(P(X_n = \cdot), \pi) = 0. \quad (6.2)$$

If one is able to generate a realization of the above HMC, its distribution at a large time n will be close to π , and one will therefore have simulated π . In this context, the transition mechanism of the chain will be called a *simulation algorithm*, or *sampling algorithm*, and

the asymptotic distribution π the *target distribution*, or *simulated distribution*, or *sampled distribution*.

The first problem is that of identifying a chain $\{X_n\}_{n \geq 0}$ with the above properties. The *Gibbs sampler* uses a strictly positive probability distribution $(q_s, s \in S)$ on S , and the transition from $X_n = x$ to $X_{n+1} = y$ is made according to the following rule.

The new state y is obtained from the old state x by changing (or not) the value of the phase at *one site only*. The site s to be changed (or not) at time n is chosen independently of the past with probability q_s . When site s has been selected, the current configuration x is changed into y as follows: $y(S \setminus s) = x(S \setminus s)$, and the new phase $y(s)$ at site s is selected with probability $\pi(y(s) \mid x(S \setminus s))$. Thus, configuration x is changed into $y = (y(s), x(S \setminus s))$ with probability $\pi(y(s) \mid x(S \setminus s))$, according to the local specification at site s . This gives for the nonzero entries of the transition matrix

$$P(X_{n+1} = y \mid X_n = x) = q_s \pi(y(s) \mid x(S \setminus s)) \mathbf{1}_{y(S \setminus s) = x(S \setminus s)}. \quad (6.3)$$

The corresponding chain is irreducible and aperiodic. To prove that π is the stationary distribution, we use the detailed balance test (Theorem 6.1 of Chapter 2). One has to check that for all $x, y \in \Lambda^S$,

$$\pi(x) P(X_{n+1} = y \mid X_n = x) = \pi(y) P(X_{n+1} = x \mid X_n = y),$$

that is, in view of (6.3), for all $s \in S$,

$$\pi(x) q_s \pi(y(s) \mid x(S \setminus s)) = \pi(y) q_s \pi(x(s) \mid x(S \setminus s)).$$

But the last equality is just

$$\pi(x) q_s \frac{\pi(y(s), x(S \setminus s))}{P(X(S \setminus s) = x(S \setminus s))} = \pi(y(s), x(S \setminus s)) q_s \frac{\pi(x)}{P(X(S \setminus s) = x(S \setminus s))}.$$

Example 6.1. Simulation of the Ising Model

The local specification at site s depends only on the local configuration $x(\mathcal{N}_s)$. Note that small neighborhoods speed up computations. Note also that the Gibbs sampler is a natural sampler, in the sense that in a piece of ferromagnetic material, for instance, the spins are randomly changed according to the local specification. When nature decides to update the orientation of a dipole, it does so according to the law of statistical mechanics. It computes the local energy for each of the two possible spins, $E_+ = E(+1, x(\mathcal{N}_s))$ and $E_- = E(-1, x(\mathcal{N}_s))$, and takes the corresponding orientation with a probability proportional to e^{E_+} and e^{E_-} , respectively. \diamond

Example 6.2. Gibbs Sampler for Neural Networks

The framework and notation are those of Examples 1.3 and 2.1, where we obtained the local characteristics

$$\pi_T(x(s) \mid x(\mathcal{N}_s)) = \frac{e^{-\frac{1}{T} \{ \sum_{t \in \mathcal{N}_s^*} (w_{ts} + w_{st}) x(t) - h_s \}} x(s)}{1 + e^{-\frac{1}{T} \{ \sum_{t \in \mathcal{N}_s^*} (w_{ts} + w_{st}) x(t) - h_s \}}}. \quad (6.4)$$

Therefore, at a given time n , the Gibbs sampler examines site s (with probability q_s) and updates its phase. It chooses phase 0 with probability

$$p_0 = \frac{1}{1 + e^{-a/T}},$$

where

$$a = \sum_{t \in \mathcal{N}_s} (w_{ts} + w_{st})x(t) - h_s, \quad (6.5)$$

and it chooses phase 1 with probability $p_1 = 1 - p_0$.

This procedure can also be described in terms of a random *threshold jitter* σ with the cumulative distribution function

$$P(\sigma \leq a) = \frac{e^{-a/T}}{1 + e^{-a/T}}. \quad (6.6)$$

The Gibbs sampler selects phase 0 if

$$\sum_{t \in \mathcal{N}_s} (w_{ts} + w_{st})x(t) < h_s + \sigma, \quad (6.7)$$

and 1 otherwise. Indeed,

$$P(\text{phase 0}) = P\left(\sigma > \sum_{t \in \mathcal{N}_s} (w_{ts} + w_{st})x(t) - h_s\right) = \frac{1}{1 + e^{-\frac{1}{T}\{\sum_{t \in \mathcal{N}_s} (w_{ts} + w_{st})x(t) - h_s\}}}.$$

One can interpret h_s as the *nominal threshold* at site s and $h_s + \sigma$ as the actual (random) threshold. Also the quantity $\sum_{t \in \mathcal{N}_s} (w_{ts} + w_{st})x(t)$ is the input into neuron s . Thus the excitation of neuron s is obtained by comparing its input to a random threshold (see Fig. 7.6.1). \diamond

Example 6.3. Gibbs Sampler for Random Vectors

Clearly, Gibbs sampling applies to any multivariate probability distribution

$$\pi(x(1), \dots, x(N))$$

on a set $E = \Lambda^N$, where Λ is countable (but this restriction is not essential, as we mentioned earlier). This trivial remark is intended to remind us that there are many applications of Gibbs sampling, and more generally of Monte Carlo Markov chain simulation (see Section 7), outside physics or image processing, and especially in statistics.

The basic step of the Gibbs sampler for the multivariate distribution π consists in selecting a coordinate number $i \in [1, N]$, at random, and choosing the new value $y(i)$ of the corresponding coordinate, given the present values $x(1), \dots, x(i-1), x(i+1), \dots, x(N)$ of the other coordinates, with probability

$$\pi(y(i) \mid x(1), \dots, x(i-1), x(i+1), \dots, x(N)).$$

One checks as above that π is the stationary distribution of the corresponding chain. \diamond




Figure 7.6.1. Jitter sampling of a neural network

Example 6.4. Periodic Gibbs Sampler

In practice, the updated sites are not chosen at random, but instead in a well-determined order $s(1), s(2), \dots, s(N)$, where $\{s(i)\}_{1 \leq i \leq N}$ is an enumeration of all the sites of S , called a *scanning policy*. The sites are visited in this order periodically. The state of the random field after the n th sweep is $Z_n = X_{nN}$, where X_k denotes the image before the k th update time. At time k , site $s(k \bmod N)$ is updated to produce the new image X_{k+1} . If $X_k = x$ and $s(k \bmod N) = s$, then $X_{k+1} = (y(s), x(S \setminus s))$ with probability $\pi(y(s) \mid x(S \setminus s))$. The Gibbs distribution π is stationary for $\{X_k\}_{k \geq 0}$, in the sense that if $P(X_k = \cdot) = \pi$, then $P(X_{k+1} = \cdot) = \pi$. In particular, π is a stationary distribution of the irreducible aperiodic Markov chain $\{Z_n\}_{n \geq 0}$, and $\lim_{n \uparrow \infty} P(Z_n = \cdot) = \pi$.

The transition matrix \mathbf{P} of $\{Z_n\}_{n \geq 0}$ is

$$\mathbf{P} = \prod_{k=1}^N \mathbf{P}_{s(k)}, \tag{6.8}$$

where $\mathbf{P}_s = \{p_{xy}^s\}_{x, y \in \Lambda^s}$, and the entry p_{xy}^s of \mathbf{P}_s is nonzero if and only if $y(S \setminus s) = x(S \setminus s)$, and then

$$p_{xy}^s = \frac{e^{-\mathcal{E}(y(s), x(S \setminus s))}}{\sum_{\lambda \in \Lambda} e^{-\mathcal{E}(\lambda, s(S \setminus s))}}. \tag{6.9}$$

◇

6.2 Convergence Rate of the Gibbs Sampler

One can, in principle, use the algebraic tools of Chapter 6 to obtain rates of convergence for the Gibbs sampler.

Example 6.5. Periodic Gibbs Sampler

Expression (6.9) will be used to produce a geometric rate of convergence of the periodic Gibbs sampler, namely,

$$|\mu \mathbf{P}^n - \pi| \leq \frac{1}{2} |\mu - \pi| (1 - e^{-N\Delta})^n \quad (6.10)$$

where

$$\Delta = \sup_{s \in S} \delta_s \quad (6.11)$$

and

$$\delta_s = \sup\{|\mathcal{E}(x) - \mathcal{E}(y)|; x(S \setminus s) = y(S \setminus s)\}. \quad (6.12)$$

The proof of (6.10) uses Theorem 7.2 of Chapter 6, which gives

$$|\mu \mathbf{P}^n - \pi| \leq \frac{1}{2} |\mu - \pi| \delta(\mathbf{P})^n.$$

Using (7.3) of Chapter 6, it follows that for any transition matrix \mathbf{P} on a finite state space E ,

$$\delta(\mathbf{P}) = 1 - \inf_{i,j \in E} \sum_{k \in E} p_{ik} \wedge p_{jk} \leq 1 - |E| \left(\inf_{i,j \in E} p_{ij} \right). \quad (6.13)$$

If we define $m_s(x) = \inf\{\mathcal{E}(y); y(S \setminus s) = x(S \setminus s)\}$, it follows from (6.9) that

$$p_{xy}^s = \frac{\exp\{-\mathcal{E}(y(s), x(S \setminus s)) - m_s(x)\}}{\sum_{z(s) \in \Lambda} \exp\{-\mathcal{E}(z(s), x(S \setminus s)) - m_s(x)\}} \geq \frac{e^{-\delta_s}}{|\Lambda|},$$

and therefore, from (6.8),

$$\min_{x,y \in \Lambda^S} p_{xy} \geq \prod_{k=1}^N \frac{e^{-\delta_{\pi(k)}}}{|\Lambda|} \geq \frac{e^{-N\Delta}}{|\Lambda|^N}.$$

Using (6.13),

$$\delta(\mathbf{P}) \leq 1 - |\Lambda|^N \frac{e^{-N\Delta}}{|\Lambda|^N} = 1 - e^{-N\Delta},$$

and (6.10) follows. \diamond

Convergence in variation of the Gibbs sampler to the target distribution takes place with geometric speed. This is a common feature of all the Monte Carlo Markov chain simulation algorithms, and more generally of all *finite* ergodic HMCs. However, the Gibbs sampler has a special feature. It turns out that the distance in variation to the target distribution *decreases* to 0, as the following result shows.

Theorem 6.1. *Monotonicity Property of the Gibbs Sampler*

Let μ be an arbitrary probability measure on Λ^S and let ν be the probability measure obtained by applying the Gibbs sampler at an arbitrary site $s \in S$. Then $d_V(\nu, \pi) \leq d_V(\mu, \pi)$.

Proof. Take $s = 1$ for definiteness. We have

$$\begin{aligned}
 \sum_x |v(x) - \pi(x)| &= \sum_x |\pi(x(1)|x(S \setminus 1))\mu(x(S \setminus 1)) - \pi(x)| \\
 &= \sum_x |\pi(x(1)|x(S \setminus 1))[\mu(x(S \setminus 1)) - \pi(x(S \setminus 1))]| \\
 &= \sum_{x(S \setminus 1)} \left\{ \sum_{x(1)} \pi(x(1)|x(S \setminus 1)) \right\} |\mu(x(S \setminus 1)) - \pi(x(S \setminus 1))| \\
 &= \sum_{x(S \setminus 1)} |\mu(x(S \setminus 1)) - \pi(x(S \setminus 1))| \\
 &= \sum_{x(S \setminus 1)} \left| \sum_{x(1)} (\mu(x) - \pi(x)) \right| \\
 &\leq \sum_x |\mu(x) - \pi(x)|. \quad \square
 \end{aligned}$$

It is quite unlikely that a good simulation algorithm chooses a short-sighted strategy in which the distance in variation is decreased at every step. It is an experimental fact that the Gibbs sampler is not the best Monte Carlo Markov chain simulation algorithm.

7 Monte Carlo Markov Chain Simulation

7.1 General Principle

We shall now take a more general approach to simulation.

In view of evaluating the expectation $E[\phi(Z)]$ of a random vector Z of dimension k with a probability density $f(x)$, where $\phi: \mathbb{R}^k \rightarrow \mathbb{R}$ and $E[|\phi(Z)|] < \infty$, the formula

$$E[\phi(Z)] = \int_{\mathbb{R}^n} \phi(x) f(x) dx$$

can be used if one is able to compute the integral analytically, which is seldom the case. A first alternative is numerical integration. A second alternative is to generate a sequence of i.i.d random vectors $\{Z_n\}_{n \geq 1}$ with the same distribution as Z , and to invoke the strong law of large numbers

$$E[\phi(Z)] = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \phi(Z_i)$$

to obtain an estimate based on a finite sample (Z_1, \dots, Z_n) , namely $\frac{1}{n} \sum_{i=1}^n \phi(Z_i)$. The choice of the sample size n can be made via the theory of confidence intervals in order to attain a given accuracy; see, for instance, (Ross, 1987), (Ripley, 1987), or (Fishman, 1996).

This method, known as *Monte Carlo*, must be compared with numerical integration. It is a fact that Monte Carlo simulation is competitive with numerical integration for large

dimension k . In some applications, the Monte Carlo method is unavoidable. One of them occurs when Z_n is a complicated functional of the n th cycle of a regenerative process.

Another situation is that in which the probability density $f(x)$ is known only up to a normalizing factor, that is, $f(x) = K\tilde{f}(x)$, and when, of course, the integral of $\tilde{f}(x)$ that gives the normalizing factor is difficult, or impossible, to compute. In physics, this is a frequent case: The partition function of a Gibbs distribution is usually uncomputable in closed form, due to the size of the state space in the corresponding models.

A third situation is that in which the standard methods of random-variable generation that we shall introduce next are not applicable, because of round-off errors, or coding problems (see the discussion after Example 7.2).

Method of the Inverse and Acceptance-Rejection

Given that one has to resort to Monte Carlo simulation, the problem is that of generating replicas of Z . The *inverse distribution* method uses a sequence of i.i.d. random variables $\{U_n\}_{n \geq 1}$ uniformly distributed on $[0,1]$, and generates $\{Z_n\}_{n \geq 1}$ by

$$Z_n = F^{-1}(U_n),$$

where $F^{-1}(u)$ is the inverse function of $u = F(x) = \int_{-\infty}^x f(y)dy$ (we consider the one-dimensional case, but this is not essential in the discussion). (See Fig. 7.7.1.) The scope of this method is limited by the fact that $F^{-1}(u)$ is often difficult to obtain, and in any case the inverse function method requires full knowledge of the probability density function $f(x)$.




Figure 7.7.1. The inverse method

Example 7.1. *Exponential Random Variable*

The inverse function of $F(x) = 1 - e^{-\theta x}$ is $F^{-1}(y) = -\frac{1}{\theta} \ln(1 - y)$. Therefore, if U is uniformly distributed on $[0, 1]$, then $\frac{1}{\theta} \ln(1 - U)$ is an exponential random variable with mean θ^{-1} . So is

$$Z = \frac{1}{\theta} \ln(U),$$

since U and $1 - U$ are identically distributed. ◇

Another method, called *acceptance–rejection* (A–R for short), brings us nearer to an efficient solution. It is, in fact, the ancestor of the *Monte Carlo Markov chain* (MCMC) method.

Suppose one can generate a sequence of i.i.d. random vectors $\{Y_n\}_{n \geq 1}$ with the probability density $g(x)$ satisfying, for all x ,

$$\frac{f(x)}{g(x)} \leq c$$

for a finite constant c . Let $\{U_n\}_{n \geq 1}$ be a sequence of i.i.d. random variables uniformly distributed on $[0, 1]$. If we define τ to be the first index $n \geq 1$ for which

$$U_n \leq \frac{f(Y_n)}{cg(Y_n)}$$

and let

$$Z = Y_\tau,$$

then (Problem 7.7.1) Z admits the probability density function $f(x)$ and

$$E[\tau] = c.$$

Example 7.2.

Fig. 7.7.2 summarizes the operations in the A–R method in an artificial example (where the A–R method is certainly not needed!). Here the target density $f(x)$ is a “triangle” of base $[0, 1]$ and summit at $(\frac{1}{2}, 2)$, and $g(x)$ is the uniform distribution on $[0, 1]$. We take the smallest possible c , here $c = 2$. The average number of tests to obtain a sample is therefore $c = 2$. ◇

Both the inverse method and the acceptance–rejection method have obvious counterparts when Z is a discrete random variable with values in a finite space $E = \{1, 2, \dots, r\}$. We then denote by π the distribution of Z . The inverse method is in this case always theoretically feasible: It consists in generating a random variable U uniformly distributed on $[0, 1]$ and letting $Z = i$ if and only if $\sum_{\ell=1}^{i-1} \pi(\ell) \leq U < \sum_{\ell=1}^i \pi(\ell)$. When the size r of the state space E is large, problems arise that are due to the small size of the intervals partitioning $[0, 1]$ and to the cost of precision in computing. In random field simulation, another, maybe more

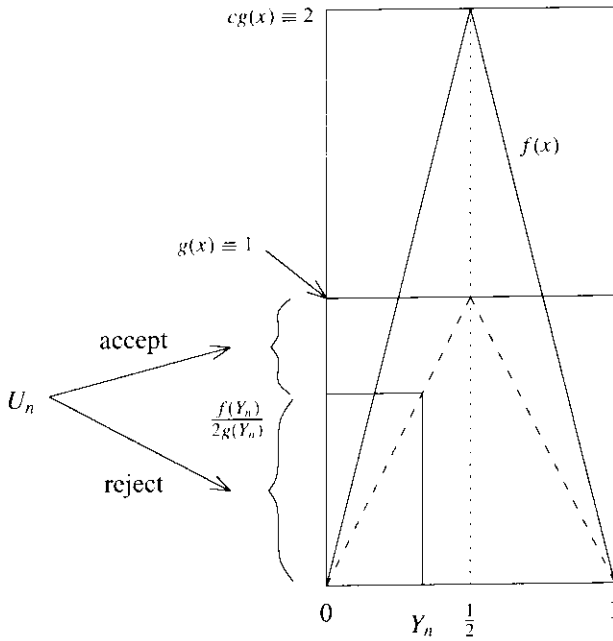


Figure 7.7.2. The acceptance-rejection method

important, reason is the necessity to enumerate the configurations, which implies coding and decoding of a mapping from the integers to the configuration space.

The A-R method uses a distribution p on E such that $\frac{\pi(i)}{p(i)} \leq c < \infty$ for all $i \in E$. The main problem is that π is often known only up to a normalizing factor (the inverse method suffers from the same ailment).

MCMC

In a number of applications, and especially in the simulation of large random fields occurring in statistical physics, all the above problems arise, and the quest for a random-variable generator without these problems is at the origin of the Monte Carlo Markov chain method. We have already seen in Section 6 the basic methodology:

One constructs an irreducible aperiodic HMC $\{X_n\}_{n \geq 0}$ with state space E with the stationary distribution π . Since E is finite, the chain is ergodic, and for any initial distribution μ and all $i \in E$,

$$\lim_{n \rightarrow \infty} P_\mu(X_n = i) = \pi(i) \tag{7.1}$$

and

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \phi(X_n) = E_\pi[\phi(X)]. \tag{7.2}$$

By (7.1), when n is “large,” we can consider that X_n has a distribution close to π . Of course, one would like to know how accurately X_n imitates an E -valued random variable Z with distribution π . For this we shall seek to obtain estimates of the form

$$|\mu \mathbf{P}^n - \pi| \leq A\alpha^n,$$

where $\alpha < 1$. The Perron–Frobenius theorem gives a theoretical answer, in terms of the second-largest eigenvalue modulus (SLEM). To make it applicable, it is necessary to find good upper bounds of the SLEM, since the eigenstructure of a transition matrix is often very hard to obtain exactly. We have given results in this direction in Chapter 6. They concern mainly reversible chains, and fortunately, this is the most frequent case in MCMC, as we shall now see.

For the time being, the basic problem is that of designing the MCMC algorithm. One must find an ergodic transition matrix \mathbf{P} on E , the stationary distribution of which π is the *target distribution*. There are infinitely many such transition matrices, and among them there are infinitely many that correspond to a reversible chain, that is, such that

$$\pi(i)p_{ij} = \pi(j)p_{ji}. \quad (7.3)$$

We seek solutions of the form

$$p_{ij} = q_{ij}\alpha_{ij} \quad (7.4)$$

for $j \neq i$, where $Q = \{q_{ij}\}_{i,j \in E}$ is an arbitrary irreducible transition matrix on E , called the *candidate-generating matrix*: When the present state is i , the next *tentative* state j is chosen with probability q_{ij} . When $j \neq i$, this new state is accepted with probability α_{ij} . Otherwise, the next state is the same state i . Hence, the resulting probability of moving from i to j when $i \neq j$ is given by (7.4). It remains to select the *acceptance* probabilities α_{ij} . A quite general form of the acceptance probabilities covering most well-known MCMC algorithms is due to Hastings (1970); see (Fishman, 1996) for additional bibliography.

$$\alpha_{ij} = \frac{s_{ij}}{1 + t_{ij}}, \quad (7.5)$$

where $\Sigma = \{s_{ij}\}_{i,j \in E}$ is a *symmetric* matrix and

$$t_{ij} = \frac{\pi(i)q_{ij}}{\pi(j)q_{ji}}. \quad (7.6)$$

Of course, Σ must be selected in such a way that the constraint $\alpha_{ij} \in [0, 1]$ is respected. One can check that the reversibility condition (7.3) is satisfied, and therefore π is the stationary distribution, by the detailed balance test (Corollary 6.1 of Chapter 2).

The particular forms of the Hastings algorithm that follow are due to Metropolis et alii (1953) and Barker (1965), respectively.

Example 7.3. *Metropolis Algorithm*

In order to satisfy the constraint $\alpha_{ij} \in [0, 1]$, one must have

$$s_{ij} \leq 1 + \min(t_{ij}, t_{ji})$$

(recall that Σ is assumed symmetric). Equality corresponds to the Metropolis algorithm

$$\alpha_{ij} = \min \left(1, \frac{\pi(j)q_{ji}}{\pi(i)q_{ij}} \right). \quad (7.7)$$

A special case is that for which the candidate-generation mechanism is purely random, that is, $q_{ij} = \text{constant}$. Then

$$\alpha_{ij} = \min \left(1, \frac{\pi(j)}{\pi(i)} \right). \quad \diamond$$

Example 7.4. Barker's Algorithm

With the choice $s_{ij} = 1$, we have Barker's algorithm, for which

$$\alpha_{ij} = \frac{\pi(j)q_{ji}}{\pi(j)q_{ji} + \pi(i)q_{ij}}. \quad (7.8)$$

In the special case of a purely random selection of the candidate,

$$\alpha_{ij} = \frac{\pi(j)}{\pi(i) + \pi(j)}.$$

As we shall see later, Barker's algorithm is connected to the Gibbs sampler. \diamond

It is important to observe that Hastings algorithms require the knowledge of the target distribution π only up to a normalizing constant, since it depends only on the ratios $\pi(j)/\pi(i)$. The latter statement is true only as long as the candidate-generating matrix Q is known. This is *not* the case in the following example.

Example 7.5. Independent Sampling

With $q_{ij} = \pi(j)$ and $\alpha_{ij} = 1$, for all $i, j \in E$ we have independent sampling, that is, $p_{ij} = \pi(j)$ for all $i, j \in E$. \diamond

7.2 Convergence Rates in MCMC

If a Monte Carlo simulation algorithm is stopped at the n th iteration, the closeness to the target distribution of the sample obtained is measured by the distance in variation $d_V(\delta_i \mathbf{P}^n, \pi)$, where i is the initial state. Upper bounds on this performance index can be in principle derived with the help of the results of Chapter 6. For instance, assuming that (\mathbf{P}, π) is reversible, Theorem 3.3 of Chapter 6 gives

$$4d_V(\delta_i \mathbf{P}^n, \pi)^2 \leq \frac{\rho^{2n}}{\pi(i)}, \quad (7.9)$$

where ρ , the SLEM of \mathbf{P} , can be bounded above using any of the techniques (Poincaré, Cheeger, ...) explained in Chapter 6.

Example 7.6.

(Diaconis and Saloff-Coste (1996)) Let π be the distribution on $E = \{1, \dots, r\}$ given by

$$\pi(i) = z(a)a^{h(i)},$$

where $a \in (0, 1)$, and where for $i \in [1, r - 1]$,

$$h(i + 1) - h(i) \geq c \geq 1. \quad (7.10)$$

We use the Metropolis algorithm with acceptance probability

$$\alpha_{ij} = \min\left(1, \frac{\pi(j)q_{ij}}{\pi(i)q_{ji}}\right),$$

where the candidate-generating matrix $Q = \{q_{ij}\}_{i,j \in E}$ corresponds to a symmetric random walk on E with holding probability $\frac{1}{2}$ at the extreme states 1 and r . Therefore, the simulation matrix \mathbf{P} is

$$p_{11} = 1 - \frac{a^{h(2)-h(1)}}{2}, \quad p_{12} = \frac{a^{h(2)-h(1)}}{2}, \quad p_{r,r-1} = p_{rr} = \frac{1}{2}$$

and for $i \in [2, r - 1]$

$$p_{i,i-1} = \frac{1}{2}, \quad p_{i,i+1} = \frac{a^{h(i+1)-h(i)}}{2}, \quad p_{ii} = 1 - p_{i,i-1} - p_{i,i+1}.$$

We are going to show that

$$\lambda_2 \leq 1 - \frac{(1 - a^{c/2})^2}{2}.$$

According to (2.18) and (2.19) of Chapter 6, it suffices to show that for all $x \in \mathbb{R}^r$,

$$\text{Var}_\pi(x) \leq A \mathcal{E}_\pi(x, x)$$

for some A such that

$$A \leq \frac{2}{(1 - a^{c/2})^2}. \quad (7.11)$$

For this purpose, one can use an adaptation of the Diaconis-Stroock geometric method. The notations are those of Section 4 of Chapter 6. Let $\theta \in (0, 1)$. Proceeding as in the proof of Theorem 4.1 of Chapter 6, with $Q(e)^\theta$ replacing $Q(e)^{\frac{1}{2}}$,

$$\begin{aligned} 2\text{Var}_\pi(x) &= \sum_{i,j \in E} \left\{ \sum_{e \in \gamma_{ij}} \frac{1}{Q(e)^\theta} Q(e)^\theta (x(e^-) - x(e^+)) \right\}^2 \pi(i)\pi(j) \\ &\leq \sum_{i,j \in E} \left\{ \sum_{e \in \gamma_{ij}} Q(e)^{2\theta} (x(e^-) - x(e^+))^2 \right\} \left\{ \sum_{e \in \gamma_{ij}} \frac{1}{Q(e)^{2\theta}} \right\} \pi(i)\pi(j), \end{aligned}$$

where the last inequality is Cauchy-Schwarz.

Therefore, defining

$$|\gamma_{ij}|_\theta = \sum_{e \in \gamma_{ij}} \frac{1}{Q(e)^{2\theta}}, \tag{7.12}$$

we have

$$\begin{aligned} 2\text{Var}_\pi(x) &\leq \sum_{i,j \in E} |\gamma_{ij}|_\theta \sum_{e \in \gamma_{ij}} Q(e)^{2\theta} (x(e^-) - x(e^+))^2 \pi(i)\pi(j) \\ &= \sum_e (x(e^-) - x(e^+))^2 Q(e)Q(e)^{2\theta-1} \sum_{\gamma_{ij} \ni e} \pi(i)\pi(j) |\gamma_{ij}|_\theta. \end{aligned}$$

With

$$A = \max_e \left\{ Q(e)^{2\theta-1} \sum_{\gamma_{ij} \ni e} \pi(i)\pi(j) |\gamma_{ij}|_\theta \right\}, \tag{7.13}$$

we finally have

$$\text{Var}_\pi(x) \leq A \mathcal{E}_\pi x, x).$$

It now remains to bound A as in (7.11), making use of the freedom in the choice of the γ_{ij} 's and of θ . First, take $\gamma_{ij} = (i, i + 1, \dots, j - 1, j)$ for all $i, j \in E, i \geq j$.

For the specific algorithm considered,

$$Q(i, i + 1) = \pi(i)p_{i,i+1} = z(a)a^{h(i)} \frac{a^{h(i+1)-h(i)}}{2} = z(a) \frac{a^{h(i+1)}}{2} = \frac{\pi(i + 1)}{2},$$

and therefore by reversibility,

$$Q(i, i + 1) = Q(i + 1, i) = \frac{\pi(i + 1)}{2}.$$

The dominating term in $|\gamma_{ij}|_\theta$ is $Q(j - 1, j)^{-2\theta} = \left(\frac{\pi(j)}{2}\right)^{-2\theta}$. Put it in factor to obtain, using (7.10),

$$|\gamma_{ij}|_\theta = \left(\left(\frac{\pi(i + 1)}{\pi(j)}\right)^{-2\theta} + \dots + \left(\frac{\pi(j)}{\pi(j)}\right)^{-2\theta} \right) \left(\frac{\pi(j)}{2}\right)^{-2\theta} \leq \frac{\pi(j)^{-2\theta}}{1 - a^{2c\theta}}.$$

Let $e = (k, k + 1)$. We have to bound (see (7.13))

$$\frac{Q(e)^{2\theta} - 1}{1 - a^{2c\theta}} \sum_{0 \leq i \leq k, k+1 \leq j \leq n} \pi(i)\pi(j)^{1-2\theta}.$$

The sum in j is bounded by $\frac{\pi(k+1)^{1-2\theta}}{1 - a^{c(1-2\theta)}}$, and the sum in i by 1. Therefore,

$$A \leq \frac{2}{1 - a^{c(1-2\theta)}(1 - a^{2c\theta})},$$

from which (7.11) follows by the choice $\theta = \frac{1}{4}$. The Gershgorin bound (see Theorem 2.3 of the Appendix) gives, for the smallest eigenvalue λ_r ,

$$\lambda_r \geq -1 + 2 \min p_{ii} \geq -1 + 2 \left(\frac{1}{2} - \frac{a^c}{2} \right) = a^c,$$

and therefore,

$$\rho = \min(\lambda_2, |\lambda_r|) \leq \min\left(1 - \frac{(1 - a^{c/2})}{2}, a^c\right). \quad \diamond$$

As we have already observed, upper bounds are a necessary recourse, since the eigenstructure of a large transition matrix, even a reversible one, is usually out of reach. There is, however, a notable exception concerning the very important Metropolis algorithm.

Theorem 7.1. *Eigenstructure of Metropolis*

(Liu (1995)) Let π and p be two strictly positive probability distributions on $E = \{1, 2, \dots, r\}$, and define

$$w(i) = \frac{\pi(i)}{p(i)}.$$

The Metropolis algorithm corresponding to the candidate-generating matrix Q given by $q_{ij} = p_j$ for all $i, j \in E$ has the transition matrix \mathbf{P} given by

$$p_{ij} = p(j) \min\left(1, \frac{w(j)}{w(i)}\right),$$

for $i \neq j$. Assume that the states of E are ordered in such a way that

$$w(1) \geq w(2) \geq \dots \geq w(r).$$

The eigenvalues λ_k , $k \in [1, r]$, and the corresponding right eigenvectors v_k , $k \in [1, r]$, of \mathbf{P} are $\lambda_1 = 1$, $v_1 = \mathbf{1}$, and for $k \geq 1$,

$$\lambda_{k+1} = \sum_{j \geq k} \pi(j) \left(\frac{1}{w(j)} - \frac{1}{w(k)} \right),$$

$$v_{k+1} = \left(0, \dots, 0, \sum_{\ell=k+1}^r \pi(\ell), -\pi(k), \dots, -\pi(k) \right)^T,$$

where the first $k - 1$ entries of v_{k+1} are null.

Proof. By inspection (Problem 7.7.2) □

Example 7.7.

(Diaconis and Saloff-Coste (1996)) Define the probability distribution π on $E = \{1, \dots, r\}$ by

$$\pi(j) = \frac{\theta^{j-1}(1 - \theta)}{1 - \theta^r}$$

where $\theta \in (0, 1)$. Generate candidates randomly, i.e., $p(j) = \frac{1}{r}$. Here the states are naturally ordered as in Theorem 7.1, $w(j) = r\pi(j)$ and

$$\lambda_2 = \sum_{j=1}^r \left(\frac{1}{r} - \frac{\pi(j)}{r\pi(1)} \right) = 1 - \frac{1}{r} \frac{1 - \theta^r}{1 - \theta}.$$

Since in this algorithm $\rho = \lambda_2$ and all the eigenvalues are nonnegative, the upper bound (7.9) gives, after some pertinent bounding,

$$4d_V(\delta_r \mathbf{P}^n, \pi)^2 \leq \theta^{-r} \left(1 - \frac{2}{r} \right)^{2n}. \quad \diamond$$

7.3 Variance of Monte Carlo Estimators

The previous subsection concerned the problem of assessing the quality of MCMC in terms of the closeness of its output from the target distribution. We now consider the problem of evaluating expectations with respect to the target distribution by ergodic estimates.

In Theorem 6.5 of Chapter 6, we obtained the formula

$$v(f, \mathbf{P}, \pi) = 2 \langle f, \mathbf{Z}f \rangle_\pi - \langle f, (I + \Pi)f \rangle_\pi \quad (7.14)$$

giving the asymptotic variance

$$v(f, \mathbf{P}, \pi) = \lim_{n \rightarrow \infty} \frac{1}{n} \text{Var}_\mu \left(\sum_{k=1}^n f(X_k) \right). \quad (7.15)$$

Here $\{X_n\}$ is an ergodic HMC with finite state space E , transition matrix \mathbf{P} , and stationary distribution π , and

$$\mathbf{Z} = (I - \mathbf{P} + \Pi)^{-1}, \quad (7.16)$$

where $\Pi = \mathbf{1} \cdot \pi^T$ is the fundamental matrix.

In this section, we consider reversible transition matrices, such as those corresponding to the MCMC simulation algorithms proposed by Hastings. We are interested in designing the best simulation algorithm in the sense that $v(f, \mathbf{P}, \pi)$ is to be minimized with respect to \mathbf{P} , uniformly in f , and of course for a fixed π . The following result of Peskun (1973) answers the question in general terms.

Theorem 7.2.

Let \mathbf{P}_1 and \mathbf{P}_2 be reversible ergodic transition matrices on the finite state space E , with the same stationary distribution π . If \mathbf{P}_1 has all its off-diagonal terms greater than or equal to the corresponding off-diagonal terms of \mathbf{P}_2 , then

$$v(f, \mathbf{P}_1, \pi) \leq v(f, \mathbf{P}_2, \pi)$$

for all $f : E \rightarrow \mathbb{R}$.

Proof. Let $k, \ell \in E$ with $k \neq \ell$. From (7.14) we have

$$\frac{\partial}{\partial p_{k\ell}} v(f, \mathbf{P}, \pi) = 2 \left\langle f, \frac{\partial \mathbf{Z}}{\partial p_{k\ell}} f \right\rangle_{\pi}.$$

From $\mathbf{Z}\mathbf{Z}^{-1} = I$, it follows that $\left(\frac{\partial}{\partial p_{k\ell}} \mathbf{Z}\right) \mathbf{Z}^{-1} + \mathbf{Z} \left(\frac{\partial}{\partial p_{k\ell}} \mathbf{Z}^{-1}\right) = 0$, and therefore

$$\frac{\partial \mathbf{Z}}{\partial p_{k\ell}} = -\mathbf{Z} \frac{\partial \mathbf{Z}^{-1}}{\partial p_{k\ell}} \mathbf{Z},$$

so that

$$\frac{\partial}{\partial p_{k\ell}} v(f, \mathbf{P}, \pi) = -2 \left\langle f, \left(\mathbf{Z} \frac{\partial \mathbf{Z}^{-1}}{\partial p_{k\ell}} \mathbf{Z} \right) f \right\rangle_{\pi}.$$

Since \mathbf{P} is autoadjoint in $\ell^2(\pi)$, so is \mathbf{Z} , and therefore

$$\frac{\partial}{\partial p_{k\ell}} v(f, \mathbf{P}, \pi) = -2 \left\langle \mathbf{Z}f, \left(\frac{\partial \mathbf{Z}^{-1}}{\partial p_{k\ell}} \right) \mathbf{Z}f \right\rangle_{\pi} = -2(\mathbf{Z}f)^T d(\Pi) \frac{\partial \mathbf{Z}^{-1}}{\partial p_{k\ell}} \mathbf{Z}f.$$

Now, from (7.16),

$$\frac{\partial \mathbf{Z}^{-1}}{\partial p_{k\ell}} = -\frac{\partial \mathbf{P}}{\partial p_{k\ell}},$$

and therefore

$$\frac{\partial}{\partial p_{k\ell}} v(f, \mathbf{P}, \pi) = 2(\mathbf{Z}f)^T d(\Pi) \frac{\partial \mathbf{P}}{\partial p_{k\ell}} \mathbf{Z}f.$$

Observe that since \mathbf{P} is a stochastic matrix and (\mathbf{P}, π) is reversible, the free parameters are $(p_{k\ell}; k < \ell)$. In view of the reversibility condition, the only nonnull elements of $d(\Pi) \frac{\partial \mathbf{P}}{\partial p_{k\ell}}$ are the (ℓ, ℓ) , (ℓ, k) , (k, ℓ) , and (k, k) elements, respectively equal to $-\pi(k)$, $+\pi(k)$, $+\pi(k)$, and $-\pi(k)$. Therefore, $d(\Pi) \frac{\partial \mathbf{P}}{\partial p_{k\ell}}$ is a negative definite symmetric matrix, and

$$\frac{\partial}{\partial p_{k\ell}} v(f, \mathbf{P}, \pi) \leq 0,$$

from which the conclusion follows. \square

Example 7.8. Optimality of the Metropolis Algorithm

In the Hastings algorithms

$$p_{ij} = q_{ij} \frac{s_{ij}}{1 + t_{ij}},$$

where t_{ij} depends on $Q = \{q_{ij}\}$ and π only. We would like to find the best MCMC algorithm in the Hastings class where Q is fixed. We have observed that from the constraints $\leq \alpha_{ij} \in (0, 1)$ and the required symmetry of $\{s_{ij}\}$,

$$s_{ij} \leq 1 + \min(t_{ij}, t_{ji}),$$

with equality for the Metropolis algorithm. It follows from Peskun's result that the Metropolis algorithm is optimal with respect to asymptotic variance in the class of Hastings algorithms with fixed candidate-generating matrix Q . \diamond

It is interesting to compare a given MCMC algorithm corresponding to a reversible pair (\mathbf{P}, π) to independent sampling for which $\mathbf{P} = \pi$ (see Example 1.3). From the variance point of view, it follows from (7.14) that an MCMC algorithm based on \mathbf{P} performs better than independent sampling uniformly in f if and only if

$$\langle f, \mathbf{Z}f \rangle_{\pi} \leq \langle f, f \rangle_{\pi} \quad (7.17)$$

for all $f : E \rightarrow \mathbb{R}$.

From (7.14), $\langle f, \mathbf{Z}f \rangle_{\pi} \geq 0$ for all f , and we have already observed that Z is self-adjoint in $\ell^2(\pi)$ (its inverse is self-adjoint). Therefore its eigenvalues are real and nonnegative. Condition (7.17) is equivalent to the fact that these eigenvalues are less than or equal to 1. Therefore, in view of (7.16), (7.17) is equivalent to $\mathbf{P} - \Pi$ having all its characteristic roots nonpositive.

Example 7.9. Barker vs. Independent Sampling

The trace of a matrix is by definition the sum of its diagonal elements. For a stochastic matrix it is therefore the sum of its elements minus the sum of its off-diagonal elements. In particular, $\text{tr}(\mathbf{P}) = r - \sum_{i>j} (p_{ij} + p_{ji})$. Since $\text{tr}(\Pi) = 1$, we have

$$\text{tr}(\mathbf{P} - \Pi) = r - 1 - \sum_{i>j} (p_{ij} + p_{ji}).$$

One can verify that for Barker's algorithm

$$\min(q_{ij}, q_{ji}) \leq p_{ij} + p_{ji} \leq \max(q_{ij}, q_{ji})$$

with equality if Q is symmetric. Therefore, in the case where Q is symmetric,

$$\text{tr}(\mathbf{P} - \Pi) = r - 1 + \sum_{i>j} q_{ij} \geq \frac{1}{2}(r - 2).$$

Thus, if $r \geq 2$, the sum of the characteristic roots of $\mathbf{P} - \Pi$ is positive, which implies that at least one characteristic root is positive.

Therefore Barker's algorithm, is not uniformly better than independent sampling. This does not mean that Barker's algorithm cannot perform better than independent sampling for a specific f . Moreover, and more importantly, the fact that an MCMC algorithm performs not as well as independent sampling is not too alarming, since MCMC algorithms are used when independent sampling cannot be implemented. \diamond

We now give a lower bound for the asymptotic variance of any MCMC estimator. Let (\mathbf{P}, π) be a reversible pair, where \mathbf{P} is irreducible. Its r (real) eigenvalues are ordered as follows:

$$\lambda_1 = 1 > \lambda_2 \geq \lambda_3 \geq \dots \geq \lambda_r \geq -1.$$

For a given f , the formula

$$v(f, \mathbf{P}, \pi) = \sum_{j=1}^r \frac{1 + \lambda_j}{1 - \lambda_j} |\langle f, v_j \rangle_\pi|^2$$

obtained in Section 6.3 of Chapter 6 fully accounts for the interaction between f and \mathbf{P} , in terms of the asymptotic variance of the ergodic estimate of $\langle f \rangle_\pi$. Since the function $x \rightarrow \frac{1+x}{1-x}$ is increasing in $(0, 1]$, and λ_2 is the second largest eigenvalue of \mathbf{P} , the worst (maximal) value of the performance index

$$\gamma(f, \mathbf{P}, \pi) = \frac{v(f, \mathbf{P}, \pi)}{\text{Var}_\pi(f)} = \frac{\sum_{j=2}^r \frac{1+\lambda_j}{1-\lambda_j} |\langle f, v_j \rangle_\pi|^2}{\sum_{j=2}^r |\langle f, v_j \rangle_\pi|^2} \quad (7.18)$$

is attained for $f = v_2$, and is then equal to

$$\gamma(\mathbf{P}, \pi) = \frac{1 + \lambda_2}{1 - \lambda_2}. \quad (7.19)$$

Let $M(\pi)$ be the collection of irreducible transition matrices \mathbf{P} such that the pair (\mathbf{P}, π) is reversible, and denote by $\lambda_2(\mathbf{P})$ the second largest eigenvalue of \mathbf{P} . Assume that

$$\pi(1) \leq \pi(2) \leq \dots \leq \pi(r).$$

In particular, $0 < \pi(1) \leq \frac{1}{2}$.

The result below is due to Frigessi, Hwang and Younès (1992).

Theorem 7.3. *Best Worst Asymptotic Variance*

Let $\mathbf{P} \in M(\pi)$. Then

$$\lambda_2(\mathbf{P}) \geq -\frac{\pi(1)}{1 - \pi(1)}, \quad (7.20)$$

and the bound is attained for some $\mathbf{P} \in M(\pi)$. In particular,

$$\inf_{\mathbf{P} \in M(\pi)} \sup_{f \neq 0} \gamma(f, \mathbf{P}, \pi) = 1 - 2\pi(1). \quad (7.21)$$

Proof. Clearly, (7.21) follows from (7.19) and the first part of the theorem.

To prove (7.20), we use Rayleigh's representation

$$\lambda_2 = \sup \left\{ \frac{\langle \mathbf{P}f, f \rangle_\pi}{\|f\|_\pi^2}; \langle f \rangle_\pi = 0, f \neq 0 \right\}$$

and exhibit a vector $f \neq 0$ such that $\langle f \rangle_\pi = 0$ and

$$\frac{\langle \mathbf{P}f, f \rangle_\pi}{\|f\|_\pi^2} = -\frac{\pi(1)}{1 - \pi(1)}.$$

We try

$$v_2 = \delta_1 - \langle \delta_1 \rangle_\pi \mathbf{1} = (1 - \pi(1), -\pi(1), \dots, -\pi(1))^T, \quad (7.22)$$

where δ_i has zero entries except for the i th one, which is equal to 1. Using the reversibility of (\mathbf{P}, π) and the fact that \mathbf{P} is a stochastic matrix, the i -th coordinate of $\mathbf{P}v_2$ is

$$\begin{aligned} p_{i1}(1 - \pi(1)) - \sum_{j=2}^r p_{ij}\pi(1) &= p_{i1}(1 - \pi(1)) - (1 - p_{i1})\pi(1) \\ &= p_{i1} - \pi(1) = \frac{\pi(1)}{\pi(i)} p_{1i} - \pi(1). \end{aligned}$$

Therefore,

$$\begin{aligned} \langle \mathbf{P}v_2, v_2 \rangle_\pi &= (p_{11} - \pi(1))(1 - \pi(1))\pi(1) - \sum_{j=2}^r \left(\frac{\pi(1)}{\pi(j)} p_{1j} - \pi(1) \right) \pi(1)\pi(j) \\ &= (p_{11} - \pi(1))\pi(1). \end{aligned}$$

Also,

$$\|v_2\|_\pi^2 = \pi(1)(1 - \pi(1)).$$

Therefore,

$$\frac{\langle \mathbf{P}v_2, v_2 \rangle_\pi}{\|v_2\|_\pi^2} = \frac{p_{11} - \pi(1)}{1 - \pi(1)} \geq -\frac{\pi(1)}{1 - \pi(1)}.$$

This proves (7.20).

We must now prove that the bound is attained, and we do this by explicitly constructing \mathbf{P} such that $\lambda_2(\mathbf{P}) = -\frac{\pi(1)}{1-\pi(1)}$. In view of the last display, this implies that

$$p_{11} = 0. \quad (7.23)$$

If $\lambda_2(\mathbf{P}) = -\frac{\pi(1)}{1-\pi(1)}$ (and for this (7.23) must hold), we necessarily have

$$\frac{\langle \mathbf{P}v_2, v_2 \rangle_\pi}{\|v_2\|_\pi^2} = \lambda_2(\mathbf{P}),$$

and therefore (Theorem 2.2 of Chapter 6) v_2 is the corresponding eigenvector, as anticipated by the notation. Writing down explicitly

$$\mathbf{P}v_2 = -\frac{\pi(1)}{(1 - \pi(1))} v_2,$$

we obtain for $i > 1$,

$$\frac{\pi(1)}{\pi(i)} p_{1i} - \pi(1) = -\frac{\pi(1)}{(1 - \pi(1))} (-\pi(1)),$$

that is,

$$p_{1i} = \frac{\pi(i)}{(1 - \pi(1))}. \quad (7.24)$$

By reversibility, this implies

$$p_{i1} = \frac{\pi(1)}{(1 - \pi(1))}. \quad (7.25)$$

In summary, imposing the second eigenvalue to be equal to $\frac{\pi(1)}{(1 - \pi(1))}$ determines the first row and first column of \mathbf{P} .

We must now complete the construction by determining the restriction \mathbf{P}_2 of \mathbf{P} to $\{2, \dots, r\}$. In view of (7.25),

$$\mathbf{P}_2 = \left(1 - \frac{\pi(1)}{1 - \pi(1)}\right) \mathbf{P}',$$

where \mathbf{P}' is a stochastic matrix. Define

$$\pi' = \left(\frac{\pi(2)}{1 - \pi(1)}, \dots, \frac{\pi(r)}{1 - \pi(1)}\right)^T.$$

Since \mathbf{P} is in detailed balance with π , \mathbf{P}' is in detailed balance with π' , and therefore π' is a stationary distribution of \mathbf{P}' .

\mathbf{P}_2 and \mathbf{P}' have the same eigenvectors. In particular, a right eigenvector of \mathbf{P}_2 different from $v'_1 = (1, \dots, 1)^T \in \mathbb{R}^{r-1}$ must be π' -orthogonal to v'_1 .

Now, finding a right eigenvector of \mathbf{P}_2 that is π' -orthogonal to v'_1 is equivalent to finding a right eigenvector of \mathbf{P} π -orthogonal to v_1 and v_2 . Indeed, if v is a right eigenvector of \mathbf{P} π -orthogonal to v_1 and v_2 , it must be π -orthogonal to $v_2 - v_1 = \delta_1$, that is, its first coordinate is 0, i.e., $v = (0, x_2, \dots, x_r)^T$, and $v' = (x_2, \dots, x_r)^T$ is π' -orthogonal to v'_1 , and the latter is easily seen to be a right eigenvector of \mathbf{P}_2 corresponding to the same eigenvalue.

Conversely, if $v' = (x_2, \dots, x_r)^T$ is a right eigenvector of \mathbf{P}_2 that is π' -orthogonal to v'_1 , then $v = (0, x_2, \dots, x_r)^T$ is a right eigenvector of \mathbf{P} corresponding to the same eigenvalue.

This shows that the second eigenvalue of \mathbf{P}_2 is the third eigenvalue of \mathbf{P} .

We choose \mathbf{P}_2 such that its second eigenvalue is the smallest possible, and therefore the third eigenvalue of \mathbf{P} is the smallest possible, under the constraint that its second eigenvalue is the smallest possible. Since the eigenvalues of \mathbf{P} and \mathbf{P}_2 are proportional, it suffices to perform on the pair (\mathbf{P}', π') the same treatment as for (\mathbf{P}, π) . This leads to

$$\lambda'_2 = -\frac{\pi'(2)}{1 - \pi'(2)} = -\frac{\pi(2)}{\pi(3) + \dots + \pi(r)},$$

and the corresponding eigenvector is

$$v'_2 = (1 - \pi'(2), -\pi'(2), \dots, -\pi'(2))^T = \left(1 - \frac{\pi(2)}{\pi(2) + \dots + \pi(r)}, -\frac{\pi(2)}{\pi(2) + \dots + \pi(r)}, \dots, -\frac{\pi(2)}{\pi(2) + \dots + \pi(r)}\right)^T.$$

Therefore

$$\lambda_3 = -\frac{\pi(2)}{\pi(3) + \dots + \pi(r)} \left(1 - \frac{\pi(1)}{\pi(2) + \dots + \pi(r)} \right)$$

and

$$v_3 = \left(0, 1 - \frac{\pi(2)}{\pi(2) + \dots + \pi(r)}, -\frac{\pi(2)}{\pi(2) + \dots + \pi(r)}, \dots, -\frac{\pi(2)}{\pi(2) + \dots + \pi(r)} \right)^T$$

and the first column of \mathbf{P}_2 is

$$(0, -\lambda_3, \dots, -\lambda_3)^T.$$

By reversibility, the first row of \mathbf{P}_2 is also determined. At this point it therefore remains to determine \mathbf{P}_3 , the restriction of \mathbf{P} to $\{3, \dots, r\}$. The iteration is now clear and the end product is a stochastic matrix with the following properties:

(i) Its eigenvalues are $1 = \lambda_1 > 0 > \lambda_2 \geq \dots \geq \lambda_r$, such that for $i \in [1, r-1]$, λ_{i+1} is the smallest $\lambda_i(\mathbf{P})$ for all $\mathbf{P} \in M(\pi)$ that already have $1, \dots, \lambda_i$ for eigenvalues.

(ii) Its diagonal elements are all null except perhaps the last one.

(iii) It has constant entries under the diagonal for each column, namely $\lambda_2, \dots, \lambda_r$.

(iv) Its right-eigenvectors are $v_1 = \mathbf{1}$ and for $k \geq 1$

$$v_{k+1} = \left(0, \dots, 0, 1 - \frac{\pi(k)}{\pi(k) + \dots + \pi(r)}, -\frac{\pi(k)}{\pi(k) + \dots + \pi(r)}, \dots, -\frac{\pi(k)}{\pi(k) + \dots + \pi(r)} \right)^T$$

where the first $k-1$ entries are null. The eigenvalues are, for $k \geq 1$,

$$\lambda_{k+1} = -\frac{\pi(k)}{\pi(k+1) + \dots + \pi(r)} \prod_{\ell=1}^{k-1} \left(1 - \frac{\pi(\ell)}{\pi(\ell) + \dots + \pi(r)} \right).$$

Under condition (i), the above matrix is unique. □

8 Simulated Annealing

8.1 Stochastic Descent and Cooling

Let E be a finite set and $U : E \rightarrow \mathbf{R}$ a function, called the *cost function*, to be minimized. More precisely, one is looking for any element $i_0 \in E$ such that

$$U(i_0) \leq U(i) \text{ for all } i \in E.$$

Such an element is called a *global minimum* of the cost function U . The so-called descent algorithms define for each state $i \in E$ a subset $N(i)$ of E , called the *neighborhood* of i ,

and proceed iteratively as follows: Suppose that at a given stage solution i is examined. At the next stage one examines a solution $j \in N(i)$ chosen according to a rule specific to each algorithm, and compares $U(i)$ to $U(j)$. If $U(i) \leq U(j)$, the procedure stops and i is the retained solution. Otherwise a new solution $k \in N(j)$ is examined and compared to j , and so on. The algorithm eventually comes to a stop and produces a solution, since E is finite. However, this solution is usually not optimal, due to the possible existence of local minima that are not global minima, where i is called a *local minimum* if

$$U(i) \leq U(j) \text{ for all } j \in N(i).$$

In most interesting situations, local optima exist and are sometimes abundant, and the algorithms often become trapped at one of these local minima.

Example 8.1. The Traveling Salesman

A salesman must find the shortest route visiting each of K cities exactly once. Here E is the set of the $K!$ admissible routes, i is one such route, and $U(i)$ is the length of route i . One popular choice for the neighborhood $N(i)$ of route i is all the routes j obtained from i by a 2-change, as explained below.




Figure 7.8.1. A 2-change of a salesman's route

If the cities are numbered $1, 2, \dots, K$ a route i can be identified with a permutation σ of $\{1, 2, \dots, K\}$, where $\sigma(\alpha)$ is the order of the visit to city α in the route corresponding to σ . By renaming the cities, there is no loss of generality in supposing that σ is the identity. The 2-change involving cities α and β consists in cutting the segments $(\alpha, \alpha + 1)$ and $(\beta - 1, \beta)$, and in replacing them by the new segments $(\alpha, \beta - 1)$ and $(\beta - 1, \alpha)$. (See Fig. 7.8.1.) In this construction $|\beta - \alpha| \geq 3$ and $K \geq 4$, and one can count exactly $K(K - 1) + 1$ neighbors of a given route. The sizes of the neighborhoods are therefore reasonable in comparison to the size of the total search space. Note also that the computation of $U(j)$ from $U(i)$ when $j \in N(i)$ involves only four intercity distances. \diamond

Typically, and this is indeed the case in the traveling salesman problem, the neighborhood structure is communicating.

Definition 8.1. *Communicating Neighborhoods*

Let E be a finite set and let $\{N(i), i \in E\}$ be a collection of subsets of E satisfying condition

$$i \notin N(i). \quad (8.1)$$

Such a collection is called a *neighborhood structure*. If for all pairs of state $i, j \in E$ there exists a *path* from i to j , that is, a sequence of states $i_1, \dots, i_m \in E$ such that $i_1 \in N(i), i_2 \in N(i_1), \dots, j \in N(i_m)$, the neighborhood structure is called communicating.

The basic idea of stochastic combinatorial optimization is to leave a possibility to escape from a local minimum trap.

A canonical form of the stochastic descent algorithm is as follows: Let $Q = \{q_{ij}\}$ be an irreducible transition matrix on E . Also, for each parameter value T , and all states $i, j \in E$, let $\alpha_{ij}(T)$ be a probability. Calling X_n the current solution at stage n , the process $\{X_n\}_{n \geq 0}$ is a homogeneous Markov chain with state space E and transition matrix $\mathbf{P}(T)$ of general off-diagonal term

$$p_{ij}(T) = q_{ij}\alpha_{ij}(T). \quad (8.2)$$

We assume the chain irreducible. It is positive recurrent, since the state space is finite (Theorem 3.3, Chapter 3). Therefore, it has a unique stationary distribution $\pi(T)$.

One possible choice of the candidate-generating matrix Q consists in first choosing a communicating neighborhood structure such as in Definition 8.1, and taking $q_{ij} > 0$ only if $i = j$ or $j \in N(i)$. The matrix Q is then irreducible. Conversely, one can associate with an irreducible transition matrix Q a communicating neighborhood structure defined by $N(i) = \{j; j \neq i, q_{ij} > 0\}$.

Example 8.2. *Metropolis Sampler*

Suppose that the current solution at stage n is i . At stage $n + 1$, a tentative solution j is selected according to the probability q_{ij} . This solution is accepted with probability

$$\alpha_{ij}(T) = \min\left\{1, e^{\frac{U(i)-U(j)}{T}}\right\} = e^{-\frac{U(i)-U(j)^+}{T}}, \quad (8.3)$$

where T is a positive constant; otherwise, j is rejected. The meaning of (8.3) is that if $U(i) \geq U(j)$, then j is accepted, whereas if $U(i) < U(j)$, a chance is left to the solution j , although it is worse than i . In this particular algorithm as well as in others, the chance left to a candidate j that is worse than i diminishes as its deviation from i , measured by $U(j) - U(i)$, increases.

Suppose that the matrix Q is symmetric. With this special structure, the stationary distribution $\pi(T)$ does not depend on Q and is given by

$$\pi_i(T) = \frac{e^{-U(i)/T}}{\sum_{k \in E} e^{-U(k)/T}}. \quad (8.4)$$

◇

Example 8.3. *Barker's Sampler*

This simulation algorithm differs from the Metropolis sampler by the choice of the acceptance probability, here

$$a_{ij}(T) = \frac{1}{1 + e^{-(U(i)-U(j))/T}}. \quad (8.5)$$

The usual reversibility argument shows that the target distribution of Barker's algorithm is the same as that of the Metropolis algorithm. \diamond

For both the Metropolis and Barker samplers, if Q is irreducible and U is not a constant, then $\mathbf{P}(T)$ is irreducible and aperiodic for all $T > 0$ (Problem 7.8.1).

Example 8.4. *Gibbs Sampler as a Stochastic Descent Algorithm*

First, we shall see how the Gibbs sampler connects to the general descent algorithm corresponding to the transition probabilities (8.2). This is only a matter of notation. For simplicity, we assume that the phase space is $\Lambda = \{0, 1\}$.

In this context, a "solution" i is a value $x \in \Lambda^S$ of the random field. The neighborhood $N(x)$ of the solution $x \in \Lambda^S$ (not to be confused with a neighborhood \mathcal{N}_s of some site $s \in S$) is the set of configurations $y \in \Lambda^S$ such that y and x agree on all sites except one:

$$y = (\bar{x}(s), x(S \setminus s)), \quad (8.6)$$

where $\bar{x}(s) = 1 - x(s)$. The configuration y in (8.6) will be denoted by \bar{x}^s . Therefore,

$$N(x) = \{\bar{x}^s; s \in S\}. \quad (8.7)$$

The Gibbs sampler does this: After examining the image $x \in \Lambda^S$ at time n , it selects $\bar{x}^s \in N(x)$ with probability q_s (with the notations of (8.2), $i = x$, $j = \bar{x}^s$, and $q_{ij} = q_s$). The image \bar{x}^s is accepted with probability $\pi(\bar{x}(s) | x(\mathcal{N}_s))$ (with the notations of (8.2), $a_{ij}(T) = \pi(\bar{x}(s) | x(\mathcal{N}_s))$). \diamond

Example 8.5. *Barker's Algorithm for Neural Networks*

The model and the notation are those of Example 1.2:

$$\pi_T(\bar{x}(s) | x(\mathcal{N}_s)) = \frac{e^{-\frac{1}{T}(\sum_{t \in \mathcal{N}_s} (w_{ts} + w_{st})x(t) - h_s)(1-x(s))}}{1 + e^{-\frac{1}{T}(\sum_{t \in \mathcal{N}_s} (w_{ts} + w_{st})x(t) - h_s)}}. \quad (8.8)$$

In the notations of (8.2), with $i = x$, $j = \bar{x}^s$, the right hand side of (8.13) is $a_{ij}(T) = a_{x, \bar{x}^s}(T)$. In this case, the simulation algorithm is a special case of Barker's algorithm. Indeed,

$$\mathcal{E}(x) - \mathcal{E}(\bar{x}^s) = \left\{ \sum_{t \in \mathcal{N}_s} (w_{ts} + w_{st})x(t) - h_s \right\} (x(s) - \bar{x}(s)).$$

Therefore, if $x(s) = 1$, then $\mathcal{E}(x) - \mathcal{E}(\bar{x}^s) = \sum_{t \in \mathcal{N}_s} ((w_{ts} + w_{st})x(t) - h_s)$, and if $x(s) = 0$, then $\mathcal{E}(x) - \mathcal{E}(\bar{x}^s) = -\sum_{t \in \mathcal{N}_s} ((w_{ts} + w_{st})x(t) - h_s)$. Therefore, the quantity on the left side of (8.8) is equal to

$$\frac{1}{1 + e^{-\frac{1}{T}(\mathcal{E}(x) - \mathcal{E}(\bar{x}^s))}}$$

if $x(s) = 1$, and to

$$\frac{e^{\frac{1}{T}(\mathcal{E}(x) - \mathcal{E}(\bar{x}'))}}{1 + e^{\frac{1}{T}(\mathcal{E}(x) - \mathcal{E}(\bar{x}'))}} = \frac{1}{1 + e^{-\frac{1}{T}(\mathcal{E}(x) - \mathcal{E}(\bar{x}'))}}$$

if $x(s) = 0$. In either case, the probability of acceptance of \bar{x}^s is equal to

$$\frac{1}{1 + e^{-\frac{1}{T}(\mathcal{E}(x) - \mathcal{E}(\bar{x}'))}}. \quad \diamond$$

Consider the chain with transition probabilities (8.2), and suppose that the corresponding HMC is irreducible, so that the stationary distribution $\pi(T)$ exists and is unique. Suppose, moreover, that the chain is aperiodic. By Theorem 2.1 of Chapter 4, it converges to steady state, in that for all $i \in E$ and all initial distributions μ ,

$$\lim_{n \rightarrow \infty} P_\mu(X_n = i) = \pi_i(T). \quad (8.9)$$

We see, in particular, that if the stationary distribution $\pi(T)$ puts most of its mass on the states minimizing the cost function $U(i)$, then if one stops the algorithm at a stage n sufficiently large, the current solution X_n will be a global minimum with high probability.

Example 8.6. Metropolis Sampler (8.2 continued)

The stationary probability is given by (8.4). Define the set of *global minima*

$$H = \{i \in E; U(i) \leq U(j) \text{ for all } j \in E\}. \quad (8.10)$$

Then clearly, $\pi_i(T)$ is maximal on $i \in H$. But there is more:

$$\lim_{T \downarrow 0} \pi_i(T) = \begin{cases} \frac{1}{|H|} & \text{if } i \in H, \\ 0 & \text{if } i \notin H. \end{cases} \quad (8.11)$$

To see this, call $m = \min_{i \in E} U(i)$, and write the right-hand side of (8.11), after division of its numerator and denominator by $e^{-\frac{m}{T}}$, as

$$\frac{e^{-\frac{(U(i)-m)}{T}}}{|H| + \sum_{k \notin H} e^{-\frac{(U(k)-m)}{T}}}$$

The result follows, since as $T \downarrow 0$, $e^{-\frac{U(k)-m}{T}}$ tends to 0 if $U(k) > m$, and to 1 if $U(k) = m$. \(\diamond\)

This observation suggests the following heuristic procedure (Kirkpatrick, Gelatt, and Vecchi (1982); see (Aarts and Cors, 1989) for additional details) *indexVecchi*. Start the algorithm with the value $T = a_0$ of the parameter, and wait a sufficiently long time for the chain to get close to its stationary regime. Then set $T = a_1 < a_0$ and again wait for the steady state. Then set $T = a_2 < a_1$, etc. At the k th change of the parameter T , the chain

will be close to the stationary regime $\pi(a_k)$, and therefore if $\lim_{k \uparrow \infty} a_k = 0$, one expects that for large n , X_n will be with very high probability in H , the set of global minima.

However, for this to happen, the times in between the parameter changes must be sufficiently long for the chain to come close to the stationary distribution corresponding to the current value of the parameter. What is “sufficiently long”?

Simulated annealing algorithms all have a *cooling schedule*, that is, a sequence $\{T_n\}_{n \geq 0}$ of positive numbers decreasing to 0 controlling the transition rates of $\{X_n\}_{n \geq 0}$: At time n , $P(X_{n+1} = j \mid X_n = i) = p_{ij}(T_n)$. The question becomes, *How slowly* must $\{T_n\}_{n \geq 0}$ converge to zero so that

$$\lim_{n \uparrow \infty} P(X_n = i) = \begin{cases} 0 & \text{if } i \notin H, \\ \frac{1}{|H|} & \text{if } i \in H. \end{cases} \quad (8.12)$$

Since $\{X_n\}_{n \geq 0}$ is now a nonhomogeneous Markov chain, the results concerning convergence in distribution of such chains of Chapter 6 will be especially useful. They will be applied in Section 6 to obtain cooling schedules with the desired convergence property in various situations.

The next example features an avatar of the simulated annealing algorithm involving on-line estimation of the cost function.

Example 8.7. Stochastic Ruler

In a number of situations the cost function U is not directly computable, but of the form

$$U(i) = E[H(i)], \quad (8.13)$$

where $H(i)$ is a nonnegative random variable. Moreover, it is assumed that one can generate at will, for any $i \in E$, a random variable $H(i)$, but that one is not able to compute its expectation. Therefore, the simulated annealing procedure as described before cannot be applied as it is. In one way or another, the expectation will have to be estimated from independent samples. The stochastic ruler algorithm of Yan and Mukai (1992) does this in an indirect manner.

It works under the assumption that the domains of the random variables $H(i)$ are all contained in a fixed interval, i.e., for all $i \in E$,

$$-\infty < a \leq H(i) \leq b < \infty,$$

and it uses at each step a random variable U , called the *ruler*, which is uniform on $[a, b]$:

$$P(U \leq x) = \frac{x - a}{b - a},$$

for all $x \in [a, b]$. The stochastic ruler algorithm is based on the observation that the simulated annealing algorithm (8.2) with the acceptance probability

$$\alpha_{ij}(T) = p(j)^{1/T} = e^{-\frac{1}{T} \{-\log p(j)\}}, \quad (8.14)$$

where

$$p(j) = P(H(j) \leq U),$$

leads, with a good cooling schedule, to the minimum of $g(i) = -\log p(i)$. But if $H(j)$ and U are independent, then

$$p(j) = E \left[\frac{1}{b-a} \int_a^b \mathbf{1}_{\{u \geq H(i)\}} du \right] = \frac{1}{b-a} E[b - H(i)] = \frac{b - U(i)}{b-a}.$$

Therefore, the simulated annealing with acceptance probability (8.14) will provide, for a good cooling schedule, a minimum of the original cost function U .

The problem now is: how to implement the acceptance rule based on (8.14), since $p(j)$ is not computable? Of course, one also has to give a cooling schedule that makes simulated annealing work. For the time being, suppose that we have a cooling schedule $\{T_k\}$ of the form $T_k = \frac{1}{M_k}$ for some integer M_k . The problem is therefore to realize for any integer M the acceptance probability $p(j)^M$ without computing $p(j)$. A solution is as follows. Generate $2M$ independent random variables $V_1, \dots, V_M, H_1(j), \dots, H_M(j)$, where the V_k are uniform on $[a, b]$ and the $H_k(j)$ have the same distribution as $H(j)$. Then

$$P(H_1(j) \leq V_1, \dots, H_M(j) \leq V_M) = P(H(j) \leq V)^M = p(j)^M.$$

The acceptance rule for j is therefore to accept if $H_k(j) \leq V_k$ for $k = 1, \dots, M$.

This procedure works, theoretically, but it requires more and more time as M increases to ∞ . As the algorithm makes progress, more and more care must be taken for the estimation of $U(i)$. \diamond

8.2 Convergence of Simulated Annealing

We begin with an example, and then proceed to the general theory.

Example 8.8. *Annealed Gibbs Sampler.*

(D. and S. Geman (1984)) We use periodic scanning as in Example 6.4, only at the n th sweep, we introduce a temperature T_n . Thus $\{Z_n\}_{n \geq 0}$ is a nonhomogeneous Markov chain, the transition matrix at time n being

$$\mathbf{P}(n) = \prod_{k=1}^N \mathbf{P}_{s(k)}^{T_n},$$

where the (x, y) -entry of \mathbf{P}_s^T is

$$\frac{\exp \left\{ -\frac{1}{T} \mathcal{E}(y(s), x(S \setminus s)) \right\}}{\sum_{\lambda \in \Lambda} \exp \left\{ -\frac{1}{T} \mathcal{E}(\lambda, x(S \setminus s)) \right\}}$$

if $y = (y(s), x(S \setminus s))$. The bound of Example 6.5 gives

$$\delta(\mathbf{P}(n)) \leq 1 - e^{-\frac{\Delta}{T_n}}.$$

In particular, by the block criterion of weak ergodicity, Theorem 8.2 of Chapter 6,

$$\sum_{n=1}^{\infty} e^{-\frac{N\Delta}{T_n}} = \infty \quad (8.15)$$

is a sufficient condition of weak ergodicity.

Now, $\mathbf{P}(n)$ admits the stationary distribution

$$\pi_{T_n}(x) = \frac{e^{-\frac{1}{T_n}\mathcal{E}(x)}}{Z_{T_n}}.$$

Also, for all $x \in \Lambda^S$, $\lim_{T \downarrow 0} \pi_T(x) = \frac{1}{|H|}$ if $x \in H$ and is 0 otherwise, where $H = \{x \in \Lambda^S; \mathcal{E}(x) = \min\}$. Moreover, it can be shown that for $x \in H$, the quantity $\pi_T(x)$ it increases as $T \downarrow 0$, whereas for $x \notin H$, it eventually decreases, and this guarantees that

$$\sum_{n=1}^{\infty} |\pi_{T(n+1)} - \pi_{T_n}| < \infty.$$

Therefore, by Theorem 8.3 of Chapter 6, if $T_n \downarrow 0$ in such a way that (8.15) is respected, then the nonhomogeneous MC $\{Z_n\}_{n \geq 0}$ is strongly ergodic, with the limit distribution uniform on H . \diamond

The general results of Chapter 6 will be applied to the simulated annealing algorithm corresponding to the transition matrix $\mathbf{P}(T)$ given by (8.2). We follow the analysis of Anily and Federgruen (1987).

The transition matrix $\mathbf{P}(T)$ is assumed *uniformly irreducible* for sufficiently small $T \in (0, 1]$. This means that for all ordered pair of states (i, j) , there is a $\mathbf{P}(T)$ -path from i to j which is independent of $T \in (0, c]$ for some $c > 0$. This is always satisfied in practice. For instance, for the Metropolis or Barker samplers, it suffices that Q is irreducible and that U is not a constant (see Problem 7.8.1).

Define

$$d = \inf \{q_{ij}; \quad j \neq i, p_{ij}(T) > 0\}, \quad (8.16)$$

a positive quantity, since the state space is finite.

The crucial assumption is the following: There exists $T^* \in (0, 1]$ such that on $(0, T^*]$,

$$\alpha_{ij}(T) \downarrow 0 \text{ as } T \downarrow 0 \text{ if } U(j) > U(i), \quad (8.17)$$

$$\alpha_{ij}(T) \uparrow 1 \text{ as } T \downarrow 0 \text{ if } U(j) < U(i), \quad (8.18)$$

and

$$\lim_{T \downarrow 0} \alpha_{ij}(T) > 0 \text{ exists if } U(i) = U(j). \quad (8.19)$$

Assumptions (8.17) and (8.18) imply, in particular, that in the vicinity of 0, the functions $\alpha_{ij}(T)$ are monotonic if $U(i) \neq U(j)$. They force the algorithm less permissive as T approaches 0, rejecting more often the nonlocally optimal solutions. Define for each $T \in (0, 1]$

$$\underline{\alpha}(T) = \inf_{i \in E, j \in N(i)} \alpha_{ij}(T). \quad (8.20)$$

Assumptions (8.17)–(8.19) imply that in the vicinity of 0,

$$\inf_{i \in E, j \neq i} \alpha_{ij}(T) = \inf_{\substack{i \in E, j \in N(i) \\ U(j) > U(i)}} \alpha_{ij}(T),$$

and therefore, in the vicinity of 0, $\underline{\alpha}(T)$ is decreasing to zero.

We are now ready for the main result.

Theorem 8.1. *Weak Ergodicity of Simulated Annealing*

Let $\{\mathbf{P}(T)\}_{T \in (0,1]}$ satisfy the above assumptions. Let $\{T_n\}_{n \geq 0}$ be a sequence of numbers in $(0, 1]$ decreasing to zero as $n \rightarrow \infty$. Then if

$$\sum_{k=0}^{\infty} (\underline{\alpha}(T_{kN}))^N = \infty, \quad (8.21)$$

$\{\mathbf{P}(T_n)\}_{n \geq 0}$ is weakly ergodic.

Proof. Define, with a slight and innocuous ambiguity, $\mathbf{P}(n) = \mathbf{P}(T_n)$. The uniform irreducibility assumption guarantees the existence, for all ordered pair of states i, j , of a path $i_0 = i, i_1, \dots, i_N = j$ such that

$$p_{i_j, i_{j+1}}(kN + j, kN + j + 1) = p_{i_j, i_{j+1}}(T_{kN+j}) > 0.$$

But $p_{kl}(T) > 0$ implies $p_{kl}(T) \geq d\underline{\alpha}(T)$, and therefore

$$p_{i_j, i_{j+1}}(kN + j, kN + j + 1) \geq d\underline{\alpha}(T_{kN+j}).$$

Since $\underline{\alpha}(T)$ is, in the vicinity of 0, monotone decreasing, then for sufficiently large k

$$p_{i_j, i_{j+1}}(kN + j, kN + j + 1) \geq d\underline{\alpha}(T_{(k+1)N}),$$

and therefore

$$p_{ij}(kN, (k+1)N) \geq d^N (\underline{\alpha}(T_{(k+1)N}))^N.$$

Therefore, in view of (7.3) of Chapter 6,

$$1 - \delta(\mathbf{P}(kN, (k+1)N)) \geq d^N (\underline{\alpha}(T_{(k+1)N}))^N.$$

Therefore, (8.21) implies

$$\sum_{k=1}^{\infty} (1 - \delta(\mathbf{P}(kN, (k+1)N))) = \infty,$$

and the conclusion follows from the block criterion, Theorem 8.2, Chapter 6. \square

Example 8.9. *Simulated Annealing with the Metropolis Sampler*

The acceptance probabilities of the Metropolis sampler are

$$\alpha_{ij}(T) = \begin{cases} e^{(U(j)-U(i))/T} & \text{if } U(j) > U(i), \\ 1 & \text{if } U(j) \leq U(i). \end{cases}$$

We see that conditions (8.17), (8.18), and (8.19) are satisfied. We have

$$\underline{\alpha}(T) = \inf_{\substack{j \in N(i) \\ U(i) < U(j)}} e^{-(U(j)-U(i))/T},$$

and therefore

$$\underline{\alpha}(T) \geq e^{-\frac{\Delta}{T}}, \tag{8.22}$$

where

$$\Delta = \sup\{U(j) - U(i); j \in N(i)\}. \tag{8.23}$$

It follows that

$$\sum_{k=0}^{\infty} \{\underline{\alpha}(T_{kN})\}^N \geq \sum_{k=0}^{\infty} e^{-\frac{N\Delta}{T_{kN}}}.$$

For a cooling schedule $\{T_k\}_{k \geq 0}$ satisfying

$$T_k \geq \frac{N\Delta}{\log(k)}, \tag{8.24}$$

we see that

$$\sum_{k=1}^{\infty} \{\underline{\alpha}(T_{kN})\}^N \geq \sum_{k=1}^{\infty} \frac{1}{kN} = \infty,$$

and therefore, $\{\mathbf{P}(T_n)\}_{n \geq 1}$ is weakly ergodic.

Therefore, in view of Theorem 8.2 and Example 8.2 of Chapter 6, $\{\mathbf{P}(T_n)\}$ is strongly ergodic. As shown in Example 8.6, the limiting probability vector puts all its mass uniformly on the set H of global minima. Therefore a cooling schedule verifying (8.24) guarantees convergence in distribution to the set of global minima. \diamond

We shall now see the effects of fast cooling. Denote by $\mathbf{P}(\text{lim})$ the transition matrix corresponding to the limit case $T \downarrow 0$. In particular, $p_{ij}(\text{lim}) = 0$ if $U(i) < U(j)$. Call R_1 the recurrent communication class of some global minimum, and R_2 the recurrent communication class of some strictly local minimum. Note that R_1 only contains global minima, and in particular, R_1 and R_2 are disjoint. Define

$$\bar{\alpha}(2, T) = \sup_{i \in R_2, j \in N(i)} \alpha_{ij}(T). \tag{8.25}$$

Since for $j \in R_2$,

$$\sum_{\ell \in R_2} p_{j,\ell}(T_k) = 1 - \sum_{\substack{\ell \notin R_2 \\ j \in N(\ell)}} q_{j\ell} \alpha_{j\ell}(T_k) \geq 1 - \bar{\alpha}(2, T_k),$$

the probability of staying in R_2 forever is bounded from below by $\prod_{k=1}^{\infty} (1 - \bar{\alpha}(2, T_k))$. This infinite product is strictly positive if $\sum_{k=1}^{\infty} \bar{\alpha}(2, T_k) < \infty$. Therefore, if the chain has at least one strictly local minimum, then under the condition

$$\sum_{k=1}^{\infty} \bar{\alpha}(2, T_k) < \infty, \quad (8.26)$$

the probability that it stays eternally in R_2 is strictly positive. In particular, since no globally optimal solution is in R_2 , with positive probability the algorithm will never visit a globally optimal state.

Example 8.10. (8.2 continued)

Define

$$\delta_2 = \inf \{U(j) - U(i); i \in R_2, j \notin R_2, j \in N(i)\} \quad (8.27)$$

and suppose that

$$\delta_2 > 0. \quad (8.28)$$

Since $\bar{\alpha}(2, T) \leq e^{-\frac{\delta_2}{T}}$, we have

$$\sum_{k=0}^{\infty} \bar{\alpha}(2, T_k) \leq \sum_{k=1}^{\infty} e^{-\delta_2/T_k}.$$

Therefore, if the cooling schedule satisfies

$$T_k \leq \frac{\delta_2 - \alpha}{\log k}$$

for some $\alpha > 0$ such that $\delta_2 - \alpha > 0$, we have

$$\sum_{k=1}^{\infty} \bar{\alpha}(2, T_k) \leq \sum_{k=1}^{\infty} e^{-(\log k)(1+\epsilon)},$$

where $1 + \epsilon = \frac{1}{1 - \frac{\alpha}{\delta_2}}$ (and therefore $\epsilon > 0$). Thus

$$\sum_{k=1}^{\infty} \bar{\alpha}(2, T_k) \leq \sum_{k=1}^{\infty} \frac{1}{k^{1+\epsilon}} < \infty,$$

which implies that the cooling schedule does not yield convergence in distribution toward the set of global minima. \diamond

For the simulated annealing algorithm based on the Metropolis sampler, a necessary and sufficient condition of convergence was found by Hajek (1988). It says that there exists a constant γ , such that a necessary and sufficient of convergence of the Metropolis simulated annealing algorithm, whatever the initial state, is

$$\sum_{k=1}^{\infty} e^{-\frac{\gamma}{T_k}} = \infty.$$

In particular, a logarithmic cooling schedule

$$T_k = \frac{a}{\ln(k + 1)}$$

yields convergence if and only if

$$a \geq \gamma.$$

The results of convergence given in the present section are of theoretical and qualitative interest only. Practical algorithms use faster than logarithmic schedules on a finite horizon. The theory and the performance evaluation of these algorithms is more difficult and outside the scope of this book.

Problems

7.1.1 Define on $S = \mathbb{Z}^2$ the two neighborhood systems of Fig. P7.1.1. Describe the corresponding cliques and give the boundary of a 3×3 square for each case.



Figure P7.1.1

7.1.2 Consider the nonoriented graph on $S = \{1, 2, 3, 4, 5, 6, 7\}$ in Fig. P7.1.2. Let the phase space be $\Lambda = \{-1, +1\}$. For a configuration $x \in \Lambda^S$, denote by $n(x)$ the number of positive *bonds*, i.e., the number of edges of the graph for which the phases of the adjacent sites coincide. Define a probability distribution π on Λ^S by $\pi(x) = \frac{e^{-\beta n(x)}}{Z}$. Give the value of Z and the local characteristics of this random field.

7.1.3 Inspired by the cellular automaton of Example 3.4, Chapter 5, give a counterexample for Theorem 1.1 if the positivity hypothesis is abandoned.

7.2.1 Let S be a finite set of sites and Λ a finite set of phases. Let $\{X(s)\}_{s \in S}$ be an MRF with values in Λ^S and admitting a Gibbsian description in terms of the neighborhood structure N on S and the Gibbs potential $\{V_C\}_{C \subset S}$. Prove that for all subsets A, B , of S such that

$$A \cap B = \emptyset$$

it holds that for all $x \in \Lambda^S$,

$$P(X(A) = x(A) \mid X(B) = x(B)) = P(X(A) = x(A) \mid X(\partial \bar{B}) = x(\partial \bar{B})).$$




Figure P7.1.2

7.2.2 Let $X = \{X(s)\}_{s \in S}$ be an MRF defined on S , with the neighborhood system (α) of Example 1.2. Define for $i \in [1, m]$, the p -dimensional vector Y_i by

$$Y_i = \begin{pmatrix} X_{(i,1)} \\ \vdots \\ X_{(i,p)} \end{pmatrix}$$

(thus Y_i is the “ i th column” of the random field X). Using the result of Problem 7.2.1, show that $\{Y_i\}_{1 \leq i \leq p}$ is a Markov chain with values in $E = \Lambda^p$, where Λ is the phase space (assumed countable).

7.2.3 In the Ising model, explain why chaotic boundaries between patches are less likely than smooth boundaries (part of the exercise is to interpret this statement).

7.2.4 Let S be a finite set of sites and Λ a finite set of phases. Let $\{X(s)\}_{s \in S}$ be an MRF with values in Λ^S and admitting a Gibbsian description in terms of the neighborhood structure N on S and the potential $\{V_C\}_{C \subset S}$. Let $A + B = S$ be a partition of the site. Fix $x(A) = \underline{x}(A)$ and define the distribution π_A on Λ^B by

$$\pi_A(x(B)) = \frac{e^{-\mathcal{E}(\underline{x}(A), x(B))}}{\sum_{y(B) \in \Lambda^B} e^{-\mathcal{E}(\underline{x}(A), y(B))}},$$

where \mathcal{E} is the energy function associated with the potential $\{V_C\}_{C \subset S}$. Show that

$$\pi_A(x(B)) = P(X(B) = x(B) \mid X(A) = \underline{x}(A))$$

and that $\pi_A(x(B))$ is a Gibbs distribution for which you will give the neighborhood system and the corresponding cliques, as well as the local characteristics. (An MRF with values in Λ^B and with the distribution π_A is called a version of $\{X_s\}_{s \in S}$, *frozen on A at value $\underline{x}(A)$, or clamped at $\underline{x}(A)$).*

7.2.5 Let $(Z(s), s \in S)$ be a family of i.i.d random variables indexed by a finite set S , with $P(Z(s) = -1) = p, P(Z(s) = +1) = q = 1 - p$. Show that

$$P(Z = z) = \prod_{s \in S} P(Z(s) = z(s)) = K e^{\gamma \sum_{s \in S} z(s)},$$

where $\gamma = \frac{1}{2} \ln \left(\frac{p}{q} \right)$ and $K = (pq)^{\frac{|S|}{2}}$.

7.2.6 In this problem, we consider a random field with a unilateral Markov description. Let $S = \mathbb{Z}_m^2$ and suppose that S is scanned as follows:

$$s_1 = (1, 1), s_2 = (1, 2), \dots, s_m = (1, m),$$

$$s_{m+1} = (2, 1), s_{m+2} = (2, 2), \dots, s_{2m} = (2, m), \dots$$

Define the past relevance structure $\{\mathcal{P}(s)\}_{s \in S}$ by

$$\mathcal{P}((i, j)) = \{(i - 1, j - 1), (i - 1, j), (i, j - 1)\} \cap S.$$

The scanning, the past relevance structure, and the neighborhood system are described in Fig. P7.2.6.




Figure P7.2.6

For simplification in the notation, call the phase at a given site $s, x(s) = x_0(s)$ or x_0 and at the adjacent sites in $\mathcal{P}(s), x_1(s)$ (or x_1), $x_2(s)$ (or x_2), $x_3(s)$ (or x_3) as indicated in Fig. P7.2.6. If site s is at the boundary of S , some of the sites in Fig. P7.2.6 may fall outside S , in which case they must be ignored in the computations below. Here $\Lambda = \{0, 1\}$, and the transitions are described as follows:

- (i) If there is a majority of 1's in $\mathcal{P}(s), x(s) = 1$ with probability α .
- (ii) If there is a majority of 0's in $\mathcal{P}(s), x(s) = 0$ with probability β .

This interpretation must, of course, be slightly modified if s is on the boundary of S .

Give the unilateral transition function (see Example 2.5). Give the Gibbsian description of this MRF, that is, its neighborhood system and potential functions.

7.2.7 Modify Problem 7.2.6 as follows: Use the same past relevance structure \mathcal{P} but replace the television scanning order by the antidiagonal scanning order

$$s_1 = (1, 1), s_2 = (1, 2), s_3 = (2, 1), s_4 = (3, 1), s_5 = (2, 2), s_6 = (1, 3), \text{ etc.}$$

Show that if we use the same transition functions $p(x(s) | x(\mathcal{P}(s)))$, we obtain the same random field. Generalize this result.

7.2.8 In Problem 7.2.6, replace the past relevance structure \mathcal{P} thereof by that shown in Fig. P7.2.8. What are the neighborhoods of the corresponding MRFs?



Figure P7.2.8

7.2.9 Let $\{X_n\}_{0 \leq n \leq N}$ be an AR(1) time series; more precisely,

$$X_{n+1} = aX_n + \epsilon_n, \quad n \in [0, N],$$

where $X_0, \epsilon_1, \dots, \epsilon_N$ are independent centered Gaussian random variables with respective variances $\alpha^2, \sigma^2, \dots, \sigma^2$. Show that the probability density π of $X = (X_0, \dots, X_N)$ is Gibbsian with respect to a neighboring system to be identified, and give the corresponding local characteristics.

7.3.1 Give the maximal clique corresponding to the neighborhood system of Fig. 7.3.2 (pixel + edge model of S. and D. Geman).

7.3.2 (i) Let R_1, R_2 , and T_{12} be random variables with values in $\{0, 1\}$ and such that

$$\begin{aligned} P(T_{12} = 1 \mid R_1 = 1, R_2 = 1) &= A, \quad 0 < A < 1, \\ P(T_{12} = 1 \mid R_1 = 1, R_2 = 0) &= B, \quad 0 < B < 1, \\ P(T_{12} = 1 \mid R_1 = 0, R_2 = 1) &= C, \quad 0 < C < 1, \\ P(T_{12} = 1 \mid R_1 = 0, R_2 = 0) &= D, \quad 0 < D < 1. \end{aligned}$$

Show that

$$P(T_{12} = t \mid R_1 = r_1, R_2 = r_2) = \exp\{Q(t, r_1, r_2)\},$$

where Q is a polynomial of degree 2 in r_1, r_2 with coefficients that are first-degree polynomials in t .

Let R be a random variable with values in $\{0, 1\}$, with $P(R) = p$ ($0 < p < 1$). Show that for all $r \in \{0, 1\}$, $P(R = r) = \exp(ar + b)$ for some coefficients a and b to be made explicit.

(ii) Let $P = \{(i, j) \in \mathbb{Z}^2; i, j \in [1, M]\}$ be a set of pixel sites. The pixel sites are denoted α, β, \dots . Let L be the collection of interpixel sites $\{(\alpha, \beta)\}$, where (α, β) is the midpoint of segment $\alpha\beta$, where α and β are neighbors (at distance 1).

Denote by $S = P \cup L$ the collection of sites (pixels and interpixels), and let $\{0, 1\}$ be the phase space for each site. A random field on S is therefore a collection of random variables $(R_\alpha; \alpha \in P)$ and $\{T_{(\alpha, \beta)}; (\alpha, \beta) \in L\}$.

Suppose that, conditionally on $(R_\alpha, \alpha \in P)$, the random variables $(T_{(\alpha, \beta)}; (\alpha, \beta) \in L)$ are independent, and that the conditional distribution of $T_{(\alpha, \beta)}$ given R_α and R_β is the same as

the conditional distribution of T_{12} given R_1 and R_2 described in the beginning of the problem, and that $T_{(\alpha, \beta)}$ is independent of $(R_\gamma; \gamma \in P, \gamma \neq \alpha, \beta)$ given R_α and R_β .

Suppose, finally, that the collection $(R_\alpha, \alpha \in P)$ is independent and that for all $\alpha \in P$, R_α has the same distribution as R above.

Compute

$$P(R_\alpha = r_\alpha; \alpha \in P \mid T_{(\gamma, \delta)} = t_{(\gamma, \delta)}; \langle \gamma, \delta \rangle \in L)$$

in Gibbs representation.

7.3.3 Consider the model of Example 3.1 with the following modifications. Firstly, the phase space is $\Lambda = \mathbb{N}$, and secondly, the potential is now

$$V_C(x) = \begin{cases} -\ln(g(x(s)) + \alpha_1 x(s)) & \text{if } C = \{s\} \in \mathcal{C}_1, \\ \alpha_j x(s)x(t) & \text{if } C = \{s, t\} \in \mathcal{C}_2, \end{cases}$$

where $\alpha_j \in \mathbb{R}$ and $g : \mathbb{N} \rightarrow \mathbb{R}$ is strictly positive. As in the autobinomial model, for any cluster C not of the type \mathcal{C}_j , $V_C \equiv 0$.

Choose the function g in such a way that

$$\pi^s(x) = e^{-\rho} \frac{\rho^{x(s)}}{x(s)!},$$

where $\rho = e^{-(\alpha, b)}$, and where (α, b) is as in Example 3.1. (This is the *auto-Poisson model*.)

7.5.1 Let

$$\pi_T(x) = \frac{1}{Z} e^{-\frac{E(x)}{kT}}$$

be a Gibbs distribution on the finite space $E = \Lambda^S$. Here Z is short for Z_T , and $E(x)$ is the energy of physics, differing from $\mathcal{E}(x)$ by the Boltzmann constant k .

For any function $f : E \rightarrow \mathbb{R}$, define

$$\langle f \rangle = \sum_{x \in E} \pi(x) f(x).$$

In particular, the *internal energy* is

$$U = \langle E \rangle = \sum_{x \in E} \pi(x) E(x).$$

The *free energy* F is defined by

$$F = -kT \ln(Z).$$

Show that

$$U = -T^2 \frac{\partial}{\partial T} \left(\frac{F}{T} \right).$$

(This is in agreement with standard thermodynamics.)

7.5.2 (continuation of Problem 7.5.1) For the Ising model, take

$$E(x) = E_0(x) + E_1(x),$$

where $E_0(x)$ is the interaction energy, assumed symmetric, i.e., $E_0(-x) = E_0(x)$, and

$$E_1(x) = -Hm(x),$$

where

$$m(x) = \sum_{i=1}^N x(i)$$

is the magnetic moment of the configuration $x = (x(1), \dots, x(N))$ (recall that $S = \{1, 2, \dots, N\}$), and H is the external magnetic field. The partition function, still denoted by Z , is now a function of T and H . The free energy *per site* is

$$f(H, T) = -kT \frac{1}{N} \ln(Z),$$

whereas the *magnetization*

$$M(H, T) = \frac{1}{N} \langle m \rangle$$

is the average magnetic moment per site.

Show that

$$M(H, T) = -\frac{\partial}{\partial H} f(H, T)$$

and

$$\frac{\partial M}{\partial H} = \frac{1}{NkT} (\langle m^2 \rangle - \langle m \rangle^2).$$

In particular,

$$\frac{\partial M}{\partial H} \geq 0.$$

7.5.3 (continuation of Problem 7.5.2) Compute $\lim_{N \uparrow \infty} M(H, T)$ for the Ising model on the torus (Example 1.4). Observe that this limit, as a function of H , is analytic, and null at $H = 0$. In particular, in this model, there is no phase transition.

7.6.1 Describe Gibbs sampler for the 2-state HMC of Example 2.4.

7.6.2 The synchronous version of the Gibbs sampler consists in changing the phases at all sites simultaneously rather than sequentially. The corresponding Markov chain therefore has the following transition matrix:

$$P(X_{n+1} = y \mid X_n = x) = \prod_{s \in S} \pi(y(s) \mid x(\mathcal{N}_s)).$$

Show that this quantity, denoted by p_{xy} , can be put in the form

$$p_{xy} = Z_p(x)^{-1} \exp\{-U(x, y)\},$$

where $U(x, y) = \sum_{s \in S} \sum_{c \in C} V_c(y(s), x(S \setminus s))$.

Show that if U is symmetric ($U(x, y) = U(y, x), \forall x, y \in \Lambda^S$), then $\{X_n\}_{n \geq 0}$ converges in distribution to μ , where

$$\mu(x) = \frac{\sum_{z \in \Lambda^S} \exp\{-U(x, z)\}}{\sum_{y \in \Lambda^S} \sum_{z \in \Lambda^S} \exp\{-U(y, z)\}}.$$

7.7.1 Prove that the A–R algorithm of Section 7.1 works, i.e., that $Z = Y_\tau$ admits the probability density $f(x)$. Show that $E[\tau] = c$.

7.7.2 Prove Theorem 7.1.

7.7.3 Show that in Theorem 7.1,

$$p_{22}(n) - \pi(2) = (1 - \pi(2))\lambda_2^n$$

and for $j \geq 3$,

$$p_{2j}(n) - \pi(j) = -\pi(j)\lambda_2^n.$$

Also,

$$d_V(\delta_2 \mathbf{P}^n, \pi) = (1 - \pi(2))\lambda_2^n.$$

7.8.1 Show that for both the Metropolis and Barker samplers, if Q is irreducible and U is not a constant, then $\mathbf{P}(T)$ is irreducible and aperiodic for all $T > 0$.

7.8.2 Discuss the convergence of the stochastic ruler algorithm of Example 8.7.

Continuous-Time Markov Models

1 Poisson Processes

1.1 Point Processes

This section introduces *random point processes* of which the simplest example is the *homogeneous Poisson process*. A random point process is, roughly speaking, a countable random set of points of the real line. In most applications to engineering and operations research, a *point* of a point process is the time of occurrence of some event, and this is why points are also called *events*. For instance, the arrival times of customers at the desk of a post office or jobs at the central processing unit of a computer are point-process events. In biology, an event can be the time of birth of an organism. In physiology, the firing time of a neuron is also an event. In general, point processes on the line appear in stochastic models where the state of a system is changed by the occurrence of some event. In this case one can use the phrase *stochastic systems driven by point processes*, and if the state of the system is discrete, one sometimes prefers to talk about *stochastic discrete event systems*. The basic examples are the Poisson process and the continuous-time Markov chain.

Definition 1.1. *Random Point Process*

A *random point process* on the positive half-line is a sequence $\{T_n\}_{n \geq 0}$ of nonnegative random variables such that, almost surely,

- (i) $T_0 \equiv 0$,
- (ii) $0 < T_1 < T_2 < \dots$,
- (iii) $\lim_{n \uparrow \infty} T_n = +\infty$.

The usual definition of a random point process is less restrictive. Condition (ii), in particular, is relaxed in the more general definition, where multiple points (simultaneous arrivals, for instance) are allowed. When condition (ii) holds, one speaks of a *simple point process*. Also, condition (iii) is not required in the more general definition, where it may occur that $P(\lim_{n \rightarrow \infty} T_n < \infty) > 0$: With positive probability there is an *explosion*, that is, an accumulation of events in finite time.

Conditions (ii) and (iii) fit the special case of homogeneous Poisson processes, the center of interest in this section.

The sequence $\{S_n\}_{n \geq 1}$ defined by

$$S_n = T_n - T_{n-1} \quad (1.1)$$

is called the *interevent* sequence, and sometimes, in the appropriate context, the *interarrival* sequence. For any interval $(a, b]$ in \mathbb{R}_+ ,

$$N((a, b]) \stackrel{\text{def}}{=} \sum_{n \geq 1} \mathbf{1}_{(a, b]}(T_n) \quad (1.2)$$

is an integer-valued random variable counting the events occurring in the time interval $(a, b]$. For typographical simplicity, it will be occasionally denoted by $N(a, b)$, omitting the external parentheses. For $t \geq 0$, let

$$N(t) \stackrel{\text{def}}{=} N(0, t].$$

In particular, $N(0) = 0$ and $N(a, b] = N(b) - N(a)$. Since the interval $(0, t]$ is closed on the right, the trajectories (or sample paths) $t \mapsto N(t, \omega)$ are right-continuous for almost all samples $\omega \in \Omega$. The sample paths are nondecreasing, have limits on the left at every time t , and jump one unit upwards at each event of the point process. The family of random variables $N = \{N(t)\}_{t \geq 0}$ is called the *counting process* of the point process $\{T_n\}_{n \geq 1}$. Since the sequence of events can be recovered from N , the latter also receives the appellation “point process.”

1.2 Counting Process of an HPP

There exist several equivalent definitions of a Poisson process. The one adopted here is the most practical.

Definition 1.2. Homogeneous Poisson Process

A point process N on the positive half-line is called an *homogeneous Poisson process* (HPP) with *intensity* $\lambda > 0$ if

(α) For all times t_i , $i \in [1, k]$, such that $0 \leq t_1 \leq t_2 \leq \dots \leq t_k$, the random variables $N(t_i, t_{i+1})$, $i \in [1, k - 1]$, are independent.

(β) For any interval $(a, b] \subset \mathbb{R}_+$, $N(a, b]$ is a Poisson random variable with mean $\lambda(b - a)$.

Thus, for all $k \geq 0$,

$$P(N(a, b) = k) = e^{-\lambda(b-a)} \frac{[\lambda(b-a)]^k}{k!}$$

and, in particular,

$$E[N(a, b)] = \lambda(b - a).$$

In this sense, λ is the average density of points.

Condition (α) is the property of *independence of increments* of Poisson processes. It implies in particular that for any interval $(a, b]$, the random variable $N(a, b)$ is independent of $(N(s), s \in (0, a])$. For this reason, Poisson processes are sometimes called *memoryless*. A more precise statement is this: The *increments* of homogeneous Poisson processes have no memory of the past.

Remark 1.1. *HPPs Are Simple*

The definition we adopted for random point processes does not allow multiple points or explosions. But suppose it did. It turns out that requirements (α) and (β) in Definition 1.2 suffice to prevent multiple points or explosions.

The proof is as follows. Since $E[N(a)] = \lambda a < \infty$ for all $a \geq 0$, $N(a) < \infty$ almost surely for all $a \geq 0$, and this is equivalent to $\lim_{n \uparrow \infty} T_n = \infty$ almost surely. Simplicity will follow from $P(D(a)) = 0$ for all $a \geq 0$, where

$$D(a) = \{\text{there exists multiple points in } (0, a)\}.$$

We prove this for $D = D(1)$ to simplify notation. The event

$$D_n = \left\{ N \left(\frac{i}{2^n}, \frac{i+1}{2^n} \right) \geq 2 \text{ for some } i \in [0, 2^n - 1] \right\}$$

decreases to D as n tends to infinity, and therefore, by monotone sequential continuity,

$$P(D) = \lim_{n \uparrow \infty} P(D_n) = 1 - \lim_{n \uparrow \infty} P(\bar{D}_n).$$

But

$$\begin{aligned} P(\bar{D}_n) &= P \left(\bigcap_{i=0}^{2^n-1} \left\{ N \left(\frac{i}{2^n}, \frac{i+1}{2^n} \right) \leq 1 \right\} \right) = \prod_{i=0}^{2^n-1} P \left(N \left(\frac{i}{2^n}, \frac{i+1}{2^n} \right) \leq 1 \right) \\ &= \prod_{i=0}^{2^n-1} e^{-\lambda 2^{-n}} (1 + \lambda 2^{-n}) = e^{-\lambda} (1 + \lambda 2^{-n})^{2^n}. \end{aligned}$$

The limit of the latter quantity is 1 as $n \uparrow \infty$, and therefore, $P(D) = 0$.

Theorem 1.1. *HPPs are i.i.d. Exponentials*

The interevent sequence $\{S_n\}_{n \geq 1}$ of an HPP on the positive half-line with the intensity $\lambda > 0$ is i.i.d., with exponential distribution of parameter λ .

The cumulative distribution function of an arbitrary interevent time is therefore,

$$P(S_n \leq t) = 1 - e^{-\lambda t}.$$

Recall that

$$E[S_n] = \lambda^{-1};$$

that is, the average number of events per unit of time equals the inverse average interevent time.

Proof. Suppose we can show that for any $n \geq 1$, the random vector $T = (T_1, \dots, T_n)$ admits the probability density function

$$f_T(t_1, \dots, t_n) = \lambda^n e^{-\lambda t_n} 1_C(t_1, \dots, t_n), \quad (1.3)$$

where $C \stackrel{\text{def}}{=} \{(t_1, \dots, t_n); 0 < t_1 < \dots < t_n\}$. Since

$$S_1 = T_1, \quad S_2 = T_2 - T_1, \dots, \quad S_n = T_n - T_{n-1},$$

the formula of smooth change of variables gives for the p.d.f. of $S = (S_1, \dots, S_n)$

$$f_S(s_1, \dots, s_n) = f_T(s_1, s_1 + s_2, \dots, s_1 + \dots + s_n) = \prod_{i=1}^n \{\lambda e^{-\lambda s_i} 1_{\{s_i > 0\}}\},$$

and this is the conclusion of Theorem 1.1. It remains to prove (1.3). The p.d.f. of T at $t = (t_1, \dots, t_n)$ is obtained as the limit as $h_1, \dots, h_n \in \mathbb{R}_+$ tend to 0 of the quantity

$$\frac{P(\cap_{i=1}^n \{T_i \in (t_i, t_i + h_i]\})}{\prod_{i=1}^n h_i}, \quad (1.4)$$

where it suffices to consider those (t_1, \dots, t_n) inside C since the points T_1, \dots, T_n are strictly ordered in increasing order. The event identity

$$\begin{aligned} & \cap_{i=1}^n \{T_i \in (t_i, t_i + h_i]\} \\ &= \{N(0, t_1] = 0\} \cap \left(\cap_{i=1}^{n-1} \{N(t_i, t_i + h_i] = 1, N(t_i + h_i, t_{i+1}] = 0\} \right) \\ & \quad \cap \{N(t_n, t_n + h_n] \geq 1\} \end{aligned}$$

holds for sufficiently small h_1, \dots, h_n , and therefore, the numerator of (1.4) equals

$$\begin{aligned} P(N(0, t_1] = 0) & \left(\prod_{i=1}^{n-1} P(N(t_i, t_i + h_i] = 1, N(t_i + h_i, t_{i+1}] = 0) \right) P(N(t_n, t_n + h_n] \geq 1) \\ &= e^{-\lambda t_1} \prod_{i=1}^{n-1} (e^{-\lambda h_i} \lambda h_i e^{-\lambda(t_{i+1} - t_i - h_i)}) (1 - e^{-\lambda h_n}) \\ &= \lambda^{n-1} e^{-\lambda t_n} h_1 \dots h_{n-1} (1 - e^{-\lambda h_n}). \end{aligned}$$

Dividing by $h_1 \dots h_n$ and taking the limit as h_1, \dots, h_n go to 0, we obtain $\lambda^n e^{-\lambda t_n}$. \square

1.3 Competing Poisson Processes

Let $\{T_n^1\}_{n \geq 1}$ and $\{T_n^2\}_{n \geq 1}$ be two independent HPPs on \mathbb{R}_+ with respective intensities $\lambda_1 > 0$ and $\lambda_2 > 0$. Their *superposition* is defined to be the sequence $\{T_n\}_{n \geq 1}$ formed by merging the two sequences $\{T_n^1\}_{n \geq 1}$ and $\{T_n^2\}_{n \geq 1}$ as explained in Figure 8.1.1. We shall prove that

- (i) the point processes $\{T_n^1\}_{n \geq 1}$ and $\{T_n^2\}_{n \geq 1}$ have no points in common, and
- (ii) the point process $\{T_n\}_{n \geq 1}$ is an HPP with intensity $\lambda = \lambda_1 + \lambda_2$.




Figure 8.1.1. Superposition, or sum, of two point processes

Indeed, defining N by

$$N(a, b] = N_1(a, b] + N_2(a, b],$$

we see that condition (α) of Definition 1.1 is satisfied, in view of the independence of N_1 and N_2 . Also, $N(a, b]$ being the sum of two independent Poisson random variables of mean $\lambda_1(b - a)$ and $\lambda_2(b - a)$ is a Poisson variable of mean $\lambda(b - a)$ where $\lambda = \lambda_1 + \lambda_2$, and therefore, condition (β) of Definition 1.1 is satisfied. This proves (ii). By Remark 1.1, N is simple, and therefore, (i) is true.

The above result can be extended to several—possibly an infinity of—homogeneous Poisson processes as follows:

Theorem 1.2. *Sum of Independent HPPs*

Let $\{N_i\}_{i \geq 1}$ be a family of independent HPPs with respective positive intensities $\{\lambda_i\}_{i \geq 1}$. Then

- (i) two distinct HPPs of this family have no points in common, and

(ii) if

$$\sum_{i=1}^{\infty} \lambda_i = \lambda < \infty,$$

then

$$N(t) \stackrel{\text{def}}{=} \sum_{i=1}^{\infty} N_i(t)$$

defines the counting process of an HPP with intensity λ .

Proof. Assertion (i) has already been proven. Observe that for all $t \geq 0$, $N(t)$ is almost surely finite, since

$$E[N(t)] = \sum_{i=1}^{\infty} E[N_i(t)] = \left(\sum_{i=1}^{\infty} \lambda_i \right) t < \infty.$$

In particular, $N(a, b]$ is almost surely finite for all $(a, b] \subset \mathbb{R}_+$. The proof of lack of memory of N is the same as in the case of two superposed Poisson processes. Finally, $N(a, b]$ is a Poisson random variable of mean $\lambda(b - a)$, since

$$\begin{aligned} P(N(a, b] = k) &= \lim_{n \uparrow \infty} P\left(\sum_{i=1}^n N_i(a, b] = k\right) \\ &= \lim_{n \uparrow \infty} e^{-(\sum_{i=1}^n \lambda_i(b-a))} \frac{[\sum_{i=1}^n \lambda_i(b-a)]^k}{k!} \\ &= e^{-\lambda(b-a)} \frac{[\lambda(b-a)]^k}{k!}. \quad \square \end{aligned}$$

The next result is central to the theory of Poisson systems (Chapter 9). It is called the *competition theorem* because it features HPPs competing for the production of the first event.

Theorem 1.3. Competition Theorem

In the situation of Theorem 1.2, where $\sum_{i=1}^{\infty} \lambda_i = \lambda < \infty$, denote by Z the first event time of $N = \sum_{i=1}^{\infty} N_i$ and by J the index of the HPP responsible for it; in particular, Z is the first event of N_J . Then

$$P(J = i, Z \geq a) = P(J = i)P(Z \geq a) = \frac{\lambda_i}{\lambda} e^{-\lambda a}. \quad (1.5)$$

In particular, J and Z are independent, $P(J = i) = \frac{\lambda_i}{\lambda}$, and Z is exponential with mean λ^{-1} .

Proof. A. We first prove the result for a finite number of processes. We have to show that if X_1, \dots, X_K are K independent exponential variables with means $\lambda_1^{-1}, \dots, \lambda_K^{-1}$, if $Z_K = \inf(X_1, \dots, X_K)$ and if J_K is defined by $X_{J_K} = Z_K$, then

$$P(J_K = i, Z_K \geq a) = \frac{\lambda_i}{\lambda_1 + \dots + \lambda_K} \exp\{-(\lambda_1 + \dots + \lambda_K)a\}. \quad (1.6)$$




Figure 8.1.2. Competition among four point processes

First observe that

$$P(Z_K \geq a) = P(\cap_{j=1}^K \{X_j \geq a\}) = \prod_{j=1}^K P(X_j \geq a) = \prod_{j=1}^K e^{-\lambda_j a} = e^{-(\lambda_1 + \dots + \lambda_K)a}.$$

Defining $U = \inf\{X_2, \dots, X_K\}$, we have (use the result of Exercise 4.2, Chapter 1)

$$\begin{aligned} P(J_K = 1, Z_K \geq a) &= P(a \leq X_1 < U) \\ &= \int_a^\infty P(U > x) \lambda_1 e^{-\lambda_1 x} dx = \int_a^\infty e^{-(\lambda_2 + \dots + \lambda_K)x} \lambda_1 e^{-\lambda_1 x} dx \\ &= \frac{\lambda_1}{\lambda_1 + \dots + \lambda_K} e^{-(\lambda_1 + \dots + \lambda_K)a}. \end{aligned}$$

This gives (1.6). Letting $a \rightarrow \infty$ yields $P(J_K = 1) = \frac{\lambda_1}{\lambda_1 + \dots + \lambda_K}$. This, together with (1.6) and the expression for $P(Z_K \geq a)$ gives (1.5), for $i = 1$, without loss of generality.

B. Suppose the result true for a finite number of HPPs. Since the event $\{J_K = 1, Z_K \geq a\}$ decreases to $\{J = 1, Z \geq a\}$ as $K \uparrow \infty$, we have

$$P(J = 1, Z \geq a) = \lim_{K \uparrow \infty} P(J_K = 1, Z_K \geq a),$$

from which (1.5) follows, using the result of part A of the proof. \square

2 Distribution of a Continuous-Time HMC

2.1 Transition Semigroup

The traditional approach to continuous-time Markov chains is based on the *transition semigroup*, and the principal mathematical object is then the *infinitesimal generator*. The tran-

sition semigroup is the continuous-time analogue of the iterates of the transition matrix in discrete time. The infinitesimal generator, however, has no analogue because it is essentially a continuous-time notion involving derivatives.

The transition semigroup approach is mainly analytical, and we shall propose in Chapter 9 a sample path approach that describes continuous-time Markov chains in terms of a stochastic differential equation, the analogue of the recurrence equation associated with a discrete-time Markov chain (Theorem 2.1 of Chapter 2). These two approaches are complementary. We begin with the semigroup approach because it makes the continuous-time theory look more like a natural extension of the discrete-time theory.

We now proceed to give the basic definitions concerning continuous-time HMCs. Let E be a countable set, called the state space, and let $\{X(t)\}_{t \geq 0}$ be an E -valued stochastic process, that is, a family of random variables $X(t)$ indexed by \mathbb{R}_+ and taking their values in E . The probability distribution of $\{X(t)\}_{t \geq 0}$ consists of the data $P(X(t_1) = i_1, \dots, X(t_k) = i_k)$, for all $t_1, \dots, t_k \geq 0$, and all $i_1, \dots, i_k \in E$.

Definition 2.1. *Continuous-Time Homogeneous Markov Chains*

The E -valued stochastic process $\{X(t)\}_{t \geq 0}$ is called a *continuous-time Markov chain* if for all $i, j, i_1, \dots, i_k \in E$, all $t, s \geq 0$, and all $s_1, \dots, s_k \geq 0$ with $s_\ell \leq s$ for all $\ell \in [1, k]$,

$$P(X(t+s) = j \mid X(t) = i, X(s_1) = i_1, \dots, X(s_k) = i_k) = P(X(t+s) = j \mid X(s) = i), \quad (2.1)$$

whenever both sides are well-defined. This continuous-time Markov chain is called *homogeneous* if the right-hand side of (2.1) is independent of s .

Let then

$$\mathbf{P}(t) \stackrel{\text{def}}{=} \{p_{ij}(t)\}_{i,j \in E}, \quad (2.2)$$

where

$$p_{ij}(t) \stackrel{\text{def}}{=} P(X(t+s) = j \mid X(s) = i). \quad (2.3)$$

The family $\{\mathbf{P}(t)\}_{t \geq 0}$ is the *transition semigroup* of the continuous-time HMC.

In the same way as for discrete-time Markov chains, one obtains the Chapman–Kolmogorov equation

$$p_{ij}(t+s) = \sum_{k \in E} p_{ik}(t)p_{kj}(s),$$

that is, in compact form,

$$\mathbf{P}(t+s) = \mathbf{P}(t)\mathbf{P}(s). \quad (2.4)$$

Also, clearly,

$$\mathbf{P}(0) = I, \quad (2.5)$$

where I is the identity matrix.

The distribution at time t of $X(t)$ is the vector $\mu(t) = \{\mu_i(t)\}_{i \in E}$ where $\mu_i(t) = P(X(t) = i)$. It is obtained from the initial distribution by the formula

$$\mu(t)^T = \mu(0)^T \mathbf{P}(t). \quad (2.6)$$

Also, for all t_1, \dots, t_k such that $0 \leq t_1 \leq t_2 \leq \dots \leq t_k$, and for all states i_0, i_1, \dots, i_k ,

$$P(\bigcap_{j=1}^k \{X(t_j) = i_j\}) = \sum_{i_0 \in E} P(X(0) = i_0) \prod_{j=1}^k p_{i_{j-1}i_j}(t_j - t_{j-1}). \tag{2.7}$$

Formulas (2.6) and (2.7) are proven in the same manner as the analogue results in discrete time. Formula (2.7) shows in particular that the probability distribution of a continuous-time HMC is entirely determined by its initial distribution and its transition semigroup.

Notation Recall that $P_i(\cdot)$ is an abbreviated notation for $P(\cdot \mid X(0) = i)$, where i is a state. If μ is a probability distribution on E , then $P_\mu(A) = \sum_{i \in E} \mu(i)P(A \mid X(0) = i)$ is the probability of A when the distribution of the initial state is μ .

Example 2.1. Poisson Is Markovian

Let N be an HPP on the positive half-line with the intensity $\lambda > 0$. The counting process $\{N(t)\}_{t \geq 0}$ is a continuous-time HMC. Indeed, writing $C = \{N(s_1) = i_1, \dots, N(s_k) = i_k\}$, we have

$$\begin{aligned} P(N(t+s) = j \mid N(s) = i, C) &= \frac{P(N(t+s) = j, N(s) = i, C)}{P(N(s) = i, C)} \\ &= \frac{P(N(s, s+t] = j-i, N(s) = i, C)}{P(N(s) = i, C)}. \end{aligned}$$

But $N(s, s+t]$ is independent of $N(s)$ and of $N(s_\ell)$ when $s_\ell \leq s$, and therefore,

$$P(N(s, s+t] = j-i, N(s) = i, C) = P(N(s, s+t] = j-i)P(N(s) = i, C),$$

so that

$$P(N(t+s) = j \mid N(s) = i, C) = P(N(s, s+t] = j-i).$$

Similarly,

$$P(N(t+s) = j \mid N(s) = i) = P(N(s, s+t] = j-i).$$

Therefore, $\{N(t)\}_{t \geq 0}$ is a continuous-time Markov chain. It is homogeneous, since $P(N(s, s+t] = j-i)$ does not depend on s . The transition semigroup $\{\mathbf{P}(t)\}_{t \geq 0}$ is given by

$$p_{ij}(t) = P(N(0, t] = j-i) = e^{-\lambda t} \frac{(\lambda t)^{j-i}}{(j-i)!},$$

for $j \geq i$. Otherwise, if $j > i$, then $p_{ij}(t) = 0$. ◇

Example 2.2. Flip-Flop

Let N be the Poisson process of intensity λ considered in the previous example, and define the *flip-flop process* $\{X(t)\}_{t \geq 0}$ with state space $E = \{+1, -1\}$ by

$$X(t) = X(0) \times (-1)^{N(t)},$$

where $X(0)$ is an E -valued random variable independent of the counting process N . Thus the flip-flop process switches between -1 and $+1$ at each event of N . The value $X(t + s)$ depends on $N(s, s + t]$ and $X(s)$. Also, $N(s, s + t]$ is independent of $X(0), N(s_1), \dots, N(s_k)$ when $s_\ell \leq s$ for all $\ell \in [1, k]$, and the latter random variables determine $X(s_1), \dots, X(s_k)$. Therefore, $X(t + s)$ is independent of $X(s_1), \dots, X(s_k)$ given $X(s)$, that is, $\{X(t)\}_{t \geq 0}$ is a Markov chain. Moreover,

$$\begin{aligned} P(X(t + s) = 1 \mid X(s) = -1) &= P(N(s, s + t] = \text{odd}) \\ &= \sum_{k=0}^{\infty} e^{-\lambda t} \frac{(\lambda t)^{2k+1}}{(2k + 1)!} = \frac{1}{2}(1 - e^{-2\lambda t}), \end{aligned}$$

that is, $p_{-1,+1}(t) = \frac{1}{2}(1 - e^{-2\lambda t})$. Similar computations give for the transition semigroup

$$\mathbf{P}(t) = \frac{1}{2} \begin{pmatrix} 1 + e^{-2\lambda t} & 1 - e^{-2\lambda t} \\ 1 - e^{-2\lambda t} & 1 + e^{-2\lambda t} \end{pmatrix}. \quad \diamond$$

Definition 2.2. Uniform Markov Chains

Let $\{\hat{X}_n\}_{n \geq 0}$ be a discrete-time HMC with countable state space E and transition matrix $\mathbf{K} = \{k_{ij}\}_{i,j \in E}$ and let $\{T_n\}_{n \geq 1}$ be an HPP on \mathbb{R}_+ with intensity $\lambda > 0$ and associated counting process N . Suppose that $\{\hat{X}_n\}_{n \geq 0}$ and N are independent. The process $\{X(t)\}_{t \geq 0}$ with values in E , defined by

$$X(t) = \hat{X}_{N(t)} \tag{2.8}$$

(see Figure 8.2.1), is called a *uniform Markov chain*. The Poisson process N is called the *clock*, and the chain $\{\hat{X}_n\}_{n \geq 0}$ is called the *subordinated chain*.




Figure 8.2.1. Uniform Markov chain

Observe that $X(T_n) = \hat{X}_n$ for all $n \geq 0$. Observe also that the discontinuity times of the uniform chain are all events of N but that not all events of N are discontinuity times, since it may well occur that $\hat{X}_{n-1} = \hat{X}_n$ (a transition of type $i \rightarrow i$ of the subordinated chain).

The process $\{X(t)\}_{t \geq 0}$ is a continuous-time HMC (Problem 8.2.6). Its transition semigroup is

$$\mathbf{P}(t) = \sum_{n=0}^{\infty} e^{-\lambda t} \frac{(\lambda t)^n}{n!} \mathbf{K}^n, \tag{2.9}$$

that is,

$$p_{ij}(t) = \sum_{n=0}^{\infty} e^{-\lambda t} \frac{(\lambda t)^n}{n!} k_{ij}(n). \tag{2.10}$$

Indeed,

$$\begin{aligned} P_i(X(t) = j) &= P_i(\hat{X}_{N(t)} = j) = \sum_{n=0}^{\infty} P_i(\hat{X}_n = j, N(t) = n) \\ &= \sum_{n=0}^{\infty} P_i(\hat{X}_n = j) P_i(N(t) = n). \end{aligned}$$

2.2 Infinitesimal Generator

Let $\{\mathbf{P}(t)\}_{t \geq 0}$ be a transition semigroup on E , i.e., for each $t, s \geq 0$,

- (a) $\mathbf{P}(t)$ is a stochastic matrix,
- (b) $\mathbf{P}(0) = I$,
- (c) $\mathbf{P}(t + s) = \mathbf{P}(t)\mathbf{P}(s)$.

Suppose, moreover, that the semigroup is *continuous* at the origin, that is,

$$\lim_{h \downarrow 0} \mathbf{P}(h) = \mathbf{P}(0) = I, \tag{2.11}$$

where the convergence therein is pointwise and for each entry. In Problem 8.2.1, the reader is invited to prove that continuity at the origin implies continuity at any time $t \geq 0$, that is,

$$\lim_{h \rightarrow 0} p_{ij}(t + h) = p_{ij}(t),$$

for all states i, j .

The result to follow is purely analytical: it does not require $\{\mathbf{P}(t)\}_{t \geq 0}$ to be the transition semigroup of some continuous-time HMC.

Theorem 2.1. Local Characteristics

Let $\{\mathbf{P}(t)\}_{t \geq 0}$ be a continuous transition semigroup on the countable state space E . For any state i , there exists

$$q_i \stackrel{\text{def}}{=} \lim_{h \downarrow 0} \frac{1 - p_{ii}(h)}{h} \in [0, \infty], \tag{2.12}$$

and for any pair i, j of different states, there exists

$$q_{ij} \stackrel{\text{def}}{=} \lim_{h \downarrow 0} \frac{p_{ij}(h)}{h} \in [0, \infty). \tag{2.13}$$

Proof. For all $t \geq 0$, all $n \geq 1$, we have $\mathbf{P}(t) = [\mathbf{P}(\frac{t}{n})]^n$, and therefore $p_{ii}(t) \geq [p_{ii}(\frac{t}{n})]^n$ for all $i \in E$. Since $\lim_{h \downarrow 0} p_{ii}(h) = 1$, there exists $\epsilon > 0$ such that $p_{ii}(h) > 0$ for all $h \in [0, \epsilon]$. For n sufficiently large, $\frac{t}{n} \in [0, \epsilon]$. Therefore, for all $t \geq 0$, $p_{ii}(t) > 0$, and one can define

$$f_i(t) = -\log p_{ii}(t).$$

The function f_i is real-valued, nonnegative, and such that $\lim_{h \downarrow 0} f_i(h) = 0$. Also, from $\mathbf{P}(t)\mathbf{P}(s) = \mathbf{P}(t+s)$, we see that $p_{ii}(t+s) \geq p_{ii}(t)p_{ii}(s)$, and therefore, f_i is subadditive, that is

$$f_i(t+s) \leq f_i(t) + f_i(s)$$

for all $s, t \in \mathbb{R}_+$. Define the (possibly infinite) nonnegative real number

$$q_i = \sup_{t>0} \frac{f_i(t)}{t}.$$

Then (see Theorem 1.11 of the Appendix)

$$\lim_{h \downarrow 0} \frac{f_i(h)}{h} = q_i. \quad (2.14)$$

From (2.14),

$$\lim_{h \downarrow 0} \frac{1 - p_{ii}(h)}{h} = \lim_{h \downarrow 0} \frac{1 - e^{-f_i(h)}}{f_i(h)} \frac{f_i(h)}{h} = q_i,$$

and this proves the first equality in (2.12). For future reference, we shall observe that

$$\frac{1 - p_{ii}(h)}{h} = \frac{1 - e^{-f_i(h)}}{f_i(h)} \frac{f_i(h)}{h} \leq \frac{f_i(h)}{h},$$

and therefore for all $h > 0$,

$$\frac{1 - p_{ii}(h)}{h} \leq q_i. \quad (2.15)$$

It now remains to prove (2.13).

Take two different states i and j . Since $p_{ii}(t)$ and $p_{jj}(t)$ tend to 1 as $t > 0$ tends to 0, for any $c \in (\frac{1}{2}, 1)$, there exists $\delta > 0$ such that for $t \in [0, \delta]$, $p_{ii}(t) > c$ and $p_{jj}(t) > c$. Let $n > 0$ be an integer and $h > 0$ be such that $0 \leq nh \leq \delta$. Denote by $\{X_n\}_{n \geq 0}$ the discrete-time HMC defined by $X_n = X(nh)$, with transition matrix $\mathbf{P}(h)$. One way to go from $X_0 = i$ to $X_n = j$ is to go from $X_0 = i$ to $X_r = i$, $r = 0, 1, \dots, n-1$, without passing through state j meanwhile, then to go from $X_r = i$ to $X_{r+1} = j$, and then from $X_{r+1} = j$ to $X_n = j$. The paths corresponding to different values of r are different, but they do not exhaust the possibilities of going from $X_0 = i$ to $X_n = j$. Therefore,

$$p_{ij}(nh) \geq \sum_{r=0}^{n-1} P(\cap_{\ell=r}^{n-1} \{X_\ell \neq j\}, X_r = i \mid X_0 = i) p_{ij}(h) P(X_n = j \mid X_{r+1} = j).$$

The parameters $\delta, n,$ and h are such that $P(X_n = j \mid X_{r+1} = j) \geq c$. Also

$$\begin{aligned} &P(\cap_{\ell=1}^{r-1} \{X_\ell \neq j\}, X_r = i \mid X_0 = i) \\ &= P(X_r = i \mid X_0 = i) - \sum_{k < r} P(\cap_{\ell=1}^{k-1} \{X_\ell \neq j\}, X_k = j \mid X_0 = i)P(X_r = i \mid X_k = j) \\ &\geq c - (1 - c) \sum_{k < r} P(\cap_{\ell=1}^{k-1} \{X_\ell \neq j\}, X_k = j \mid X_0 = i) \geq c - (1 - c) = 2c - 1, \end{aligned}$$

where we observed that for $i \neq j,$

$$P(X_r = i \mid X_k = j) + P(X_r = j \mid X_k = j) \leq 1,$$

and therefore,

$$P(X_r = i \mid X_k = j) \leq 1 - P(X_r = j \mid X_k = j) \leq 1 - c.$$

Therefore,

$$p_{ij}(nh) \geq c(2c - 1)np_{ij}(h).$$

Let now $t < \delta$ and $h < \delta,$ and take for n the integer part of t/h . From the last inequality we obtain

$$\frac{p_{ij}(h)}{h} \leq \frac{1}{c(2c - 1)} \frac{p_{ij}(nh)}{nh},$$

and since $\lim_{h \downarrow 0} nh = t,$ we see that $\lim_{h \downarrow 0} \frac{p_{ij}(nh)}{nh} = \frac{p_{ij}(t)}{t},$ so that from the last inequality,

$$\limsup_{h \downarrow 0} \frac{p_{ij}(h)}{h} \leq \frac{1}{c(2c - 1)} \frac{p_{ij}(t)}{t} < \infty,$$

which in turn gives

$$\limsup_{h \downarrow 0} \frac{p_{ij}(h)}{h} \leq \frac{1}{c(2c - 1)} \liminf_{t \downarrow 0} \frac{p_{ij}(t)}{t} < \infty.$$

Since c can be chosen arbitrarily close to 1, we have

$$\limsup_{h \downarrow 0} \frac{p_{ij}(h)}{h} \leq \liminf_{t \downarrow 0} \frac{p_{ij}(t)}{t} < \infty,$$

and this implies the existence of $\lim_{h \downarrow 0} \frac{p_{ij}(h)}{h}$ and the finiteness of this limit. □

Let for each state i

$$q_{ii} \stackrel{\text{def}}{=} -q_i. \tag{2.16}$$

Definition 2.3. *Infinitesimal Generator*

The numbers q_{ij} are called the *local characteristics* of the semigroup, or of the corresponding continuous-time HMC. The matrix

$$\mathbf{A} = \{q_{ij}\}_{i,j \in E} \tag{2.17}$$

is called the *infinitesimal generator* of the semigroup, or of the continuous-time HMC .

In compact notation,

$$\mathbf{A} = \lim_{h \downarrow 0} \frac{\mathbf{P}(h) - \mathbf{P}(0)}{h}, \quad (2.18)$$

where the meaning of (2.18) is given by (2.12) and (2.13). Thus, in this sense, the infinitesimal generator \mathbf{A} is the derivative at 0 of the matrix function $t \mapsto \mathbf{P}(t)$.

Example 2.3. *Infinitesimal Generator of the Uniform HMC*

From the expression (2.9), or (2.10), of the transition semigroup of the uniform HMC of Definition 2.2, we easily obtain its infinitesimal generator (Problem 8.2.7)

$$\mathbf{A} = \lambda(\mathbf{K} - \mathbf{I}), \quad (2.19)$$

that is,

$$q_i = \lambda(1 - k_{ii}), \quad (2.20)$$

and for $i \neq j$,

$$q_{ij} = \lambda k_{ij}. \quad (2.21)$$

◇

Definition 2.4. *Stability and Conservation*

If for all states i ,

$$q_i < \infty, \quad (2.22)$$

the semigroup $\{\mathbf{P}(t)\}$ is called *stable*. If for all states i ,

$$q_i = \sum_{\substack{j \in E \\ j \neq i}} q_{ij}, \quad (2.23)$$

it is called *conservative*.

The reason for the last appellation comes from the conservation equality

$$\sum_{j \in E} p_{ij}(h) = 1,$$

or equivalently,

$$\frac{1 - p_{ii}(h)}{h} = \sum_{\substack{j \in E \\ j \neq i}} \frac{p_{ij}(h)}{h},$$

which yields

$$q_i = \lim_{h \downarrow 0} \sum_{\substack{j \in E \\ j \neq i}} \frac{p_{ij}(h)}{h}.$$

And if the interchange of summation and limit is allowed, we obtain (2.23).

Interchange of sums and limits is always possible if E is finite. In this case (2.23) holds, and consequently, (2.22) holds because q_{ij} is finite for all pairs of different sites i, j . Also,

the semigroup of an uniform HMC is stable and conservative, even when E is infinite (use (2.20) and (2.21)).

As a matter of fact, a very general class of Markov chains, namely *regular jump* Markov chains, are stable and conservative.

Definition 2.5. *Regular Jump HMC*

A stochastic process $\{X(t)\}_{t \geq 0}$ taking its values in the (not necessarily countable) state space E is called a *jump process* if for almost all $\omega \in \Omega$ and all $t \geq 0$, there exists $\epsilon(t, \omega) > 0$ such that

$$X(t + s, \omega) = X(t, \omega) \text{ for all } s \in [t, t + \epsilon(t, \omega)).$$

It is called a *regular jump process* if in addition, for almost all $\omega \in \Omega$, the set $A(\omega)$ of discontinuities of the function $t \mapsto X(t, \omega)$ is σ -discrete, that is, for all $c \geq 0$,

$$|A(\omega) \cap [0, c]| < \infty,$$

where the notation $|B|$ represents the number of elements in the set B . A *regular jump homogeneous Markov chain* is by definition a continuous-time HMC that is also a regular jump process.

Observe that for a jump process (not necessarily regular), there exists a sequence of times $\{\tau_n\}_{n \geq 0}$ where

$$\tau_0 = 0 < \tau_1 < \tau_2 < \tau_3 < \dots$$

and a sequence $\{X_n\}_{n \geq 0} \in E$ such that

$$X(t) = X_n \text{ if } \tau_n \leq t < \tau_{n+1}.$$

This describes $\{X(t)\}_{t \geq 0}$ on the interval $[0, \tau_\infty)$, where

$$\tau_\infty = \lim \uparrow \tau_n$$

is the *explosion time*. If, moreover, the process is regular, then $\tau_\infty = \infty$, and $\{X(t)\}_{t \geq 0}$ is right-continuous.

Theorem 2.2.

A regular jump HMC is stable and conservative.

Proof. Postponed to Section 1 of Chapter 9.

There exist bona fide continuous-time HMCs that are not regular jump HMCs, however, for which all the states are unstable! Fortunately, such Markov chains are of little concern for applications and are sometimes called pathological, although it is not hard to construct one. The reader is directed to (Chung, 1960) for a thorough discussion of this issue.

When the chain is stable and conservative, we have

$$P(X(t + h) = i \mid X(t) = i) = 1 - q_i h + o(h), \quad (2.24)$$

and if $i \neq j$,

$$P(X(t+h) = j \mid X(t) = i) = q_{ij}h + o(h). \quad (2.25)$$

(Recall that the $o(h)$ symbol represents a function defined in a neighborhood of 0 for which $\lim_{h \rightarrow 0} \frac{|o(h)|}{h} = 0$.)

Definition 2.6. *Continuous-Time Birth-and-Death Processes*

A continuous-time birth-and-death process is a regular jump HMC taking its values in \mathbb{N} , and with an infinitesimal generator of the form

$$q_{i,i+1} = \lambda_i, \quad q_{i,i-1} = \mu_i \mathbf{1}_{\{i \geq 1\}}, \quad (2.26)$$

and $q_{ij} = 0$ if $j \notin \{i-1, i, i+1\}$.

The parameters λ_i and μ_i are the *birth* and *death* parameters, respectively. Indeed, in view of (2.25),

$$P(X(t+h) = i+1 \mid X(t) = i) = \lambda_i h + o(h) \quad (2.27)$$

and

$$P(X(t+h) = i-1 \mid X(t) = i) = \mu_i \mathbf{1}_{\{i \geq 1\}} h + o(h). \quad (2.28)$$

By conservation,

$$P(X(t+h) = i \mid X(t) = i) = 1 - (\lambda_i + \mu_i \mathbf{1}_{\{i \geq 1\}})h + o(h).$$

Birth-and-death processes are important models in biology (where the terminology comes from, obviously), but also in operations research, and in particular in queuing theory, where they appear as $M/M/1/\infty$, $M/M/K/0$ queues, among many other models of waiting lines (see Chapter 9).

Observe that we have included in the definition the regularity of birth-and-death processes. However, in modeling, it is the birth-and-death parameters that are given. The question is, Given the birth-and-death parameters, does there exist a regular jump HMC satisfying (2.27) and (2.28)? This problem is nontrivial when the birth-and-death parameters are not uniformly bounded. This issue will be examined in Section 4.

3 Kolmogorov's Differential Systems

3.1 Finite State Space

In view of the semigroup properties, for all $t \geq 0$ and all $h \geq 0$

$$\frac{\mathbf{P}(t+h) - \mathbf{P}(t)}{h} = \mathbf{P}(t) \frac{\mathbf{P}(h) - I}{h} = \frac{\mathbf{P}(h) - I}{h} \mathbf{P}(t). \quad (3.1)$$

Therefore, if the passage to the limit in (3.1) is allowed, which is the case when the state space E is finite, we obtain the differential system

$$\frac{d}{dt} \mathbf{P}(t) = \mathbf{P}(t) \mathbf{A} = \mathbf{A} \mathbf{P}(t), \quad (3.2)$$

where \mathbf{A} is the infinitesimal generator. The equation

$$\frac{d}{dt} \mathbf{P}(t) = \mathbf{A} \mathbf{P}(t) \quad (3.3)$$

can be written explicitly. For all $i, j \in E$,

$$\frac{d}{dt} p_{ij}(t) = -q_i p_{ij}(t) + \sum_{\substack{k \in E \\ k \neq i}} q_{ik} p_{kj}(t).$$

System (3.3) is Kolmogorov's *backward* differential system. The *forward* differential system is

$$\frac{d}{dt} \mathbf{P}(t) = \mathbf{P}(t) \mathbf{A}, \quad (3.4)$$

that is, for all $i, j \in E$,

$$\frac{d}{dt} p_{ij}(t) = -p_{ij}(t) q_j + \sum_{\substack{k \in E \\ k \neq j}} p_{ik}(t) q_{kj}.$$

When the state space is finite, a solution of (3.3) or (3.4) with the initial condition $\mathbf{P}(0) = \mathbf{I}$ is

$$\mathbf{P}(t) = e^{t\mathbf{A}}, \quad (3.5)$$

where the exponential of a finite-dimensional matrix \mathbf{C} is defined by

$$e^{\mathbf{C}} \stackrel{\text{def}}{=} \sum_{n=0}^{\infty} \frac{\mathbf{C}^n}{n!} \quad (3.6)$$

(see Theorem 2.1 of the Appendix). This is the *unique* solution with the given initial data; see (Hirsch and Smale, 1974) for more information on linear systems of differential equations.

Example 3.1. Diagonalizable Infinitesimal Generator

Suppose that the infinitesimal generator has the form

$$\mathbf{A} = \mathbf{V} \Lambda \mathbf{U}^T,$$

where $\Lambda = \text{diag}\{\lambda_1, \dots, \lambda_r\}$ and $r = |E|$. Then (see Example 2.1 of the Appendix)

$$\mathbf{P}(t) = e^{t\mathbf{A}} = \mathbf{V} \text{diag}\{e^{t\lambda_1}, \dots, e^{t\lambda_r}\} \mathbf{U}^T.$$

For the general case, the reader is directed to (Hirsch and Smale, 1974). ◇

3.2 General Case

When the state space is not finite, difficulties may arise in the passage to the limit $h \downarrow 0$ in (3.1) because of the possibly infinite sums involved. However, for the backward system, there is a positive result:

Theorem 3.1. *Backward Kolmogorov System*

If the continuous semigroup $\{\mathbf{P}(t)\}_{t \geq 0}$ is stable and conservative, Kolmogorov's backward differential system (3.3) is satisfied.

Proof. Write the equality involving the extreme terms of (3.1):

$$\frac{p_{ij}(t+h) - p_{ij}(t)}{h} = \frac{p_{ii}(h) - 1}{h} p_{ij}(t) + \sum_{\substack{k \in E \\ k \neq i}} \frac{p_{ik}(h)}{h} p_{kj}(t). \quad (3.7)$$

It will be shown that the limit as $h \downarrow 0$ of the sum appearing in the above equality exists and is equal to

$$\sum_{\substack{k \in E \\ k \neq i}} q_{ik} p_{kj}(t).$$

This means that $p_{ij}(t)$ has a right-hand derivative equal to

$$-q_i p_{ij}(t) + \sum_{\substack{k \in E \\ k \neq i}} q_{ik} p_{kj}(t).$$

Since for all $k, j \in E$, $p_{kj}(t)$ is a continuous function, and since $\sum_{\substack{k \in E \\ k \neq i}} q_{ik} < \infty$, the right-hand derivative is a continuous function, by the dominated convergence theorem for series (see Theorem 1.6 of the Appendix). It is therefore also the left-hand derivative, since a continuous function with a continuous right-hand derivative is differentiable.

To prove that the limit as $h \downarrow 0$ of the sum in (3.7) exists, we start from the inequality

$$\sum_{\substack{k \in E \\ k \neq i}} \frac{p_{ik}(h)}{h} p_{kj}(t) \geq \sum_{\substack{k=1 \\ k \neq i}}^N \frac{p_{ik}(h)}{h} p_{kj}(t),$$

where N is an arbitrary integer (E is identified with the set of integers). Therefore, the \liminf as $h \downarrow 0$ of the left-hand side is larger than or equal to the limit of the right-hand side. Letting then N go to ∞ after the passage to the limit $h \downarrow 0$ yields

$$\liminf_{h \downarrow 0} \sum_{\substack{k \in E \\ k \neq i}} \frac{p_{ik}(h)}{h} p_{kj}(t) \geq \sum_{\substack{k \in E \\ k \neq i}} q_{ik} p_{kj}(t). \quad (3.8)$$

Next, observe that for $N > i$,

$$\begin{aligned} \sum_{\substack{k \in E \\ k \neq i}} \frac{p_{ik}(h)}{h} p_{kj}(t) &\leq \sum_{\substack{k=0 \\ k \neq i}}^N \frac{p_{ik}(h)}{h} p_{kj}(t) + \sum_{k > N} \frac{p_{ik}(h)}{h} \\ &= \sum_{\substack{k=0 \\ k \neq i}}^N \frac{p_{ik}(h)}{h} p_{kj}(t) + \frac{1 - p_{ii}(h)}{h} - \sum_{\substack{k=0 \\ k \neq i}}^N \frac{p_{ik}(h)}{h}, \end{aligned}$$

and therefore,

$$\limsup_{h \downarrow 0} \sum_{\substack{k \in E \\ k \neq i}} \frac{p_{ik}(h)}{h} p_{kj}(t) \leq \sum_{\substack{k=0 \\ k \neq i}}^N q_{ik} p_{kj}(t) + q_i - \sum_{\substack{k=0 \\ k \neq i}}^N q_{ik}.$$

By letting N go to ∞ , and in view of the stability and conservation hypothesis, we have

$$\limsup_{h \downarrow 0} \sum_{\substack{k \in E \\ k \neq i}} \frac{p_{ik}(h)}{h} p_{kj}(t) \leq \sum_{\substack{k \in E \\ k \neq i}} q_{ik} p_{kj}(t). \quad (3.9)$$

Equality (3.3) then follows from (3.8) and (3.9). \square

For the forward system, in the absence of regularity assumptions on the trajectories of the Markov chain, the result is considerably less general.

Theorem 3.2. *Forward Kolmogorov System*

Under the assumptions of Theorem 3.1, and if, moreover, for all states i and all $t \geq 0$,

$$\sum_{k \in E} p_{ik}(t) q_k < \infty, \quad (3.10)$$

then Kolmogorov's forward differential system (3.4) is satisfied.

Proof. Using estimate (2.15), we see that

$$\frac{p_{ki}(h)}{h} \leq \frac{1 - p_{kk}(h)}{h} \leq q_k.$$

The first equality in (3.1) reads

$$\frac{p_{ij}(t+h) - p_{ij}(t)}{h} = \sum_{\substack{k \in E \\ k \neq i}} p_{ik}(t) \frac{p_{kj}(h) - \delta_{kj}}{h}.$$

Each term of the series on the right-hand side is bounded by the corresponding term of the convergent series $\sum_{k \in E} p_{ik}(t) q_k$. Passage to the limit $h \downarrow 0$ is therefore allowed by Lebesgue's dominated convergence theorem, and this leads to the announced result, again after observing that the continuous right-hand derivative of a continuous function is also the left-hand derivative. \square

The distribution at time t of the chain is the column vector $\mu(t) = \{\mu_i(t)\}_{i \in E} = \{P(X(t) = i)\}_{i \in E}$. It satisfies, for all $t, s \geq 0$,

$$\mu^T(t+s) = \mu^T(t) \mathbf{P}(s), \quad (3.11)$$

that is, for all $i, j \in E$,

$$\mu_i(t+s) = \sum_{j \in E} \mu_j(t) p_{ji}(s).$$

Theorem 3.3. *Global Balance*

Under the assumptions of Theorem 3.1, and if, moreover, for all $t \geq 0$,

$$\sum_{i \in E} q_i \mu_i(t) < \infty, \quad (3.12)$$

then the Kolmogorov's global differential system

$$\frac{d}{dt} \mu^T(t) = \mu^T(t) \mathbf{A} \quad (3.13)$$

is satisfied, that is, for all $i \in E$,

$$\frac{d}{dt} \mu_i(t) = -\mu_i(t) q_i + \sum_{\substack{j \in E \\ j \neq i}} \mu_j(t) q_{ji}.$$

Proof. From (3.11), for all $i \in E$, all $t \geq 0$, and all $h > 0$,

$$\mu_i(t+h) = \mu_i(t) p_{ii}(h) + \sum_{\substack{j \in E \\ j \neq i}} \mu_j(t) p_{ji}(h),$$

that is, after rearrangement,

$$\frac{\mu_i(t+h) - \mu_i(t)}{h} = -\mu_i(t) \frac{1 - p_{ii}(h)}{h} + \sum_{\substack{j \in E \\ j \neq i}} \mu_j(t) \frac{p_{ji}(h)}{h}.$$

The rest of the proof is then the same as that of Theorem 3.2. □

Remark 3.1. Condition (3.10) is trivially satisfied when the state space is finite, or when

$$\sup_{i \in E} q_i < \infty. \quad (3.14)$$

If this condition is satisfied, the chain has the same distribution as a uniform chain, as we shall prove later, in Corollary 4.1. ◇

Definition 3.1. *Stationary Distribution*

A *stationary distribution* of the semigroup $\{\mathbf{P}(t)\}_{t \geq 0}$ is any probability π on E such that for all $t \geq 0$,

$$\pi^T \mathbf{P}(t) = \pi^T. \quad (3.15)$$

If $X(0)$ is distributed according to π , then so is $X(t)$ for all $t > 0$, because from (3.11),

$$\mu^T(t) = \pi^T \mathbf{P}(t) = \pi^T.$$

The chain is then said to be *in a stationary regime* or *in equilibrium* because for all times t_1, \dots, t_k, t such that $t \geq 0$ and $0 \leq t_1 \leq t_2 \leq \dots \leq t_k$, and all states $i_1, \dots, i_k \in E$,

$$P(X(t_1 + t) = i_1, \dots, X(t_k + t) = i_k) = P(X(t_1) = i_1, \dots, X(t_k) = i_k).$$

This is proved by computing the left-hand side using Bayes's sequential rule, to obtain

$$\pi(i_1)p_{i_1 i_2}(t_2 - t_1) \cdots p_{i_{k-1} i_k}(t_k - t_{k-1}),$$

and by observing that this quantity does not depend upon t .

In general, as for discrete-time HMCs, a stationary distribution need not exist, and if it exists, need not be unique. Continuous time brings something new into the picture, namely a local characterization of the stationary distributions, that is often more practical than the global characterization (3.15).

If π is a stationary distribution of a stable and conservative continuous-time HMC, and if $\sum_{i \in E} \pi(i)q_i < \infty$, then, according to Theorem 3.3, π satisfies the global balance equation

$$\pi^T \mathbf{A} = 0, \quad (3.16)$$

that is, in expanded form,

$$\pi(i)q_i = \sum_{\substack{j \in E \\ j \neq i}} \pi(j)q_{ji}.$$

The above result is not stated as a theorem because it is too weak. First we need to assume stability and conservation, then we must introduce an assumption on the stationary distribution itself. Also, we lack the principal result, namely the converse: If a probability distribution π satisfies (3.16), then it is a stationary distribution. However, there are interesting cases where a satisfying result can be obtained at low cost.

Example 3.2. Finite State Space

If the state space is finite, we know that the system is stable and conservative, and that both the forward and backward Kolmogorov systems are satisfied, as well as (3.13). Also, if there exists a stationary distribution π , (3.12) is satisfied, and therefore, necessarily $\pi^T \mathbf{A} = 0$.

Now suppose that $\pi^T \mathbf{A} = 0$ for some probability distribution π on E . Then $\mu(t) = \pi$ is a solution of the system (3.13). But (3.13) is a finite linear system of differential equations and it has a unique solution $\mu(t)$ such that $\mu(0) = \pi$. Therefore, if $X(0)$ is distributed according to π , then so is $X(t)$ for all $t \geq 0$. Thus π is a stationary distribution.

In summary: For a continuous-time HMC on the *finite* state space E with infinitesimal generator \mathbf{A} , the condition $\pi^T \mathbf{A} = 0$ is necessary and sufficient for the probability distribution π on E to be a stationary distribution. \diamond

Example 3.3. Stationary Uniform HMC

For the uniform Markov chain of Definition 2.2 and Example 2.3, $q_i = \lambda(1 - k_{ii})$, and therefore $\sup q_i < \infty$. Therefore, the conditions of Theorems 3.1, 3.2, and 3.3 are satisfied.

Therefore, a stationary distribution π , if it exists, is a solution of $\pi^T \mathbf{A} = 0$. In view of (2.19), $\pi^T \mathbf{A} = 0$ is equivalent to

$$\pi^T = \pi^T \mathbf{K}. \quad (3.17)$$

(This shows, in particular, that π is a stationary distribution of the subordinated chain.)

Now, suppose that $\pi^T \mathbf{A} = 0$, and therefore (3.17) holds. From (2.9),

$$\pi^T \mathbf{P}(t) = \sum_{n=0}^{\infty} e^{-\lambda t} \frac{(\lambda t)^n}{n!} \pi^T \mathbf{K}^n = \left(\sum_{n=0}^{\infty} e^{-\lambda t} \frac{(\lambda t)^n}{n!} \right) \pi^T = \pi^T.$$

In other words, π is a stationary distribution of the uniform HMC $\{X(t)\}_{t \geq 0}$. Therefore, in this case also, $\pi^T \mathbf{A} = 0$ is a necessary and sufficient condition for the probability distribution π to be a stationary distribution. \diamond

3.3 Regular Jumps

In the practice of operations research and of engineering, and in the biological and social sciences also, the Markov chains that are likely to be encountered are regular jump processes. For such chains, the situation is just as simple as for finite state space or uniform Markov chains.

Theorem 3.4.

Let $\{X(t)\}_{t \geq 0}$ be a regular jump HMC with countable state space E and transition semigroup $\{\mathbf{P}(t)\}_{t \geq 0}$. Then

(i) The semigroup $\{\mathbf{P}(t)\}$ is continuous, stable, and conservative.

(ii) The backward and forward differential systems of Kolmogorov are satisfied, as well as the global balance differential systems.

(iii) A necessary and sufficient condition for a probability distribution π on E to be a stationary distribution of the chain is $\pi^T \mathbf{A} = 0$, where \mathbf{A} is the infinitesimal generator.

The proof is a little bit technical and is omitted; see Chapter 1 of (Anderson, 1991). In Section 1 of Chapter 9, we shall obtain a very close result, which is sufficient for practical purposes.

Example 3.4. Global Balance for Birth-and-Death Processes

Recall Definition 2.6, where we defined a birth-and-death process with parameters λ_n and μ_n to be a regular jump HMC taking its values in \mathbb{N} and with an infinitesimal generator of the form

$$\mathbf{A} = \begin{pmatrix} -\lambda_0 & \lambda_0 & 0 & 0 & \cdots & \cdots \\ \mu_1 & -(\lambda_1 + \mu_1) & \lambda_1 & 0 & 0 & \cdots \\ 0 & \mu_2 & -(\lambda_2 + \mu_2) & \lambda_2 & 0 & \cdots \\ 0 & 0 & \mu_3 & -(\lambda_3 + \mu_3) & \lambda_3 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \quad (3.18)$$

The differential system satisfied by $p_n(t) = P(X(t) = n)$ is

$$\begin{aligned}\dot{p}_0(t) &= -\lambda_0 p_0(t) + \mu_1 p_1(t), \\ \dot{p}_n(t) &= \lambda_{n-1} p_{n-1}(t) - (\lambda_n + \mu_n) p_n(t) + \mu_{n+1} p_{n+1}(t) \quad (\text{for } n \geq 1).\end{aligned}$$

A necessary and sufficient condition for a probability π to be a stationary distribution of the birth-and-death process is

$$\begin{aligned}0 &= \lambda_0 \pi(0) - \mu_1 \pi(1), \\ 0 &= \lambda_{n-1} \pi(n-1) - (\lambda_n + \mu_n) \pi(n) + \mu_{n+1} \pi(n+1) \quad (\text{for } n \geq 1).\end{aligned}$$

For fixed $\pi(0)$, there exists one and only one solution for this system, namely,

$$\pi(n) = \pi(0) \frac{\lambda_0 \lambda_1 \cdots \lambda_{n-1}}{\mu_1 \mu_2 \cdots \mu_n} \quad (3.19)$$

for $n \geq 1$. If we require π to be a probability distribution, we must have

$$\pi(0) \left(1 + \sum_{n=1}^{\infty} \frac{\lambda_0 \lambda_1 \cdots \lambda_{n-1}}{\mu_1 \mu_2 \cdots \mu_n} \right) = 1, \quad (3.20)$$

and this is possible if and only if

$$1 + \sum_{n=1}^{\infty} \frac{\lambda_0 \lambda_1 \cdots \lambda_{n-1}}{\mu_1 \mu_2 \cdots \mu_n} < \infty, \quad (3.21)$$

in which case the unique stationary probability is given by (3.19) and (3.20). \diamond

4 The Regenerative Structure

4.1 Strong Markov Property

A regular jump HMC has the strong Markov property. For a precise statement of this result, we need to introduce stopping times in continuous time. The definition is completely analogous to the discrete-time definition.

Definition 4.1. Stopping Times

Let $\{X(t)\}_{t \geq 0}$ be a stochastic process with values in E . A random variable τ taking its values in \mathbb{R}_+ is called a *stopping time* with respect to $\{X(t)\}$ if for all $t \geq 0$, the event $\{\tau \leq t\}$ is expressible in terms of $\{X(s), s \in [0, t]\}$. We denote this by

$$\{\tau \leq t\} \in X_0^t. \quad (4.1)$$

Example 4.1. *Escape and Return Times*

Let $\{X(t)\}_{t \geq 0}$ be a regular jump HMC with countable state space E . The *escape time* from state i , denoted by E_i is defined by

$$E_i = \inf \{t \geq 0; X(t) \neq i\}, \tag{4.2}$$

with $E_i = \infty$ if $X(t) = i$ for all $t \geq 0$. The *return time* to i , denoted by R_i , is defined by

$$R_i = \inf \{t > 0; t > E_i \text{ and } X(t) = i\}, \tag{4.3}$$

with $R_i = \infty$ if $E_i = \infty$ or $X(t) \neq i$ for all $t \geq E_i$. Then E_i and R_i are stopping times with respect to $\{X(t)\}_{t \geq 0}$ (exercise). \diamond

We define the process after the random time τ just as we did in discrete time:

$$\{X(t + \tau)\}_{t \geq 0}, \tag{4.4}$$

with the convention $X(\infty) = \Delta$, where Δ is an arbitrary element not in E . The process before τ is

$$\{X(t \wedge \tau)\}_{t \geq 0}. \tag{4.5}$$

Theorem 4.1. *Strong Markov Property*

Let $\{X(t)\}_{t \geq 0}$ be a regular jump HMC with countable state space $E = \mathbb{N}$ (without loss of generality) and transition semigroup $\{\mathbf{P}(t)\}_{t \geq 0}$, and let τ be a stopping time with respect to $\{X(t)\}_{t \geq 0}$. Let $k \in E$ be an arbitrary state. Then, given that $X(\tau) = k$,

(α) the chain after τ and the chain before τ are independent, and

(β) the chain after τ is a regular jump HMC with transition semigroup $\{\mathbf{P}(t)\}_{t \geq 0}$.

Proof. It suffices to show that for all states k , all positive times $t_1, \dots, t_n, s_1, \dots, s_p$, all real numbers $u_1, \dots, u_n, v_1, \dots, v_p$, and all initial distributions μ

$$\begin{aligned} E_\mu \left[\exp \left\{ i \sum_{\ell=1}^n u_\ell X(\tau + t_\ell) + i \sum_{m=1}^p v_m X(\tau \wedge s_m) \right\} \mathbf{1}_{\{X(\tau)=k\}} \right] \\ = E_k \left[\exp \left\{ i \sum_{\ell=1}^n u_\ell X(t_\ell) \right\} \right] E_\mu \left[\exp \left\{ i \sum_{m=1}^p v_m X(\tau \wedge s_m) \right\} \mathbf{1}_{\{X(\tau)=k\}} \right] \end{aligned} \tag{4.6}$$

Indeed, if (4.6) is satisfied, fixing $v_1 = \dots = v_p = 0$, we obtain

$$E_\mu \left[\exp \left\{ i \sum_{\ell=1}^n u_\ell X(\tau + t_\ell) \right\} \mathbf{1}_{\{X(\tau)=k\}} \right] = E_k \left[\exp \left\{ i \sum_{\ell=1}^n u_\ell X(t_\ell) \right\} \right] P_\mu(X(\tau) = k),$$

that is,

$$\begin{aligned} E_\mu \left[\exp \left\{ i \sum_{\ell=1}^n u_\ell X(\tau + t_\ell) \right\} \mid X(\tau) = k \right] &= E_k \left[\exp \left\{ i \sum_{\ell=1}^n u_\ell X(t_\ell) \right\} \right] \\ &= E_\mu \left[\exp \left\{ i \sum_{\ell=1}^n u_\ell X(t_\ell) \right\} \mid X(0) = k \right], \end{aligned}$$

and this shows that given $X(\tau) = k$, $\{X(\tau + t)\}_{t \geq 0}$ had the same distribution as $\{X(t)\}_{t \geq 0}$ given $X(0) = k$. We therefore will have proven (β) . For (α) , it suffices to rewrite (4.6) as follows, using the previous equality:

$$\begin{aligned} & E_\mu \left[\exp \left\{ i \sum_{\ell=1}^n u_\ell X(\tau + t_\ell) + i \sum_{m=1}^n v_m X(\tau \wedge s_m) \right\} \mid X(\tau) = k \right] \\ &= E_\mu \left[\exp \left\{ i \sum_{\ell=1}^n u_\ell X(\tau + t_\ell) \right\} \mid X(\tau) = k \right] \\ &\quad \times E_\mu \left[\exp \left\{ i \sum_{m=1}^n v_m X(\tau \wedge s_m) \right\} \mid X(\tau) = k \right]. \end{aligned}$$

It remains to prove (4.6). For the sake of simplicity, we do the case where $n = m = 1$, and let $u_1 = u, t_1 = t, v_1 = v, s_1 = s$.

We first treat the case where τ takes a countable number of values, denoted by a_j . Then

$$\begin{aligned} & E_\mu[\exp\{iuX(\tau + t) + ivX(\tau \wedge s)\} 1_{\{X(\tau)=k\}}] \\ &= \sum_{j=1}^{\infty} E_\mu[\exp\{iuX(a_j + t) + ivX(a_j \wedge s)\} 1_{\{X(a_j)=k\}} 1_{\{\tau=a_j\}}]. \end{aligned}$$

For all $j \geq 1$,

$$\begin{aligned} & E_\mu[\exp\{iuX(a_j + t) + ivX(a_j \wedge s)\} 1_{\{X(a_j)=k\}} 1_{\{\tau=a_j\}}] \\ &= E_\mu[\exp\{iuX(a_j + t)\} 1_{\{X(a_j)=k\}} \exp\{ivX(a_j \wedge s)\} 1_{\{\tau=a_j\}}] \\ &= E_\mu[\exp\{iuX(a_j + t)\} \mid X(a_j) = k] E_\mu[\exp\{ivX(a_j \wedge s)\} 1_{\{\tau=a_j\}} 1_{\{X(a_j)=k\}}], \end{aligned}$$

where for the last equality, we have used the fact that $1_{\{\tau=a_j\}}$ is a function $X(u), 0 \leq u \leq a_j$, and the Markov property at time a_j .

Therefore, for all $j \geq 1$,

$$\begin{aligned} & E_\mu[\exp\{iuX(a_j + t) + ivX(a_j \wedge s)\} 1_{\{X(a_j)=k\}} 1_{\{\tau=a_j\}}] \\ &= E_k[\exp\{iuX(t)\}] E_\mu[\exp\{ivX(a_j \wedge s)\} 1_{\{X(a_j)=k\}} 1_{\{\tau=a_j\}}]. \end{aligned}$$

Summing with respect to j , we obtain the equality corresponding to (4.6). To pass from the case where the stopping time τ takes a countable number of values to the general case, define for each $n \geq 1$ the following approximation $\tau(n)$ to the arbitrary stopping time τ (with respect to $\{X(t)\}_{t \geq 0}$):

$$\tau(n, \omega) = \begin{cases} 0 & \text{if } \tau(\omega) = 0, \\ \frac{k+1}{2^n} & \text{if } \frac{k}{2^n} < \tau(\omega) \leq \frac{k+1}{2^n}, \\ +\infty & \text{if } \tau(\omega) = \infty. \end{cases}$$

Then $\tau(n)$ is a stopping time (exercise) with a countable number of values, and therefore, (4.6) is satisfied for $\tau(n)$. Now, $\lim_{n \uparrow \infty} \tau(n, \omega) = \tau(\omega)$, and therefore

$$\lim_{n \uparrow \infty} X(\tau(n) \wedge a) = X(\tau \wedge a), \lim_{n \uparrow \infty} X(\tau(n) + b) = X(\tau + b), \lim_{n \uparrow \infty} 1_{\{X(\tau(n))=k\}} = 1_{\{X(\tau)=k\}}$$

(use the fact that a regular jump process is right-continuous). Therefore, letting n go to ∞ in (4.6) with τ replaced by $\tau(n)$, we obtain the result for τ itself, by dominated convergence (see Theorem 3.2 of the Appendix). \square

4.2 Embedded Chain and Transition Times

Let $\{\tau_n\}_{n \geq 0}$ be the nondecreasing sequence of transition times of the regular jump process $\{X(t)\}_{t \geq 0}$, where $\tau_0 = 0$ and $\tau_n = \infty$ if there are strictly fewer than n transitions in $(0, \infty)$.

For each $n \geq 0$, τ_n is a stopping time with respect to $\{X(t)\}_{t \geq 0}$ (exercise).

The process $\{X_n\}_{n \geq 0}$ with values in $E_\Delta = E \cup \{\Delta\}$, where Δ is an arbitrary element not in E , is defined by

$$X_n = X(\tau_n), \quad (4.7)$$

with the convention $X(\infty) = \Delta$, and it is called the *embedded process* of the jump process. If $\{X(t)\}_{t \geq 0}$ is a regular jump HMC, it follows from the strong Markov property that given $X(\tau_n) = k \in E$, $\{X(\tau_n + t)\}_{t \geq 0}$ is independent of $\{X(\tau_n \wedge t)\}_{t \geq 0}$, and therefore, given $X_n = k$ the variables $(X_{n+1}, X_{n+2}, \dots)$ are independent of (X_0, \dots, X_n) , that is, $\{X_n\}_{n \geq 0}$ is a Markov chain. It is clearly homogeneous because the distribution of $\{X(\tau_n + t)\}_{t \geq 0}$ given $X(\tau_n) = k$ is independent of n , being identical with the distribution of $\{X(t)\}_{t \geq 0}$ given $X(0) = k$.

More precisely:

Theorem 4.2. Regenerative Structure

Let $\{X(t)\}_{t \geq 0}$ be a regular jump HMC, with infinitesimal generator \mathbf{A} , transition times sequence $\{\tau_n\}_{n \geq 0}$, and embedded process $\{X_n\}_{n \geq 0}$. Then

(α) $\{X_n\}_{n \geq 0}$ is a discrete-time HMC with state space $E_\Delta = E \cup \{\Delta\}$ with transition matrix given by $p_{\Delta\Delta} = 1$, $p_{i\Delta} = 1$ if $i \in E$ and $q_i = 0$; by $p_{i\Delta} = 0$ if $i \in E$ and $q_i > 0$; and if $q_i > 0$ and $j \neq i$, by

$$p_{ij} = \frac{q_{ij}}{q_i}. \quad (4.8)$$

(β) Given $\{X_n\}_{n \geq 0}$, the sequence $\{\tau_{n+1} - \tau_n\}_{n \geq 0}$ is independent, and for all $n \geq 0$ and all $a \in \mathbb{R}_+$,

$$P(\tau_{n+1} - \tau_n \leq a \mid \{X_k\}_{k \geq 0}) = 1 - e^{-q_{X_n} a}. \quad (4.9)$$

Proof. The complete proof of Theorem 4.2 will be given in Chapter 9, Section 1. We shall for the time being begin with the following partial result.

(α') $\{X_n\}_{n \geq 0}$ is a discrete-time HMC on $E_\Delta = E \cup \{\Delta\}$.

(β') Given $\{X_n\}_{n \geq 0}$, the sequence $\{\tau_{n+1} - \tau_n\}_{n \geq 0}$ is independent.

(β'') There exists for each $i \in E$ a real number $\lambda(i) \geq 0$ such that for all $n \geq 0$ and all $a \in \mathbb{R}_+$,

$$P(\tau_{n+1} - \tau_n \leq a \mid \{X_k\}_{k \geq 0}) = 1 - e^{-\lambda(X_n) a}.$$

We have already proven (α') . Call $p_{ij} = P_i(X(\tau_1) = j)$ the transition probability of $\{X_n\}_{n \geq 1}$ from i to j . To prove (β') and (β'') , it suffices to show that

$$\begin{aligned} P_i(X_1 = i, \dots, X_n = i_n, \tau_1 - \tau_0 > a_1, \dots, \tau_n - \tau_{n-1} > a_n) \\ = e^{-\lambda(i)a_1} p_{ii_1} e^{-\lambda(i_1)a_2} p_{i_1 i_2} \dots e^{-\lambda(i_{n-1})a_n} p_{i_{n-1} i_n} \end{aligned}$$

for all $i, i_1, \dots, i_n \in E$, $a_1, \dots, a_n \in \mathbb{R}_+$, and for some function $\lambda : E \rightarrow \mathbb{R}_+$. By the strong Markov property, it suffices, in fact, to show that for all $i, j \in E$, $a \in \mathbb{R}_+$, there exists $\lambda(i) \geq 0$ such that

$$P_i(X_1 = j, \tau_1 - \tau_0 > a) = P_i(X_1 = j) e^{-\lambda(i)a}. \quad (4.10)$$

Define $g(t) = P_i(\tau_1 > t)$. For $t, s \geq 0$, using the obvious set identities,

$$\begin{aligned} g(t+s) &= P_i(\tau_1 > t+s) \\ &= P_i(\tau_1 > t+s, \tau_1 > t, X(t) = i) \\ &= P_i(X(t+u) = i(u \in [0, s]), \tau_1 > t, X(t) = i). \end{aligned}$$

The last expression is, in view of the Markov property at time t and using the fact that $\{\tau_1 > t\}$ is expressible in terms of $\{X(v), v \in [0, t]\}$,

$$\begin{aligned} P_i(X(t+u) = i(u \in [0, s]) \mid X(t) = i) P_i(\tau_1 > t, X(t) = i) \\ = P_i(X(u) = i(u \in [0, s]) \mid X(0) = i) P_i(\tau_1 > t) \\ = P_i(\tau_1 > s) P_i(\tau_1 > t) \end{aligned}$$

where the last two equalities again follow from the obvious set identities. Therefore, for all $s, t \geq 0$,

$$g(t+s) = g(t)g(s).$$

Also, $g(t)$ is nonincreasing, and $\lim_{t \downarrow 0} g(t) = 1$ (the semi group is continuous). It follows (see Theorem 1.11 of the Appendix) that there exists $\lambda(i) \in [0, \infty)$ such that $g(t) = e^{-\lambda(i)t}$, i.e., $P_i(\tau_1 > t) = e^{-\lambda(i)t}$, for all $t \geq 0$.

Now, using the Markov property and appropriate set identities,

$$\begin{aligned} P_i(X_1 = j, \tau_1 > t) &= P_i(X(\tau_1) = j, \tau_1 > t, X(t) = i) \\ &= P_i(\text{first jump of } \{X(t+s)\}_{s \geq 0} \text{ is } j, \tau_1 > t, X(t) = i) \\ &= P_i(\text{first jump of } \{X(t+s)\}_{s \geq 0} \text{ is } j \mid X(t) = i) P_i(\tau_1 > t, X(t) = i) \\ &= P_i(\text{first jump of } \{X(s)\}_{s \geq 0} \text{ is } j \mid X(0) = i) P_i(\tau_1 > t) \\ &= P_i(X(\tau_1) = j) P_i(\tau_1 > t), \end{aligned}$$

and this is (4.10). □

Definition 4.2. *Essential, Permanent*

A state $i \in E$ such that $q_i = 0$ is called *permanent*; otherwise, it is called *essential*.

In view of (4.9), if $X(\tau_n) = i$, a permanent state, then $\tau_{n+1} - \tau_n = \infty$; that is, there is no more transition at finite distance, hence the terminology.

Example 4.2. *Regenerative Structure of Uniform MC*

For the uniform MC of Definition 2.2 and Example 2.3, the embedded process $\{X_n\}_{n \geq 0}$ is an HMC with state space $E_\Delta = E \cup \{\Delta\}$, and if $i \in E$ is not permanent (i.e., in this case if $k_{ii} < 1$), then for $j \neq i$,

$$p_{ij} = \frac{k_{ij}}{1 - k_{ii}}.$$

Indeed, $\{X_n\}$ is obtained from $\{\hat{X}_n\}$ by considering only the “real” transitions (see Problem 2.7.1). ◊

An immediate corollary of Theorem 4.2 is:

Theorem 4.3. *The Infinitesimal Generator Characterizes the Semigroup*

Two regular jump HMCs with the same infinitesimal generator and the same initial distribution are probabilistically equivalent.

Another way to state this is as follows: Two regular jump HMCs with the same infinitesimal generator have the same transition semigroup.

Corollary 4.1.

A regular jump HMC with infinitesimal generator \mathbf{A} such that $\sup_{i \in E} q_i \leq \infty$ has the same transition semigroup as a uniform chain.

Proof. Select any real number $\lambda > \sup_{i \in E} q_i$, and define the transition matrix \mathbf{K} by (2.20) and (2.21). The uniform chain corresponding to (λ, \mathbf{K}) has the infinitesimal generator \mathbf{A} . □

We shall say that a continuous time HMC with an infinitesimal generator such that

$$\sup_{i \in E} q_i < \infty, \tag{4.11}$$

is *uniformizable*. Any pair (λ, \mathbf{K}) as in the proof of corollary 4.1 gives rise to a uniform version of the chain. The *minimal* uniform version is that with $\lambda = \sup_{i \in E} q_i$.

4.3 Explosions

Definition 4.3. *Generators*

Let $\mathbf{A} = \{q_{ij}\}_{i,j \in E}$ be a matrix with entries satisfying for all $i, j \in E$,

$$q_i \in [0, \infty), \quad q_{ij} \in [0, \infty), \quad \sum_{\substack{k \in E \\ k \neq i}} q_{ik} = q_i, \tag{4.12}$$

where $q_i = -q_{ii}$. This matrix is called a stable and conservative *generator* on E ; and it is called an *essential generator* if, moreover, $q_i > 0$ for all $i \in E$.

Note that no reference is made to a continuous-time HMC.

Theorem 4.2 suggests a way of constructing a regular jump HMC $\{X(t)\}_{t \geq 0}$ with values in a countable state space E and admitting a *given* generator $\mathbf{A} = \{q_{ij}\}$ that is stable and conservative. We shall also suppose (for simplicity) that it is *essential*. One would construct a sequence $\tau_0 = 0, X_0, \tau_1 - \tau_0, X_1, \tau_2 - \tau_1, X_2, \dots$ according to

$$\begin{cases} P(\tau_{n+1} - \tau_n \leq x \mid X_0, \dots, X_n, \tau_0, \dots, \tau_n) = 1 - e^{-q_{X_n} x} \\ P(X_{n+1} = j \mid X_0, \dots, X_n, \tau_0, \dots, \tau_{n+1}) = q_{X_n j} / q_{X_n}, \end{cases} \quad (4.13)$$

the initial state X_0 being chosen at random, with arbitrary distribution. The value of $X(t)$ for $\tau_n \leq t < \tau_{n+1}$ is then X_n .

The problem with the above regenerative construction is that

$$\tau_\infty \stackrel{\text{def}}{=} \lim_{n \uparrow \infty} \tau_n \quad (4.14)$$

may not be almost surely finite. We say that the generator \mathbf{A} is *non explosive* if

$$P_\mu(\tau_\infty = \infty) = 1, \quad (4.15)$$

for any initial distribution μ and all regular jump HMC associated to a semigroup with generator \mathbf{A} . It will be proven in Chapter 9 that in the nonexplosive case, the above construction indeed produces a regular jump HMC.

Theorem 4.4. Reuter's Criterion

Let \mathbf{A} be a stable and conservative generator on E . It is nonexplosive if and only if for any real $\lambda > 0$, the system of equations

$$(\lambda + q_i)x_i = \sum_{\substack{j \in E \\ j \neq i}} q_{ij}x_j, \quad i \in E, \quad (4.16)$$

admits no nonnegative bounded solution other than the trivial one.

Proof. The number

$$g_i(\lambda) = E_i \left[\exp \left\{ -\lambda \sum_{k=1}^{\infty} S_k \right\} \right]$$

is uniformly bounded in $\lambda > 0$ and $i \in E$, and if $P_i(\tau_\infty = \infty) < 1$, it is strictly positive. Also, $x_i = g_i(\lambda), i \in E$, is a solution of (4.16), as follows from the calculations below:

$$\begin{aligned} g_i(\lambda) &= E_i \left[\exp\{-\lambda S_1\} \exp \left\{ -\lambda \sum_{k=2}^{\infty} S_k \right\} \right] \\ &= \left(\int_0^\infty e^{-\lambda t} q_i e^{-q_i t} dt \right) E_i \left[\exp \left\{ -\lambda \sum_{k=2}^{\infty} S_k \right\} \right] \\ &= \frac{q_i}{\lambda + q_i} E_i \left[\exp \left\{ -\lambda \sum_{k=2}^{\infty} S_k \right\} \right] \end{aligned}$$

and, by first-step analysis,

$$E_i \left[\exp \left\{ -\lambda \sum_{k=2}^{\infty} S_k \right\} \right] = \sum_{\substack{j \in E \\ j \neq i}} E_j \left[\exp \left\{ -\lambda \sum_{k=2}^{\infty} S_k \right\} \right] \frac{q_{ij}}{q_i} = \sum_{\substack{j \in E \\ j \neq i}} g_j(\lambda) \frac{q_{ij}}{q_i}.$$

Therefore, if \mathbf{A} is explosive there exists a non trivial bounded solution of (4.16). We now prove the converse. Call $\{g_i(\lambda)\}_{i \in E}$ a bounded solution of (4.16) for a fixed real $\lambda > 0$. We have

$$g_i(\lambda) = E[\exp\{-\lambda S_1\} g_{X_1}(\lambda) \mid X_0 = i], \quad (4.17)$$

since the right-hand side is equal to that of (4.16), as first-step analysis shows. We prove by induction that

$$g_i(\lambda) = E \left[\exp \left\{ -\lambda \sum_{k=1}^n S_k \right\} g_{X_n}(\lambda) \mid X_0 = i \right]. \quad (4.18)$$

For this, rewrite (4.17) as

$$g_i(\lambda) = E[\exp\{-\lambda S_{n+1}\} g_{X_{n+1}}(\lambda) \mid X_n = i],$$

that is,

$$g_{X_n}(\lambda) = E[\exp\{-\lambda S_{n+1}\} g_{X_{n+1}}(\lambda) \mid X_n].$$

If this expression of $g_{X_n}(\lambda)$ is used in (4.18), then it follows that

$$g_i(\lambda) = E \left[\exp \left\{ -\lambda \sum_{k=1}^{n+1} S_k \right\} g_{X_{n+1}}(\lambda) \mid X_0 = i \right]. \quad (4.19)$$

Therefore, (4.18) implies (4.19) (the forward step in the induction argument). Since (4.18) is true for $n = 1$ (Eqn. (4.17)), it is true for all $n \geq 1$, and therefore, since $K \stackrel{\text{def}}{=} |g_i(\lambda)| < \infty$,

$$|g_i(\lambda)| \leq K E_i \left[\exp \left\{ -\lambda \sum_{k=1}^{\infty} S_k \right\} \right].$$

Therefore, if $\{g_i(\lambda)\}_{i \in E}$ is not trivial, for some $i \in E$ it must hold that $P_i(\sum_{k=1}^{\infty} S_k < \infty) > 0$, or equivalently, $P_i(\tau_{\infty} = \infty) < 1$. \square

Problems 8.4.1 to 8.4.3 give sufficient conditions of nonexplosion that are satisfied in many cases. We shall see in the next chapter that for continuous-time HMCs arising in a queuing context, nonexplosion is easy to prove. This seems to diminish the impact of Reuter's criterion. However, there exists a very important situation where it is irreplaceable:

Theorem 4.5. Reuter's Criterion for Birth-and-Death Generators

Let \mathbf{A} be generator on $E = \mathbb{N}$ defined by $q_{n,n+1} = \lambda_n$ and $q_{n,n-1} = \mu_n 1_{n \geq 1}$, where the birth parameters λ_n are strictly positive. A necessary and sufficient condition of nonexplosion of this generator is

$$\sum_{n=1}^{\infty} \left[\frac{1}{\lambda_n} + \frac{\mu_n}{\lambda_n \lambda_{n-1}} + \cdots + \frac{\mu_n \cdots \mu_1}{\lambda_n \cdots \lambda_1 \lambda_0} \right] = \infty. \quad (4.20)$$

Proof. We start with preliminary remarks concerning the system of equations (4.16) in the particular case of a birth-and-death generator

$$\begin{cases} \lambda x_0 = -\lambda_0 x_0 + \lambda_0 x_1, \\ \lambda x_k = \mu_k x_{k-1} - (\lambda_k + \mu_k) x_k + \lambda_k x_{k+1} \quad (k \geq 1). \end{cases} \quad (4.21)$$

For any fixed x_0 , this system admits a unique solution, that is identically null if and only if $x_0 = 0$. If $x_0 \neq 0$, the solution is such that x_k/x_0 does not depend on x_0 , and therefore, only the case where $x_0 = 1$ needs to be treated.

Writing $y_k = x_{k+1} - x_k$, we obtain from (4.21)

$$y_k = \frac{\lambda}{\lambda_k} x_k + \frac{\mu_k}{\lambda_k} \frac{\lambda}{\lambda_{k-1}} x_{k-1} + \cdots + \frac{\mu_k \cdots \mu_2}{\lambda_k \cdots \lambda_2} \frac{\lambda}{\lambda_1} x_1 + \frac{\mu_k \cdots \mu_1}{\lambda_k \cdots \lambda_1} y_0 \quad (4.22)$$

and $y_0 = \frac{\lambda}{\lambda_0}$. From this we deduce that if $\lambda > 0$, then $y_k > 0$ and therefore $\{x_k\}_{k \geq 0}$ is a strictly increasing sequence.

Therefore, using $y_0 = \frac{\lambda}{\lambda_0}$ in (4.22), we have

$$y_k \geq \lambda \left[\frac{1}{\lambda_k} + \frac{\mu_k}{\lambda_k \lambda_{k-1}} + \cdots + \frac{\mu_k \cdots \mu_1}{\lambda_k \cdots \lambda_1 \lambda_0} \right].$$

Thus, we see that a necessary condition for $\{x_k\}_{k \geq 0}$ to be bounded is that the left-hand-side of (4.20) be finite. This proves the sufficiency of (4.20) for nonexplosion.

We now turn to the proof of necessity. For $i \leq k$, bounding in (4.22) x_i by x_k yields the majoration

$$y_k \leq \left[\frac{\lambda}{\lambda_k} + \cdots + \frac{\mu_k \cdots \mu_1 \lambda}{\lambda_k \cdots \lambda_1 \lambda_0} \right] x_k,$$

and therefore, since $y_k = x_{k+1} - x_k$,

$$\begin{aligned} x_{k+1} &\leq \left[1 + \frac{\lambda}{\lambda_k} + \cdots + \frac{\mu_k \cdots \mu_1 \lambda}{\lambda_k \cdots \lambda_1 \lambda_0} \right] x_k \\ &\leq x_k \exp \left\{ \lambda \left[\frac{1}{\lambda_k} + \cdots + \frac{\mu_k \cdots \mu_1}{\lambda_k \cdots \lambda_0} \right] \right\}. \end{aligned}$$

Since $x_0 = 1$, this leads to

$$x_n \leq \exp \left\{ \lambda \sum_{k=1}^n \left[\frac{1}{\lambda_k} + \cdots + \frac{\mu_k \cdots \mu_1}{\lambda_k \cdots \lambda_0} \right] \right\}.$$

Therefore, a sufficient condition for the solution $\{x_n\}_{n \geq 0}$ to be bounded is that the left-hand side of (4.20) be finite. This proves necessity in Theorem 4.5. \square

Example 4.3. Pure Birth

A pure birth generator \mathbf{A} is a birth-and-death generator with all $\mu_n = 0$. The necessary and sufficient condition of nonexplosion (4.20) reads in this case

$$\sum_{n=0}^{\infty} \frac{1}{\lambda_n} = \infty. \quad (4.23)$$

◇

Example 4.4. Absorbing Birth-and-Death Processes.

Let \mathbf{A} be a birth-and-death generator with parameters λ_n and μ_n , where $\lambda_n > 0$ and $\mu_n > 0$ for all $n \geq 1$, but

$$\lambda_0 = 0. \quad (4.24)$$

Starting from the initial state $i \geq 1$, we seek to obtain the probability of absorption u_i in state 0 (indeed condition (4.24) makes 0 an absorbing state).

Clearly, this probability is the same as the probability of absorption by 0 of the embedded chain, with transition matrix

$$\mathbf{P} = \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots \\ \frac{\mu_1}{\lambda_1 + \mu_1} & 0 & \frac{\lambda_1}{\lambda_1 + \mu_1} & 0 & \cdots \\ 0 & \frac{\mu_2}{\lambda_2 + \mu_2} & 0 & \frac{\lambda_2}{\lambda_2 + \mu_2} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix},$$

By first-step analysis, for $i \geq 1$

$$u_i = \frac{\mu_i}{\lambda_i + \mu_i} u_{i-1} + \frac{\lambda_i}{\lambda_i + \mu_i} u_{i+1} \quad (4.25)$$

with the boundary condition $u_0 = 1$. Equation (4.25) is, for $i \geq 1$,

$$u_{i+1} - u_i = \frac{\mu_i}{\lambda_i} (u_i - u_{i-1}).$$

It follows from this and the boundary condition that for $n > 1$,

$$u_n - u_1 = -(1 - u_1) \sum_{i=1}^{n-1} \gamma_i, \quad (4.26)$$

where for $i \geq 1$,

$$\gamma_0 = 1, \quad \gamma_i = \frac{\mu_1 \mu_2 \cdots \mu_i}{\lambda_1 \lambda_2 \cdots \lambda_i}. \quad (4.27)$$

Since u_n is bounded by 1 for all $i \geq 0$ (it is a probability), we see from (4.26) that if

$$\sum_{i=1}^{\infty} \gamma_i = \infty, \quad (4.28)$$

then necessarily $u_1 = 1$, and consequently, $u_n = u_1 = 1$ for all $n \geq 0$. This means that under condition (4.28), the process is ultimately absorbed by state 0, whatever the initial state. Suppose now that $u_1 \in (0, 1)$, and consequently,

$$\sum_{i=1}^{\infty} \gamma_i < \infty. \quad (4.29)$$

By (4.26), u_n is decreasing to some $u \geq 0$ as $n \rightarrow \infty$. Necessarily, $u = 0$, since otherwise $u_n \geq u > 0$ for all n , and a Borel-Cantelli argument (see Problem 1.8.2) then shows that $u_n = 1$ for all n , a contradiction with $u_1 \in (0, 1)$.

Equation (4.26) gives at the limit $n \rightarrow \infty$,

$$u_1 = \frac{\sum_{i=1}^{\infty} \gamma_i}{1 + \sum_{i=1}^{\infty} \gamma_i},$$

and therefore, from (4.26) again,

$$u_n = \frac{\sum_{i=n}^{\infty} \gamma_i}{1 + \sum_{i=1}^{\infty} \gamma_i}. \quad (4.30)$$

We now compute m_i , the mean time to extinction (absorption by 0) when starting from state $i \geq 0$, supposing that absorption takes place almost surely, i.e., condition (4.28) is satisfied. When in state $i \geq 1$, the process remains there for an exponential time with mean $\frac{1}{\lambda_i + \mu_i}$, and then jumps to $i + 1$ (resp., $i - 1$) with the probability $\frac{\lambda_i}{\lambda_i + \mu_i}$ (resp., $\frac{\mu_i}{\lambda_i + \mu_i}$). First-step analysis therefore gives for $i \geq 1$

$$m_i = \frac{1}{\lambda_i + \mu_i} + \frac{\mu_i}{\lambda_i + \mu_i} m_{i-1} + \frac{\lambda_i}{\lambda_i + \mu_i} m_{i+1} \quad (4.31)$$

with the boundary condition $m_0 = 0$. For $i \geq 1$, (4.31) can be written

$$m_i - m_{i+1} = \frac{1}{\lambda_i} + \frac{\mu_i}{\lambda_i} (m_{i-1} - m_i).$$

Straightforward manipulations yield, for $n \geq 1$,

$$\frac{1}{\gamma_n} (m_n - m_{n+1}) = \sum_{i=1}^n \frac{1}{\lambda_i \gamma_i} - m_1. \quad (4.32)$$

Observe that $m_n \leq m_{n+1}$ for all n . If

$$\sum_{i=1}^{\infty} \frac{1}{\lambda_i \gamma_i} = \infty, \quad (4.33)$$

then necessarily $m_1 = \infty$. Consequently, $m_n = \infty$ for all $n \geq 1$. Suppose now that

$$\sum_{i=1}^{\infty} \frac{1}{\lambda_i \gamma_i} < \infty. \quad (4.34)$$

Then necessarily $m_1 < \infty$. Consequently, $m_n < \infty$ for all $n \geq 1$. From (4.32), letting $n \uparrow \infty$ therein,

$$m_1 = \sum_{r=1}^{\infty} \frac{1}{\lambda_r \gamma_r} - \lim_{n \rightarrow \infty} \frac{1}{\gamma_n} (m_n - m_{n-1}).$$

It can be shown that

$$\lim_{n \rightarrow \infty} \frac{1}{\gamma_n} (m_n - m_{n-1}) = 0. \quad (4.35)$$

This gives

$$m_1 = \sum_{i=0}^{\infty} \frac{1}{\lambda_i \gamma_i},$$

and going back to (4.32), for $n \geq 1$,

$$m_n = \sum_{i=1}^{\infty} \frac{1}{\lambda_i \gamma_i} + \sum_{k=1}^{n-1} \gamma_k \sum_{j=k+1}^{\infty} \frac{1}{\lambda_j \gamma_j}. \quad (4.36)$$

◇

Example 4.5. Linear Birth-and-Death Process

We shall apply the above results to the linear birth-and-death process for which

$$\lambda_n = n\lambda, \quad \mu_n = n\mu, \quad (4.37)$$

and $\lambda_0 = 0$, where $\lambda > 0$ and $\mu > 0$. In this case $\gamma_i = \left(\frac{\mu}{\lambda}\right)^i$, and therefore $\sum_{i=1}^{\infty} \gamma_i < \infty$, if and only if $\mu < \lambda$.

Therefore, $\mu \geq \lambda$ implies eventual extinction, as expected. For $\mu < \lambda$, the probability of extinction when starting from state n is u_n given by formula (4.30). In this case, for $n \geq 0$,

$$u_n = \left(\frac{\mu}{\lambda}\right)^n. \quad (4.38)$$

In the case where $\mu \geq \lambda$, which implies eventual extinction from any initial state n , formula (4.36) gives the mean time to extinction. We do the case with a single ancestor

$$m_1 = \sum_{i=1}^{\infty} \frac{1}{\lambda_i \gamma_i} = \frac{1}{\lambda} \sum_{i=1}^{\infty} \frac{1}{i} \left(\frac{\lambda}{\mu}\right)^i,$$

and therefore, $m_1 = \infty$ if $\lambda = \mu$, and if $\mu > \lambda$,

$$m_1 = \frac{1}{\lambda} \log \left(\frac{\mu}{\mu - \lambda} \right), \quad (4.39)$$

where we have used the expansion

$$\sum_{i=1}^{\infty} \frac{x^i}{i} = \log \left(\frac{1}{1-x} \right),$$

valid for $x \in (0, 1)$.

◇

5 Recurrence

5.1 Stationary Distribution Criterion of Ergodicity

We shall now define irreducibility, recurrence, transience, and positive recurrence for a regular jump HMC.

Definition 5.1. Irreducibility

A regular jump HMC is called *irreducible* if and only if the embedded discrete-time HMC is irreducible.

Definition 5.2. Recurrence

A state i is called *recurrent* if and only if it is recurrent for the embedded chain. Otherwise, it is called *transient*.

A recurrent state $i \in E$ is called t -positive recurrent if and only if $E_i[R_i] < \infty$, where R_i is the return time to state i (Example 4.1). Otherwise, it is called t -null recurrent.

Remark 5.1. We shall soon see that t -positive recurrence and n -positive recurrence (the latter is positive recurrence of the embedded chain) are not equivalent concepts. Also, observe that recurrence of a given state implies that this state is essential. Finally, in the same vein, note that irreducibility implies that all states are essential. \diamond

Definition 5.3. Recurrence and t -Invariant Measure

A t -invariant measure is a nontrivial vector $\nu = \{\nu(i)\}_{i \in E}$ such that for all $t \geq 0$,

$$\nu^T \mathbf{P}(t) = \nu^T. \quad (5.1)$$

Of course, an n -invariant measure is, by definition, an invariant measure for the embedded chain.

Theorem 5.1. t -Invariant Measure

Let the regular jump HMC $\{X(t)\}_{t \geq 0}$ with infinitesimal generator \mathbf{A} be irreducible and recurrent. Then there exists a unique (up to a multiplicative factor) t -invariant measure such that $\nu(i) > 0$ for all $i \in E$. Moreover, ν is obtained in one of the following ways. (i):

$$\nu(i) = E_0 \left[\int_0^{R_0} 1_{\{X(s)=i\}} ds \right], \quad (5.2)$$

where 0 is an arbitrary state and R_0 is the return time to state 0. (ii):

$$\nu(i) = \frac{\mu(i)}{q_i} = \frac{E_0 \left[\sum_{n=1}^{T_0} 1_{\{X_n=i\}} \right]}{q_i}, \quad (5.3)$$

where μ is the canonical invariant measure of the embedded chain relative to state 0, and T_0 is the return time to 0 of the embedded chain (eqn. (2.3) of Chapter 3). (iii):

$$\nu^T \mathbf{A} = 0. \quad (5.4)$$

Proof.

(α) We first show that (5.2) defines an invariant measure, that is, for all $j \in E$ and all $t \geq 0$,

$$v(j) = \sum_{k \in E} v(k) p_{kj}(t).$$

The right-hand side of the above equality is equal to

$$\begin{aligned} A &= \sum_{k \in E} E_0 \left[\int_0^\infty \mathbf{1}_{\{X(s)=k\}} \mathbf{1}_{\{s \leq R_0\}} ds \right] p_{kj}(t) \\ &= \int_0^\infty \sum_{k \in E} P_0(X(t+s) = j \mid X(s) = k) P_0(X(s) = k, s \leq R_0) ds. \end{aligned}$$

From the Markov property, since $\{s \leq R_0\}$ is an event expressible in terms of $(X(u), u \in [0, s])$,

$$P_0(X(t+s) = j \mid X(s) = k) = P_0(X(t+s) = j \mid X(s) = k, s \leq R_0).$$

Therefore,

$$\begin{aligned} A &= \int_0^\infty \sum_{k \in E} P_0(X(t+s) = j \mid X(s) = k, s \leq R_0) P_0(X(s) = k, s \leq R_0) ds \\ &= \int_0^\infty P_0(X(t+s) = j, s \leq R_0) ds = \int_0^\infty E_0 \left[\mathbf{1}_{\{X(t+s)=j\}} \mathbf{1}_{\{s \leq R_0\}} \right] ds \\ &= E_0 \left[\int_0^{R_0} \mathbf{1}_{\{X(t+s)=j\}} ds \right] = E_0 \left[\int_t^{t+R_0} \mathbf{1}_{\{X(u)=j\}} du \right] \\ &= E_0 \left[\mathbf{1}_{\{t \leq R_0\}} \int_t^{R_0} \mathbf{1}_{\{X(u)=j\}} du \right] - E_0 \left[\mathbf{1}_{\{t > R_0\}} \int_{R_0}^t \mathbf{1}_{\{X(u)=j\}} du \right] \\ &\quad + E_0 \left[\int_{R_0}^{R_0+t} \mathbf{1}_{\{X(u)=j\}} du \right]. \end{aligned}$$

From the strong Markov property applied to $\tau = R_0$,

$$E_0 \left[\int_{R_0}^{R_0+t} \mathbf{1}_{\{X(u)=j\}} du \right] = E_0 \left[\int_0^t \mathbf{1}_{\{X(u)=j\}} du \right].$$

Therefore,

$$\begin{aligned} A &= E_0 \left[\mathbf{1}_{\{t \leq R_0\}} \int_t^{R_0} \cdots \right] - E_0 \left[\mathbf{1}_{\{t > R_0\}} \int_{R_0}^t \cdots \right] + E_0 \left[\int_0^t \cdots \right] \\ &= E_0 \left[\int_0^{R_0} \mathbf{1}_{\{X(u)=j\}} du \right] = v(j). \end{aligned}$$

(β) We now show uniqueness. For this consider the *skeleton* chain $\{X(n)\}_{n \geq 0}$, which is irreducible, since the continuous-time HMC is irreducible (Problem 8.5.3). It is also

recurrent. Indeed for a fixed state i , consider the sequence Z_1, Z_2, \dots of successive sojourn times in state i of the state process. This sequence is infinite because the embedded chain is recurrent, i.i.d., exponential with mean $\frac{1}{q_i}$. Therefore, the event $\{Z_n > 1\}$ occurs infinitely often, and this implies that $\{X(n) = i\}$ also occurs infinitely often. Therefore, the skeleton being irreducible and recurrent has one and only one invariant measure. Since an invariant measure of the continuous-time chain is an invariant measure of the skeleton, uniqueness follows.

(γ) We now show that if ν is the (essentially unique) invariant measure of the continuous-time HMC, then μ defined by (5.3) is the (essentially unique) invariant measure of the embedded chain. For this, call T_0 the return time to 0 of the embedded chain. Then

$$\begin{aligned} \nu(i) &= E_0 \left[\int_0^{R_0} \mathbf{1}_{\{X(s)=i\}} ds \right] = E_0 \left[\sum_{n=0}^{T_0-1} S_{n+1} \mathbf{1}_{\{X_n=i\}} \right] \\ &= E_0 \left[\sum_{n=0}^{\infty} S_{n+1} \mathbf{1}_{\{X_n=i\}} \mathbf{1}_{\{n < T_0\}} \right] = \sum_{n=0}^{\infty} E_0[S_{n+1} \mathbf{1}_{\{X_n=i\}} \mathbf{1}_{\{n < T_0\}}] \\ &= \sum_{n=0}^{\infty} E_0[S_{n+1} \mid X_n = i, n < T_0] E_0[\mathbf{1}_{\{X_n=i\}} \mathbf{1}_{\{n < T_0\}}] \\ &= \sum_{n=0}^{\infty} E_0[S_{n+1} \mid X_n = i] E_0[\mathbf{1}_{\{X_n=i\}} \mathbf{1}_{\{n < T_0\}}], \end{aligned}$$

where the last equality follows from the strong Markov property at time τ_n . But $E_0[S_{n+1} \mid X_n = i] = \frac{1}{q_i}$, since conditionally on $X_n = i$, S_{n+1} is exponential of mean $1/q_i$. Therefore,

$$\nu(i) = \frac{1}{q_i} \sum_{n=0}^{\infty} E_0[\mathbf{1}_{\{X_n=i\}} \mathbf{1}_{\{n < T_0\}}] = \frac{1}{q_i} E_0 \left[\sum_{n=0}^{T_0-1} \mathbf{1}_{\{X_n=i\}} \right] = \frac{\mu(i)}{q_i}.$$

(δ) We show that ν given by (5.3) where μ is an invariant measure of the embedded chain, also satisfies (5.4). The transition matrix \mathbf{P} of the embedded chain is

$$p_{ii} = 0, \quad p_{ij} = \frac{q_{ij}}{q_i} \text{ if } i \neq j,$$

and the balance equation $\mu^T = \mu^T \mathbf{P}$ reads

$$\frac{\mu(i)}{q_i} q_i = \sum_{\substack{j \in E \\ j \neq i}} \frac{\mu(j)}{q_j} q_{ji},$$

which is just (5.4). □

Theorem 5.2. Positive vs. Null Recurrence

An irreducible recurrent regular jump HMC with invariant measure ν is t -positive recurrent if and only if

$$\sum_{i \in E} \nu(i) < \infty.$$

In this case the stationary probability π is related to the mean return times by

$$q_i \pi_i E_i[R_i] = 1. \quad (5.5)$$

Proof. Observing that

$$\sum_{i \in E} \nu(i) = E_0[R_0], \quad \nu(0) = E_0[\tau_1] = \frac{1}{q_0} \quad (5.6)$$

the proof is the same as for the corresponding result in discrete time. \square

Remark 5.2. We can give consistency to Remark 5.1 by using the relationship between the t -invariant measure ν and the n -invariant measure μ of the embedded chain in (5.3), which shows that there exist versions of the invariant measures μ of the embedded chain and ν of the continuous-time chain related by

$$\mu(i) = q_i \nu(i).$$

We see that all 4 possibilities concerning the convergence or divergence of the series $\sum_{i \in E} \mu(i)$ and $\sum_{i \in E} \nu(i)$ are open. The mean sojourn times q_i^{-1} make the difference, as is natural since they embody the deformation of the time scale when passing from the discrete time of the embedded chain to the continuous time of the regular jump HMC. \diamond

Definition 5.4. Ergodic

An irreducible regular jump HMC is called *ergodic* if it is t -positive recurrent.

Note that the notion of periodicity is irrelevant in continuous time. Note also that the embedded chain of an ergodic regular jump HMC may well be non ergodic, for two reasons. The first is the possibility that the embedded chain is null recurrent, as explained in Remark 5.2, while the second is the possibility that the embedded chain is periodic.

Theorem 5.3. Criterion of Ergodicity

The irreducible regular jump HMC with infinitesimal generator \mathbf{A} is ergodic if and only if there exists a probability π on E such that

$$\pi^T \mathbf{A} = 0. \quad (5.7)$$

Proof. In view of Theorem 5.1, only sufficiency has to be proven. Therefore, suppose that (5.7) holds for a probability distribution π . We shall prove that $\pi^T \mathbf{P}(t) = \pi^T$ for all $t \geq 0$. For this, define

$$p_{ij}^{(n)}(t) = P_i(X(t) = j, t < \tau_n),$$

where the τ_n are the transition times. A trajectory starting from state i contributes to $p_{ij}^{(n)}(t)$ either if it has no jump before t and $i = j$, or if it has a last jump (say from k to j) at a time $s \leq t$, and therefore at most $n - 1$ jumps before s . Therefore (Problem 8.5.4),

$$p_{ij}^{(n)}(t) = \delta_{ij} \exp(-q_i t) + \int_0^t \sum_{k \in E} p_{ik}^{(n-1)}(s) q_{kj} \exp(-q_j(t-s)) ds. \quad (5.8)$$

Thus

$$\sum_{i \in E} \pi(i) p_{ij}^{(n)}(t) = \pi(j) \exp(-q_j t) + \int_0^t \exp(-q_j(t-s)) \sum_{k \in E} q_{kj} \sum_{i \in E} \pi(i) p_{ik}^{(n-1)}(s) ds.$$

Now,

$$\sum_{i \in E} \pi(i) p_{ik}^{(1)}(s) = \pi(k) \exp(-q_k s) \leq \pi(k)$$

that is, $\pi^T \mathbf{P}^{(1)}(s) \leq \pi$, with the obvious notation. By induction, $\pi^T \mathbf{P}^{(n)}(s) \leq \pi^T$. Indeed, if the latter is true, then

$$\begin{aligned} \sum_{i \in E} \pi(i) p_{ij}^{(n+1)}(s) &\leq \pi(j) \exp(-q_j t) + \int_0^t \exp(-q_j(t-s)) \sum_{k \in E} q_{kj} \pi(k) ds \\ &= \pi(j) \exp(-q_j t) + \pi(j) q_j \int_0^t \exp(-q_j s) ds = \pi(j). \end{aligned}$$

Since the process is nonexplosive,

$$\lim_{n \uparrow \infty} p_{ij}^{(n)}(t) = p_{ij}(t) \text{ and } \sum_{j \in E} p_{ij}(t) = 1.$$

Therefore, by dominated convergence,

$$\sum_{i \in E} \pi(i) p_{ij}(t) \leq \pi(j).$$

Summation of both sides of the inequality shows that equality must hold. Therefore, π is a stationary distribution.

If the chain is transient, $\lim_{t \uparrow \infty} 1_{X(t)=j} = 0$, and therefore, by dominated convergence,

$$\lim_{t \rightarrow \infty} p_{ij}(t) = 0. \quad (5.9)$$

In particular, by dominated convergence, $(\pi^T \mathbf{P}(t))_j$ would tend to 0 as $t \rightarrow \infty$, a contradiction with $\pi^T \mathbf{P}(t) = \pi^T$. The chain is therefore recurrent. It is t -positive recurrent by Theorem 5.2. \square

5.2 Time Reversal

The concept of reversibility in continuous time is basically the same as in discrete time.

Let $\{X(t)\}_{t \geq 0}$ be a regular jump HMC with countable state space E , transition semigroup $\{\mathbf{P}(t)\}_{t \geq 0}$, and infinitesimal generator \mathbf{A} . Assume, moreover, that this chain is irreducible and ergodic, and in equilibrium, with the stationary distribution π . For arbitrary but fixed $T > 0$, consider the process $\{\tilde{X}(t)\}_{t \in [0, T]}$ obtained by time reversal of $\{X(t)\}_{t \in [0, T]}$, i.e.,

$$\tilde{X}(t) = X(T - t).$$

Then $\{\tilde{X}(t)\}$ is an HMC with transition semigroup $\{\tilde{\mathbf{P}}(t)\}_{t \geq 0}$ given by

$$\pi(i)\tilde{p}_{ij}(t) = \pi(j)p_{ji}(t) \quad (5.10)$$

(the proof is analogous to that of the corresponding result in discrete time), and therefore with infinitesimal generator $\tilde{\mathbf{A}}$, where

$$\pi(i)\tilde{q}_{ij} = \pi(j)q_{ji}, \quad (5.11)$$

as differentiation in (5.10) immediately shows.

In order not to depend on the choice of T , let us extend the definition of $X(t)$ to negative times by letting $\{X(-t)\}_{t \geq 0}$ be an HMC with transition semigroup $\{\tilde{\mathbf{P}}(t)\}_{t \geq 0}$ and independent of $\{X(t)\}_{t \geq 0}$ given $X(0)$. In this manner, we obtain, as in discrete time, a stationary HMC $\{X(t)\}_{t \in \mathbb{R}}$ with transition semigroup $\{\mathbf{P}(t)\}_{t \geq 0}$ and infinitesimal generator \mathbf{A} . (It will be assumed that the process has been modified so as to guarantee that the trajectories are right-continuous. In general, we shall take whatever convention is suitable as to right- or left-continuity, since this does not change the semigroup or the infinitesimal generator.)

The reversed process $\{X(-t)\}_{t \in \mathbb{R}}$ is therefore a regular jump HMC with transition semigroup $\{\tilde{\mathbf{P}}(t)\}_{t \in \mathbb{R}}$ and infinitesimal generator $\tilde{\mathbf{A}}$ given by (5.10) and (5.11), respectively.

Theorem 5.4. Reversibility Theorem

Let $\{X(t)\}_{t \in \mathbb{R}}$ be a stationary regular jump HMC on the countable state space E , and assume that it is irreducible and ergodic, with stationary distribution π . If its infinitesimal generator \mathbf{A} satisfies the detailed balance equations

$$\pi(i)q_{ij} = \pi(j)q_{ji}, \quad (5.12)$$

then the reversed process $\{X(-t)\}_{t \in \mathbb{R}}$ properly modified to be right-continuous is distributionwise equivalent to the direct process $\{X(t)\}_{t \in \mathbb{R}}$.

Proof. Comparing (5.11) and (5.12), we see that $\tilde{\mathbf{A}} \equiv \mathbf{A}$, and therefore, the direct and the reversed processes have the same infinitesimal generator. The conclusion follows from Theorem 4.3. \square

We shall not give examples of reversible chains, since we shall see a number of them when queuing networks are treated in Chapter 9.

For future reference, the reversal test analogous to Theorem 6.1 of Chapter 2 will be given because it will be needed in Chapter 9, again in connection with queuing networks.

Theorem 5.5. Reversal Test

Let \mathbf{A} be a stable and conservative generator on the countable state space E , and let π be a strictly positive probability distribution on E . Define $\tilde{\mathbf{A}}$ by (5.11). If for all $i \in E$

$$\sum_{j \in E, j \neq i} \tilde{q}_{ij} = q_i, \quad (5.13)$$

then $\pi^T \mathbf{A} = \pi$.

The proof is the same as for Theorem 6.1 of Chapter 2.

If $\{X(t)\}_{t \geq 0}$ is a regular jump HMC on E that is irreducible and stationary, with stationary distribution π and infinitesimal generator \mathbf{A} satisfying the detailed balance equation (5.13), then it is called reversible. Sometimes, the pair (\mathbf{A}, π) is called reversible. The idea of reversibility finds nice applications in queuing theory (see Chapter 9).

6 Long-Run Behavior

The limiting behavior of continuous-time HMCs follows from that of discrete-time HMCs in a usually straightforward manner. We treat the ergodic case and then the absorbing case.

6.1 Ergodic Chains

Theorem 6.1. *Long Run Behavior of Ergodic Regular Jump HMCs*

Let $\{X(t)\}_{t \geq 0}$ be an ergodic regular jump HMC with state space E and transition semigroup $\{\mathbf{P}(t)\}_{t \geq 0}$. Then, for all $i, j \in E$,

$$\lim_{t \rightarrow \infty} p_{ij}(t) = \pi(j), \quad (6.1)$$

where π is the (unique) stationary distribution.

Proof. We use the method of skeletons. Observe that the *skeleton* $\{X(n)\}_{n \geq 0}$ is an irreducible recurrent HMC (see the argument given in the proof of Theorem 5.1). It is positive recurrent, since it has π for a stationary distribution. Even though the embedded chain might be periodic, the skeleton $\{X(n)\}$ is not, because the sojourn times of the continuous-time chain in a given state i are i.i.d. exponentials, and this eliminates periodic behavior for the skeleton (Problem 8.6.2).

It is clear that two independent continuous-time HMCs with the same transition semigroup $\{\mathbf{P}(t)\}_{t \geq 0}$, but possibly different initial distributions, will meet at a finite integer random time, since their skeletons do. From this observation the result follows by the same coupling argument as in the discrete time case. \square

Theorem 6.2. *Ergodic Theorem*

Let $\{X(t)\}$ be ergodic and let π be its stationary distribution. Then

$$\lim_{t \uparrow \infty} \frac{1}{t} \int_0^t f(X(s)) ds = \sum_{i \in E} f(i) \pi(i), \quad P_\mu \text{ a.s.} \quad (6.2)$$

for all initial distributions μ and all $f : E \rightarrow \mathbb{R}$ such that $\sum_{i \in E} |f(i)| \pi(i) < \infty$.

Proof. The proof is analogous to that of the discrete-time case and is left to the reader. \square

Any method giving an estimate of the geometric rate of convergence for a discrete time chain with finite state space can in principle be used to obtain an analogous estimate for a

continuous time chain with finite state space. The trick is uniformization. Indeed, suppose $\{\mathbf{P}(t)\}_{t \geq 0}$ is the transition semigroup of an irreducible ergodic continuous-time HMC with infinitesimal generator \mathbf{A} and stationary distribution π . If the state space is finite, the chain is uniformizable, that is, its infinitesimal generator can be written

$$\mathbf{A} = \lambda(\mathbf{K} - I), \quad (6.3)$$

where $\lambda > 0$ and \mathbf{K} is an irreducible transition matrix. Moreover, the semigroup takes the form

$$\mathbf{P}(t) = \sum_{n=0}^{\infty} e^{-\lambda t} \frac{(\lambda t)^n}{n!} \mathbf{K}^n. \quad (6.4)$$

Also, \mathbf{K} has π as stationary distribution. Suppose that for an initial distribution μ it holds that for some $\alpha \in (0, 1)$ and some constant $C > 0$,

$$d_V(\mu^T \mathbf{K}^n, \pi) \leq C \alpha^n. \quad (6.5)$$

Since

$$\mu^T \mathbf{P}(t) - \pi^T = \sum_{n=0}^{\infty} e^{-\lambda t} \frac{(\lambda t)^n}{n!} (\mu^T \mathbf{K}^n - \pi^T),$$

it follows from the result of Problem 4.1.1 that

$$d_V(\mu^T \mathbf{P}(t), \pi) \leq \sum_{n=0}^{\infty} e^{-\lambda t} \frac{(\lambda t)^n}{n!} d_V(\mu^T \mathbf{K}^n, \pi),$$

and therefore

$$d_V(\mu^T \mathbf{P}(t), \pi) \leq \sum_{n=0}^{\infty} e^{-\lambda t} \frac{(\lambda t)^n}{n!} C \alpha^n.$$

Finally,

$$d_V(\mu^T \mathbf{P}(t), \pi) \leq C e^{-\lambda(1-\alpha)t}. \quad (6.6)$$

6.2 Absorbing Chains

We consider finite state space continuous-time HMCs, which can always be assumed to have the uniform structure, that is, with a generator of the form

$$\mathbf{A} = \lambda(\mathbf{K} - I), \quad (6.7)$$

where λ is the intensity of the clock and \mathbf{K} is the transition matrix of the subordinated chain (see Definition 2.2).

Absorption is discussed in terms of the subordinated chain, since it has the same structure as the continuous-time chain in terms of recurrent, transient, and absorbing states. In particular, the subordinated chain and the continuous-time chain have the same absorption probabilities. The difference arises when time intervenes. However, the adaptation of the results to the continuous-time chain is straightforward.

For instance, call $\Gamma = \{\gamma_{ij}\}_{i,j \in E}$ the potential matrix of the continuous-time chain, where

$$\gamma_{ij} = E_i \left[\int_0^\infty 1_{\{X(s)=j\}} ds \right]. \tag{6.8}$$

Since

$$\int_0^\infty 1_{\{X(s)=j\}} ds = \sum_{n=0}^\infty (T_{n+1} - T_n) 1_{\{\hat{X}_n=j\}},$$

we have, using the independence of the clock and of the subordinated chain, and $E[T_{n+1} - T_n] = \frac{1}{\lambda}$,

$$\gamma_{ij} = \frac{1}{\lambda} g_{ij}, \tag{6.9}$$

where \mathbf{G} is the potential matrix of the subordinated chain.

Suppose for simplicity that the state space E consists of one absorbing state, 1, and of the collection $\{2, \dots, r\}$ of transient states. Let us decompose the infinitesimal generator according to the partition $\{1\} + \{2, \dots, r\}$:

$$\mathbf{A} = \begin{pmatrix} 0 & 0 \\ a & \mathbf{T} \end{pmatrix}. \tag{6.10}$$

From the analogous decomposition of \mathbf{K} ,

$$\mathbf{K} = \begin{pmatrix} 1 & 0 \\ b & \mathbf{Q} \end{pmatrix},$$

and (6.7), it follows that

$$b = \frac{1}{\lambda} a, \quad \mathbf{Q} = I + \frac{\mathbf{T}}{\lambda}. \tag{6.11}$$

This identification permits us to adapt the results obtained for the discrete time case to continuous time. For instance, if we call S and Σ the restrictions of \mathbf{G} and Γ , respectively, to $\{2, \dots, r\}$, then

$$\Sigma = \frac{1}{\lambda} S = \frac{1}{\lambda} (I - \mathbf{Q})^{-1} = \frac{1}{\lambda} \left(-\frac{\mathbf{T}}{\lambda} \right)^{-1},$$

that is,

$$\Sigma = -\mathbf{T}^{-1}. \tag{6.12}$$

Calling τ the time to absorption in $\{2, \dots, r\}$ and supposing that $X(0)$ is distributed on the transient set according to ν , we have (Problem 8.6.1)

$$E_\nu[e^{iu\tau}] = 1 - iu\nu(\mathbf{T} + iuI)^{-1}, \tag{6.13}$$

and consequently,

$$P_\nu(\tau \leq t) = 1 - |\nu e^{t\mathbf{T}}|. \tag{6.14}$$

We refer to (Asmussen, 1987) for a discussion, based on the last formula, of the approximation of an arbitrary nonnegative random variable by a Markov chain absorption time.

Problems

8.1.1 Let (Y_1, \dots, Y_k) be a vector of independent exponential random variables with parameters $\lambda_1, \dots, \lambda_k$, respectively, that is independent of and has the same distribution as some other random vector (X_1, \dots, X_k) . Let $Z = \inf (X_1, \dots, X_k)$ and let J be the random index such that $X_J = Z$. Define R_1, \dots, R_k by $R_J = Y_J$, $R_i = X_i - X_J (i \neq J)$. Show that R_1, \dots, R_k are independent exponential random variables with parameters $\lambda_1, \dots, \lambda_k$, respectively, and are independent of (Z, J) .

8.1.2 Let $\{T_n\}_{n \geq 0}$ be an HPP on \mathbb{R}_+ , with intensity $\lambda > 0$. Recall that $T_0 \equiv 0$. For fixed $t > 0$, define $T_-(t) = \sup_{n \geq 0} \{T_n; T_n \leq t\}$ and $T_+(t) = \inf_{n \geq 0} \{T_n; T_n > t\}$. Define the *forward recurrence time* $F(t) = T_+(t) - t$ and the *backward recurrence time* $B(t) = t - T_-(t)$. Give the probability distribution of the vector $(B(t), F(t))$, and see what happens when $t \rightarrow \infty$. In particular, compare $\lim_{t \rightarrow \infty} E[B(t) + F(t)]$ with $E[S_n] = E[T_n - T_{n-1}]$ for fixed $n \geq 1$. Does this surprise you?

8.1.3 On the interval $[0, T]$ place n points independently and uniformly. Compute for the disjoint intervals $[a_1, b_1], \dots, [a_k, b_k] \subset [0, T]$ the distribution of (X_1, \dots, X_k) , where X_j is the number of points in $[a_j, b_j]$. What does this distribution become when n and T tend simultaneously to infinity in such a way that $\frac{n}{T}$ remains constant ($= \lambda > 0$) and, of course, the a_i and b_i remain fixed?

8.2.1 Show that a transition semigroup that is continuous at the origin is continuous at all times $t \geq 0$. Show that the transition semigroup of a regular jump HMC is continuous.

8.2.2 Consider the uniform Markov chain (see Definition 2.2 and Example 2.3) with state space $E = \{0, 1\}$; transition matrix $K = \begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix}$, where $\alpha, \beta \in (0, 1)$; and intensity $\lambda > 0$ for the underlying HPP. Find the transition semigroup $\{\mathbf{P}(t)\}_{t \geq 0}$. Suppose that $X(0) = 0$. Give the joint probability density of $(\tau_1, \tau_2 - \tau_1, \dots, \tau_n - \tau_{n-1})$, where τ_1, τ_2, \dots are the successive times when $\{X(t)\}_{t \geq 0}$ switches from one value to the other.

8.2.3 Let π be any probability distribution on the countable state space E , and for each $t \in \mathbb{R}_+$, let $X(t)$ be distributed according to π . Also suppose that the family $\{X(t)\}_{t \geq 0}$ is independent. Show that $\{X(t)\}_{t \geq 0}$ is a homogeneous Markov chain. Give its transition semigroup. What about its local characteristics? In particular, is $q_{ij} < \infty$? How does this fit with the main results of Section 2? Is the semigroup continuous?

8.2.4 Show that the flip-flop process is a special case of uniform HMC. Show that its infinitesimal generator is

$$\mathbf{A} = \begin{pmatrix} -\lambda & \lambda \\ \lambda & -\lambda \end{pmatrix}.$$

Show that the infinitesimal generator of a Poisson counting process of intensity λ is

$$\mathbf{A} = \begin{pmatrix} -\lambda & \lambda & 0 & 0 & \cdots \\ 0 & -\lambda & \lambda & 0 & \cdots \\ 0 & 0 & -\lambda & \lambda & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

8.2.5 Construct a uniform HMC that is a birth-and-death process with parameters satisfying $\sup_{i \in E} (\lambda_i + \mu_i) < \infty$.

8.2.6 Prove that the stochastic process in Definition 2.2 is indeed a HMC.

8.2.7 Prove (2.19).

8.3.1 Let $\{X(t)\}_{t \geq 0}$ be a *pure birth process*, that is, a regular jump HMC with state space $E = \mathbb{N}$ and infinitesimal generator of the form

$$\mathbf{A} = \begin{pmatrix} -\lambda_0 & \lambda_0 & 0 & 0 & \cdots \\ 0 & -\lambda_1 & \lambda_1 & 0 & 0 & \cdots \\ 0 & 0 & -\lambda_2 & \lambda_2 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

Suppose that $\lambda_n > 0$ for all $n \geq 0$. Compute $P_0(X(t) = n)$ for $n \geq 0, t \geq 0$.

8.3.2 Let $\{X(t)\}_{t \geq 0}$ be a birth-and-death process with the parameters $\lambda_n = n\lambda + a$ and $\mu_n = n\mu$ where $\lambda > 0, \mu > 0, a > 0$. Suppose that $X(0) = i$ is fixed. Show that $M(t) = E[X(t)]$ satisfies a differential equation and solve this equation. Do the same for the variance $S(t) = E[X(t)^2] - M(t)^2$.

8.4.1 Let \mathbf{A} be an essential generator on E . Define $\mathbf{P} = \{p_{ij}\}_{i, j \in E}$ by $p_{ii} = 0, p_{ij} = \frac{q_{ij}}{q_i}$ if $i \neq j$. Show that if the transition matrix \mathbf{P} is irreducible and recurrent, the generator is nonexplosive.

8.4.2 Same question as in Problem 8.4.1, only assuming that the state space is finite (\mathbf{P} is not supposed irreducible).

8.4.3 Same question as in Problem 8.4.1, without irreducibility and recurrence assumptions, but assuming $\sup_{i \in E} q_i < \infty$.

8.4.4 Let $\{X(t)\}_{t \geq 0}$ be a birth-and-death process with parameters $\lambda_n > 0$ and $\mu_n > 0$. Let w_n be the average time required to pass from state n to state $n + 1$. Find a recurrence equation for $\{w_n\}_{n \geq 0}$. Deduce from it the average time required to reach state $n \geq 1$ when starting from state 0.

8.4.5 Show that the birth-and-death generator with birth-and-death parameters $n\lambda$ and $n\mu$, respectively, where λ and μ are positive, is nonexplosive.

8.5.1 Consider an irreducible recurrent birth-and-death process with state space \mathbb{N} . Discuss the nature (positive or null) of recurrence of this process and of its embedded chain in terms of the birth-and-death parameters.

- 8.5.2 Show that an irreducible ergodic birth-an-death process in steady state is reversible.
- 8.5.3 Show that the skeleton of an irreducible regular jump HMC is irreducible.
- 8.5.4 Prove (5.8) formally.
- 8.6.1 Prove the two last results of Section 6.
- 8.6.2 Prove that the skeleton of a regular jump HMC cannot be periodic.

Poisson Calculus and Queues

1 Continuous-Time Markov Chains as Poisson Systems

1.1 Strong Markov Property of HPPs

The phrase *Poisson system* is an abbreviation of *dynamical system driven by Poisson processes*. Poisson systems form a special class of the so-called stochastic *discrete-event dynamical systems*. Their dynamics are expressed in a manner quite analogous in spirit to the recurrence equation of a discrete-time HMC, Theorem 2.1, Chapter 2. However, for Poisson systems, the *white noise* is a family of independent Poisson processes, and the functional, which associates to the present value of the state process and to the present value of the noise the future value of the state process, does not have—in general—a closed expression but takes the form of an algorithm.

To be more explicit, suppose that there is just one driving Poisson process, N , and let $X(t)$ be the state process at time t . In a time-homogeneous Poisson system, the state at time $t + a$ is given by

$$X(t + a) = \phi(X(t), N_{(t, t+a]}), \quad (1.1)$$

where $N_{(t, t+a]}$ is the restriction of N to $(0, a]$ and ϕ is an appropriate functional.

The analogy with the recurrence equation $X_{n+1} = f(X_n, Z_{n+1})$ satisfied by any discrete-time HMC justifies that we call $N_{(t, t+a]}$ the *noise in $(t, t + a]$* . It is intuitively obvious at this point that the state process of a homogeneous Poisson system is a continuous-time HMC. Conversely, as we shall see in the current section, a regular jump HMC can always be represented *distributionwise* as a homogeneous Poisson system. This means that one can associate with a given regular jump HMC a homogeneous Poisson system that has a state process with the same distribution as the original regular jump HMC. This representation is very useful in theory because the distribution of a continuous-time HMC is rarely given

by its transition semigroup, nor by its infinitesimal generator, and this is especially true for Markov chains arising in queuing theory. There, the description of the state process is in terms of a natural description featuring servers, processors, service disciplines, priorities, etc., and it is easily translated in terms of Poisson systems. The infinitesimal generator then immediately follows. Usually, the generator of such complex systems is obtained heuristically, and it is safe to rely on a methodical way of justifying intuition in case of controversy.

Note, however, that the Poisson system description of a continuous-time HMC has many advantages besides the fact that it is the only mathematically rigorous way to obtain the infinitesimal generator of most systems involving a complex mechanism for generating the state process. In particular, this type of description is the one used in the simulation of discrete event systems.

The Poisson systems approach basically sees the state process as the solution of a differential equation. This observation has deep consequences, and this *dynamical* approach has found a wider range of applications in the study of *stochastic systems driven by point processes*.

In order to develop the theory of Poisson systems independently from the previous chapter, we need a direct proof of the strong Markov property for HPPs. Actually, we want a version slightly stronger than the one that can be obtained from the strong Markov property of regular jump HMCs and the observation that a homogeneous Poisson process is a regular jump HMC.

The framework of this technical result is the following: There is an arbitrarily indexed family $\{N_\ell\}_{\ell \in L}$ of independent HPPs with the respective intensities $\{\lambda_\ell\}_{\ell \in L}$, and an arbitrary collection of random variables $\underline{Y} = \{Y_m\}_{m \in M}$ independent of $\{N_\ell\}_{\ell \in L}$.

For any random variable τ with values in $[0, +\infty]$ and arbitrary index $\ell \in L$, one defines N_ℓ^τ to be the restriction of N_ℓ to $\mathbb{R}_+ \cap [0, \tau]$: For all intervals $[a, b] \subset \mathbb{R}_+$,

$$N_\ell^\tau(a, b) = N_\ell((a, b] \cap [0, \tau]).$$

The point process N_ℓ^τ is “ N_ℓ before τ .” One also defines the point process $S_\tau N_\ell$, or “ N_ℓ after τ ,” by

$$S_\tau N_\ell(a, b) = N_\ell(\tau + a, \tau + b)$$

for all intervals $[a, b] \subset \mathbb{R}_+$.

Definition 1.1. *Extended Stopping Times of Point Processes*

The random variable τ with values in $\mathbb{R}_+ \cup \{+\infty\}$ is called a stopping time with respect to $\{N_\ell\}_{\ell \in L}$ and \underline{Y} if for all $t \in \mathbb{R}_+$, the event $\{\tau \leq t\}$ is expressible in terms of \underline{Y} and $\{N_\ell^t\}_{\ell \in L}$.

Theorem 1.1. *Extended Strong Markov Property of HPPs*

Let $\{N_\ell\}_{\ell \in L}$ and \underline{Y} be as above, and let τ be a stopping time with respect to $\{N_\ell\}_{\ell \in L}$ and \underline{Y} . Then, given $\{\tau < \infty\}$,

$$(\alpha) \{N_\ell^\tau\}_{\ell \in L} \text{ and } \underline{Y} \text{ are independent of } \{S_\tau N_\ell\}_{\ell \in L}.$$

(β) $\{S_\tau N_\ell\}_{\ell \in L}$ is a family of independent HPPs with respective intensities $\{\lambda_\ell\}_{\ell \in L}$.

Proof. It suffices to show that for all integers J , for all $u_1, \dots, u_J, v_1, \dots, v_J, w \in \mathbb{R}$, for all bounded intervals $[a_j, b_j]$ and $[c_j, d_j], j \in [1, J]$, for all real-valued random variables Z that are functionals of \underline{Y} , and for all indices $\ell_j \in L, j \in [1, J]$, it holds that

$$\begin{aligned} & E \left[\exp \left\{ i \left(\sum_{j=1}^J u_j (S_\tau N_{\ell_j})(a_j, b_j) + \sum_{j=1}^J v_j N_{\ell_j}^\tau(c_j, d_j) + wZ \right) \right\} 1_{\{\tau < \infty\}} \right] \\ &= E \left[\exp \left\{ i \sum_{j=1}^J u_j N_{\ell_j}(a_j, b_j) \right\} \right] E \left[\exp \left\{ i \sum_{j=1}^J v_j N_{\ell_j}^\tau(c_j, d_j) + wZ \right\} 1_{\{\tau < \infty\}} \right]. \end{aligned} \tag{1.2}$$

Indeed, letting $v_1 = \dots = v_J = w = 0$ in the above identity, we obtain

$$\begin{aligned} & E \left[\exp \left\{ i \sum_{j=1}^J u_j (S_\tau N_{\ell_j})(a_j, b_j) \right\} 1_{\{\tau < \infty\}} \right] \\ &= E \left[\exp \left\{ i \sum_{j=1}^J u_j N_{\ell_j}(a_j, b_j) \right\} \right] P(\tau < \infty), \end{aligned} \tag{1.3}$$

and this implies that given $\{\tau < \infty\}$, $\{S_\tau N_\ell\}_{\ell \in L}$ and $\{N_\ell\}_{\ell \in L}$ have the same distribution. Thus (β) is proven. Next, using (1.3), identity (1.2) becomes

$$\begin{aligned} & E \left[\exp \left\{ i \sum_{j=1}^J u_j (S_\tau N_{\ell_j})(a_j, b_j) + i \sum_{j=1}^J v_j N_{\ell_j}^\tau(c_j, d_j) + iwZ \right\} 1_{\{\tau < \infty\}} \right] P(\tau < \infty) \\ &= E \left[\exp \left\{ i \sum_{j=1}^J u_j (S_\tau N_{\ell_j})(a_j, b_j) \right\} 1_{\{\tau < \infty\}} \right] \\ &\quad \times E \left[\exp \left\{ i \sum_{j=1}^J v_j N_{\ell_j}^\tau(c_j, d_j) + iwZ \right\} 1_{\{\tau < \infty\}} \right]. \end{aligned}$$

This proves statement (α), by Theorem 5.4, Chapter 1.

We shall now proceed to the proof of (1.2). We do a simplified version that contains all the ingredients of the complete proof and saves us from notational hell. Calling N one of the N_ℓ 's, we shall prove that

$$\begin{aligned} & E[\exp\{iuS_\tau N(a, b) + ivN^\tau(c, d) + iwZ\}1_{\{\tau < \infty\}}] \\ &= E[\exp\{iuN(a, b)\}]E[\exp\{ivN^\tau(c, d) + iwZ\}1_{\{\tau < \infty\}}]. \end{aligned} \tag{1.4}$$

The left-hand side is

$$E[\exp\{iuN(\tau + a, \tau + b) + ivN((0, \tau] \cap (c, d)) + iwZ\}1_{\{\tau < \infty\}}],$$

that is, in the special case where τ takes a countable set of values $t_k \in \mathbb{R}_+, k \geq 1$, and maybe also the value $+\infty$,

$$\sum_{k=1}^{\infty} E[\exp\{iuN(t_k + a, t_k + b) + ivN((0, t_k] \cap (c, d)) + iwZ\}1_{\{\tau=t_k\}}].$$

But $1_{\{\tau=t_k\}}$ is a functional of Z and $\{N_\ell^t\}_{\ell \in L}$ and therefore the above expression becomes, in view of the independence and homogeneity properties of HPPs,

$$\begin{aligned} & \sum_{k=1}^{\infty} E[\exp\{iuN(t_k + a, t_k + b)\}]E[\exp\{ivN((0, t_k] \cap (c, d)) + iwZ\}1_{\{\tau=t_k\}}] \\ &= E[\exp\{iuN(a, b)\}] \left(\sum_{k=1}^{\infty} E[\exp\{ivN((0, t_k] \cap (c, d)) + iwZ\}1_{\{\tau=t_k\}}] \right), \end{aligned}$$

and this is the right-hand side of (1.4), as announced. The passage to an arbitrary stopping time is done in exactly the same manner as in the proof of the strong Markov property of regular jump HMCs (Theorem 4.1, Chapter 8). \square

1.2 From Generator to Markov Chain

We now consider the problem of *realization*, that is, of associating to a generator \mathbf{A} a continuous-time HMC admitting \mathbf{A} as an infinitesimal generator.

More precisely, let E be a countable space, and let $\mathbf{A} = \{q_{ij}\}_{i,j \in E}$ be a matrix of real numbers indexed by E , such that for all $i, j \in E$ with $i \neq j$,

$$q_i \in [0, \infty], \quad q_{ij} \in [0, \infty), \tag{1.5}$$

and

$$q_i < \infty, \quad q_i = \sum_{\substack{k \in E \\ k \neq i}} q_{ik}, \tag{1.6}$$

where $q_i = -q_{ii}$. When \mathbf{A} is the infinitesimal generator of a continuous transition semigroup on E , only (1.5) is guaranteed, whereas if it is the transition semigroup of a regular jump HMC, the stability and conservation properties (1.6) are also satisfied. Note that we did not prove the last point yet, and we shall therefore be forbidden to use this result until we have obtained it without getting into logical circles. Note also that when we speak about a generator \mathbf{A} satisfying by definition (1.5) and (1.6), we do not refer to any semigroup or any homogeneous continuous-time Markov chain. This is the difference between “generator” and “infinitesimal generator.”

Actually, we are going to *construct* a continuous-time HMC $\{X(t)\}_{t \geq 0}$ admitting \mathbf{A} as an infinitesimal generator. For this, we use a family $\{N_{i,j}\}_{\substack{i,j \in E \\ i \neq j}}$ of independent HPPs with respective intensities $\{q_{ij}\}_{\substack{i,j \in E \\ i \neq j}}$, and an initial state $X(0)$, independent of the above family of HPPs, and taking its values in E . The process is constructed as a jump process:

$$X(t) = X_n \text{ for } t \in [\tau_n, \tau_{n+1}),$$

where $\{\tau_n, X_n\}_{n \geq 0}$ is defined recursively as follows (see Figure 9.1.1). Let Δ be any dummy element not in E .

First $\tau_0 \equiv 0$ and $X_0 = X(0)$.

If $\tau_n < \infty$ and $X_n = X(\tau_n) = i \in E$, then $\tau_{n+1} - \tau_n$ is the first event of the competing HPPs $\{S_{\tau_n} N_{ij}\}_{j \in E, j \neq i}$; $\tau_{n+1} - \tau_n = \infty$ if and only if $q_i = 0$ (no events!), in which case the construction ends by setting $X_{n+m} = \Delta$ and $\tau_{n+m} = \infty$ for all $m \geq 1$. Otherwise, if $\tau_{n+1} - \tau_n < \infty$, X_{n+1} is the index $k \neq i$ such that $S_{\tau_n} N_{ik}$ is the first among the competing HPPs to produce an event.




Figure 9.1.1. First transition, obtained by competition

The process is defined by this procedure up to the explosion time $\tau_\infty = \lim_{n \uparrow \infty} \tau_n$. If $\tau_\infty = \infty$ P-a.s, it is a regular jump process on E .

Theorem 1.2.

If $\tau_\infty = \infty$ P-a.s, the process so constructed is a regular jump HMC with infinitesimal generator **A**.

Proof. The τ_n 's form a sequence of stopping times with respect to $\{N_{ij}\}_{\substack{i, j \in E \\ j \neq i}}$ and $X(0)$. In view of the strong Markov property for HPPs,

$$P(X_{n+1} = j, \tau_{n+1} - \tau_n > a \mid X_0, \dots, X_n, \tau_1, \dots, \tau_n) = P(X_{n+1} = j, \tau_{n+1} - \tau_n > a \mid X_n),$$

and by the competition theorem,

$$P(X_{n+1} = j, \tau_{n+1} - \tau_n > a \mid X_n = i) = e^{-q_i a} \frac{q_{ij}}{q_i}$$

(where $\frac{q_{ij}}{q_i} = 0$ if $q_i = 0$).

By construction, for a given time t , the process after time t depends only upon $X(t)$ and the HPPs $\{N_{ij}\}_{\substack{i, j \in E \\ i \neq j}}$ after t , and is therefore independent of the process before t , given $X(t)$,

since the process before t depends only upon $X(0)$ and the HPPs $\{N_{ij}\}_{\substack{i,j \in E \\ i \neq j}}$ before t . This proves the Markov property. Also, for all $t, s \geq 0$,

$$X(t + s) = \psi(X(t), \{S_t N_{k\ell}\}_{\substack{k,\ell \in E \\ k \neq t}}, s),$$

for a functional ψ that is given explicitly by the construction. Therefore, on $\{X(t) = i\}$,

$$X(t + s) = \psi(i, \{S_t N_{ij}\}, s)$$

has the same distribution as $\psi(i, \{N_{k,\ell}\}_{\substack{k,\ell \in E \\ k \neq t}}, s)$. This proves homogeneity.

We can now check that $\{X(t)\}_{t \geq 0}$ admits \mathbf{A} as an infinitesimal generator. For instance, we check that $\lim_{t \downarrow 0} \frac{1}{t} P_i(X(t) = j) = q_{ij}$. For this, observe that if $X(t) \neq X(0)$, necessarily $\tau_1 < t$, and write

$$\begin{aligned} P_i(X(t) = j) &= P_i(\tau_2 \leq t, X(t) = j) + P_i(\tau_2 > t, X(t) = j) \\ &= P_i(\tau_2 \leq t, X(t) = j) + P_i(\tau_2 > t, X_1 = j, \tau_1 < t) \\ &= P_i(\tau_2 \leq t, X(t) = j) + P_i(X_1 = j, \tau_1 < t) - P_i(\tau_2 \leq t, X_1 = j, \tau_1 \leq t). \end{aligned}$$

Now,

$$P_i(X_1 = j, \tau_1 < t) = (1 - e^{-q_i t}) \frac{q_{ij}}{q_i},$$

and therefore

$$\lim_{t \downarrow 0} \frac{1}{t} P(X_1 = j, \tau_1 < t) = q_{ij}.$$

It therefore remains to show that $P_i(\tau_2 \leq t, X(t) = j)$ and $P_i(\tau_2 \leq t, X_1 = j, \tau_1 \leq t)$ are $o(t)$. Both terms are bounded by $P_i(\tau_2 \leq t)$, and

$$\begin{aligned} P_i(\tau_2 \leq t) &\leq P_i(\tau_1 \leq t, \tau_2 - \tau_1 \leq t) = \sum_{\substack{k \in E \\ k \neq j}} P_i(\tau_1 \leq t, X_1 = k, \tau_2 - \tau_1 \leq t) \\ &= \sum_{\substack{k \in E \\ k \neq j}} (1 - e^{-q_i t}) p_{ik} (1 - e^{-q_k t}) \\ &= (1 - e^{-q_i t}) \sum_{\substack{k \in E \\ k \neq i}} p_{ik} (1 - e^{-q_k t}). \end{aligned}$$

The conclusion follows because $(1 - e^{-q_i t})$ is $O(t)$, or identically zero if $q_i = 0$, and because $\lim_{t \downarrow 0} \sum_{\substack{k \in E \\ k \neq i}} p_{ik} (1 - e^{-q_k t}) = 0$, by dominated convergence. \square

Example 1.1. *Existence of Birth-and-Death processes*

Recall that we defined birth-and-death processes as regular jump HMCs on $E = \mathbb{N}$ with infinitesimal generators of the form

$$\mathbf{A} = \begin{pmatrix} -\lambda_0 & \lambda_0 & 0 & \dots \\ \mu_1 & -(\lambda_1 + \mu_1) & \lambda_1 & 0 & \dots \\ 0 & \mu_2 & -(\lambda_2 + \mu_2) & \lambda_2 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

The above construction shows the existence of such a process under Reuter's necessary and sufficient condition of nonexplosion (Theorem 4.5, Chapter 8). If Reuter's condition is not satisfied, we obtain a nonregular jump HMC. \diamond

Paying Two Debts

We have now reached a position that enables us to pay two debts: we can finish the proof of Theorem 4.2 of Chapter 8, and we prove Theorem 2.2 of Chapter 8.

First debt: the proof of Theorem 4.2 of Chapter 8. Recall that the starting point was a regular jump HMC $\{X(t)\}_{t \geq 0}$ with infinitesimal generator \mathbf{A} . We wanted to prove (α) and (β) of Theorem 4.2, Chapter 8, but we were only able to prove something close, namely (α') and (β') where q_i was replaced by $\lambda(i) \in \mathbb{R}_+$ (only known to exist but not yet identified with q_i) and where $\frac{q_{ij}}{q_i}$ was replaced by $P_i(X(\tau_1) = j)$ (not yet identified with $\frac{q_{ij}}{q_i}$). We shall now proceed to the required identifications.

Define the generator \mathbf{A}' on E by

$$q'_i = \lambda(i), \quad q'_{ij} = \lambda(i)P_i(X(\tau_1) = j).$$

Construct $\{X'(t)\}_{t \geq 0}$, the regular jump HMC associated with \mathbf{A}' via the construction of Section 1. Then $\{X'(t)\}_{t \geq 0}$ and $\{X(t)\}_{t \geq 0}$ are probabilistically equivalent, by construction. In particular, $\mathbf{A}' = \mathbf{A}$, which is the identification we wanted.

Second debt: Theorem 2.2 of Chapter 8. We claimed that the infinitesimal generator of a regular jump HMC is necessarily stable and conservative.

Indeed, a regular jump HMC is strongly Markovian, and therefore has the regenerative structure of the weak version of Theorem 4.2 of Chapter 8, where $\lambda(i)$ replaced q_i and $\lambda(i)P_i(X(\tau_1) = j)$ replaced q_{ij} (we cannot use Theorem 4.2 directly because there is an implicit claim that $q_i < \infty$ and $q_i = \sum_{j \neq i} q_{ij}$ in it; the weak version does not have this a priori). But the identification $\mathbf{A}' = \mathbf{A}$ we made above is $q_i = \lambda(i)$, $q_{ij} = \lambda(i)P_i(X(\tau_1) = j)$, and therefore $q_i < \infty$, $\sum_{j \neq i} q_{ij} = q_i$. \square

2 Stochastic Calculus of Poisson Processes

2.1 Counting Integrals and the Smoothing Formula

The so-called *stochastic calculus of Poisson processes* (more briefly: *Poisson calculus*) is based on the notion of integral with respect to a counting process, an elementary notion that has little to do with probability, although the counting processes involved in the Poisson calculus are stochastic processes.

Let $\{T_n\}_{n \geq 1}$ be a simple and nonexplosive point process on $(0, \infty]$ and let $\{N(t)\}_{t \geq 0}$ be the associated counting function.

Let $(a, b) \subset \mathbb{R}_+$ be an arbitrary finite interval and let $\{Z(t)\}_{t \geq 0}$ be an arbitrary complex-valued stochastic process. The following equality is a *definition* of its left-hand side:

$$\int_{(a,b)} Z(t) dN(t) \stackrel{\text{def}}{=} \sum_{n \geq 1} Z(T_n) 1_{(a,b]}(T_n). \quad (2.1)$$

This quantity is called the *Stieltjes–Lebesgue integral*, or the *counting integral*, of Z with respect to N on $(a, b]$. Clearly, this definition could have been given with any bounded set $C \subset \mathbb{R}_+$ replacing $(a, b]$.

For bounded intervals, the definition of the integral does not present difficulties because the sum involved has a finite number of terms. For unbounded intervals, one may have to restrict the integrands in order to have a meaningful sum. For instance,

$$\int_{(0,\infty)} Z(t) dN(t) = \sum_{n \geq 1} Z(T_n) 1_{(0,\infty)}(T_n)$$

is always defined if Z is nonnegative, even though it may then be infinite.

Note that if $Z(t)$ is defined on $(a, b]$ by

$$Z(t) = c_j \in \mathbb{C} \text{ on } (t_j, t_{j+1}], \quad (2.2)$$

for all $j \in [1, K]$, where $a = t_0 < t_1 < \dots < t_{K+1} = b$, then

$$\int_{(a,b)} Z(t) dN(t) = \sum_{j=0}^K c_j \{N(t_{j+1}) - N(t_j)\}. \quad (2.3)$$

Although there is not much in the definition of a stochastic integral with respect to a point process, the stochastic calculus that such integrals generate is rather powerful. It is based on the result below.

Let $\{N_i\}_{i \in I}$ be an arbitrarily indexed family of homogeneous Poisson processes on \mathbb{R}_+ , of respective intensities $\{\lambda_i\}_{i \in I}$. Let \underline{Y} be an arbitrarily indexed family of random variables independent of $\{N_i\}_{i \in I}$.

Theorem 2.1. Smoothing Formula

Fix $j \in I$. Let $\{Z(t)\}_{t \geq 0}$ be a complex stochastic process with left-continuous trajectories, and such that for all $t \geq 0$, the random variable $Z(t)$ is a functional of \underline{Y} , of $\{N_k\}_{k \in I, k \neq j}$, and of the restriction of N_j to $(0, t]$.

Suppose at least one of the two conditions below is satisfied

- (i) $\{Z(t)\}_{t \geq 0}$ is nonnegative.
- (ii) $\{Z(t)\}_{t \geq 0}$ is complex-valued, and $E \left[\int_0^\infty |Z(t)| N_j(dt) \right]$ or $E \left[\int_0^\infty |Z(t)| \lambda_j dt \right]$ is finite.

Then the *smoothing formula* holds true:

$$E \left[\int_0^\infty Z(s) dN_j(s) \right] = E \left[\int_0^\infty Z(s) \lambda_j ds \right]. \quad (2.4)$$

Proof. A. We first treat the case where $\{Z(t)\}_{t \geq 0}$ is real, nonnegative, and bounded, and prove that

$$E \left[\int_0^T Z(s) dN_j(s) \right] = E \left[\int_0^T Z(s) \lambda_j ds \right], \tag{2.5}$$

where $T < \infty$. The hypothesis that it has left-continuous trajectories means that for all $\omega \in \Omega$, the function $t \rightarrow Z(t, \omega)$ is left-continuous. In particular, if one defines for all $n \geq 1$, all $\omega \in \Omega$ and all $t \geq 0$,

$$Z_n(t, \omega) = \sum_{k=0}^{2^n-1} Z \left(\frac{kT}{2^n}, \omega \right) 1_{(\frac{kT}{2^n}, \frac{(k+1)T}{2^n}]}(t),$$

then for all $\omega \in \Omega$ and all $t \in [0, T]$,

$$\lim_{n \uparrow \infty} Z_n(t, \omega) = Z(t, \omega).$$

We first check that (2.5) is true when $\{Z(t)\}_{t \geq 0}$ is replaced by its approximation $\{Z_n(t)\}_{t \geq 0}$. Indeed, the left-hand side of this equality is then

$$\begin{aligned} E \left[\int_0^T Z_n(t) dN(t) \right] &= E \left[\sum_{k=0}^{2^n-1} Z \left(\frac{kT}{2^n} \right) \left\{ N_j \left(\frac{(k+1)T}{2^n} \right) - N_j \left(\frac{kT}{2^n} \right) \right\} \right] \\ &= \sum_{k=0}^{2^n-1} E \left[Z \left(\frac{kT}{2^n} \right) \left\{ N_j \left(\frac{(k+1)T}{2^n} \right) - N_j \left(\frac{kT}{2^n} \right) \right\} \right]. \end{aligned}$$

But $Z \left(\frac{kT}{2^n} \right)$ is a functional of \underline{Y} , $\{N_k\}_{k \in I, k \neq j}$ and the restriction of N_j to $(0, \frac{kT}{2^n}]$, and therefore it is independent of $N_j \left(\frac{(k+1)T}{2^n} \right) - N_j \left(\frac{kT}{2^n} \right)$, so that the last term of the above chain of equalities is equal to

$$\begin{aligned} \sum_{k=0}^{2^n-1} E \left[Z \left(\frac{kT}{2^n} \right) \right] E \left[N_j \left(\frac{(k+1)T}{2^n} \right) - N_j \left(\frac{kT}{2^n} \right) \right] &= \sum_{k=0}^{2^n-1} E \left[Z \left(\frac{kT}{2^n} \right) \right] \lambda_j \frac{T}{2^n} \\ &= E \left[\sum_{k=0}^{2^n-1} Z \left(\frac{kT}{2^n} \right) \lambda_j \frac{T}{2^n} \right] \\ &= E \left[\int_0^T Z_n(s) \lambda_j ds \right]. \end{aligned}$$

Therefore,

$$E \left[\int_0^T Z_n(s) dN(s) \right] = E \left[\int_0^T Z_n(s) \lambda_j ds \right]. \tag{2.6}$$

Calling K the upper bound of $Z(t, \omega)$, we see that $\int_0^t Z_n(t) \lambda_j dt \leq K \lambda_j T$ and $\int_0^T Z_n(t) dN_j(t) \leq K N_j(T)$. Also, $E[N_j(T)] = \lambda_j T < \infty$. Therefore, we may apply the dominated convergence theorem to both sides of (2.6) to obtain (2.5).

B. we now treat the nonnegative unbounded case. Defining

$$g_K(x) = \begin{cases} 1 & \text{if } x \leq K, \\ 0 & \text{if } x \geq K + 1, \\ -x + K + 1 & \text{if } K \leq x \leq K + 1, \end{cases}$$

the process $\{g_K(Z(t))\}_{t \geq 0}$ verifies the conditions for (2.5) and therefore

$$E \left[\int_0^T g_K(Z(s)) dN_j(s) \right] = E \left[\int_0^T g_K(Z(s)) \lambda_j ds \right]$$

which yields (2.5) as $K \rightarrow \infty$, by monotone convergence.

It suffices to let $T \uparrow \infty$ and to invoke the monotone convergence theorem again to obtain (2.4) in the nonnegative unbounded case.

Note that the quantities in (2.4) can now very well be infinite.

C. The real-valued case follows easily, by first considering separately the positive and negative parts of the integrand. The complex case is a direct consequence of the real case when one considers separately the real and imaginary parts. \square

2.2 Kolmogorov's Forward System via Poisson Calculus

Our first application of the smoothing formula is to the proof that the forward integral system of Kolmogorov holds true for all regular jump HMCs:

$$\mathbf{P}(t) = I + \int_0^t \mathbf{P}(s) \mathbf{A} ds. \quad (2.7)$$

Proof. We return to the construction, in Section 1, of a regular jump HMC $\{X(t)\}_{t \geq 0}$ with values in E and with a stable, conservative, and nonexplosive generator \mathbf{A} by means of a family of independent HPPs $\{N_{ij}\}_{i,j \in E}$ of respective intensities $\{q_{ij}\}_{i,j \in E}$. Let $f : E \rightarrow \mathbb{R}$ be a nonnegative function. Since the trajectories of $\{f(X(t))\}_{t \geq 0}$ are right-continuous step processes, whose discontinuity times are discontinuity times of $\{X(t)\}_{t \geq 0}$, we can write

$$\begin{aligned} f(X(t)) - f(X(0)) &= \sum_{s \in (0,t]} \{f(X(s)) - f(X(s-))\} \\ &= \sum_{\substack{i,j \in E \\ i \neq j}} \int_{(0,t]} \{f(X(s)) - f(X(s-))\} dN_{ij}(s) \\ &= \sum_{\substack{i,j \in E \\ i \neq j}} \int_{(0,t]} \{f(j) - f(i)\} 1_{\{X(s-)=i\}} dN_{ij}(s). \end{aligned}$$

Therefore, reorganizing this equality and taking expectations, we obtain from the smoothing formula

$$E_\mu[f(X(t))] + \sum_{\substack{i,j \in E \\ i \neq j}} E_\mu \left[\int_{(0,t]} f(i) 1_{\{X(s-)=i\}} q_{ij} ds \right]$$

$$= E_\mu[f(X(0))] + \sum_{\substack{i, j \in E \\ i \neq j}} E_\mu \left[\int_{(0, t]} f(j) 1_{\{X(s-) = i\}} q_{ij} ds \right],$$

where μ is the initial distribution of the chain. Since

$$E_\mu \left[\int_{(0, t]} 1_{\{X(s-) = k\}} ds \right] = \int_{(0, t]} E_\mu [1_{\{X(s) = k\}}] ds = \int_0^t P_\mu(X(s) = k) ds,$$

we have, using the conservation equality $q_i = \sum_{j \in E, j \neq i} q_{ij}$,

$$\begin{aligned} E_\mu[f(X(t))] + \sum_{i \in E} f(i) q_i \int_0^t P_\mu(X(s) = i) ds \\ = E_\mu[f(X(0))] + \sum_{\substack{i, j \in E \\ i \neq j}} f(j) q_{ij} \int_0^t P_\mu(X(s) = i) ds. \end{aligned} \quad (2.8)$$

Taking $f(j) = 1_{\{j=k\}}$ for a fixed state k , and $\mu = \delta_\ell$, the Dirac distribution at ℓ , the above equalities reduce to

$$p_{\ell k}(t) + q_k \int_0^t p_{\ell k}(s) ds = \delta_{\ell k} + \sum_{\substack{i \in E \\ i \neq k}} q_{ik} \int_0^t p_{\ell i}(s) ds.$$

Since the left-hand side is finite, so is the right-hand side, and therefore one can reorganize this equality to obtain

$$p_{\ell k}(t) = \delta_{\ell k} + \int_0^t \left\{ -p_{\ell k}(s) q_k + \sum_{\substack{i \in E \\ i \neq k}} p_{\ell i}(s) q_{ik} \right\} ds. \quad (2.9)$$

This is true for all $\ell, k \in E$, and the corresponding equations are summarized by (2.7).

This is the integral form of the forward Kolmogorov system. Note that (2.9) is of the form

$$f(t) = a + \int_0^t g(s) ds,$$

where g is locally integrable. Therefore, by Theorem 8.17 of (Rudin, 1966), $f'(t) = g(t)$ almost everywhere. Thus Kolmogorov's forward differential system holds almost everywhere. Similarly, we find from (2.8) that for any initial distribution $\mu = \{\mu(i)\}_{i \in E}$, the vector $\mu(t) = \{\mu_i(t)\}_{i \in E}$, where $\mu_i(t) = P(X(t) = i)$, satisfies almost-everywhere

$$\frac{d}{dt} \mu^T(t) = \mu^T(t) \mathbf{A}. \quad (2.10)$$

In particular, if a stationary distribution π exists, it satisfies the global balance equation

$$\pi^T \mathbf{A} = 0. \quad (2.11)$$

Note that in the course of the proof of Kolmogorov's forward system, we have proven the following result:

Theorem 2.2. *Dynkin's Formula*

Let $\{X(t)\}_{t \geq 0}$ be a regular jump HMC with state space E , infinitesimal generator A , and initial distribution μ . Let $f : E \rightarrow \mathbb{R}_+$ satisfy $E_\mu[f(X_0)] < \infty$.

If either one of the following two conditions is satisfied,

$$(i) E_\mu \left[\int_0^t q_{X(s)} f(X(s)) ds \right] < \infty,$$

$$(ii) E_\mu[f(X(t))] < \infty,$$

then the other is satisfied and

$$E_\mu[f(X(t))] = E_\mu[f(X(0))] + E_\mu \left[\int_0^t (Af)(X(s)) ds \right] \quad (2.12)$$

for all $t \geq 0$, where

$$(Af)(j) = \sum_{\substack{i \in E \\ i \neq j}} q_{ij} f(i) - q_j f(j).$$

2.3 Watanabe's Characterization of Poisson Processes

We shall now give another application of the smoothing formula that is a typical example of the Poisson calculus. It is a characterization of Poisson processes also called the *martingale characterization* of Poisson processes.

Theorem 2.3. *Watanabe's Theorem*

Let $\{N_i\}_{i \in I}$ be an arbitrarily indexed family of simple nonexplosive point processes on \mathbb{R}_+ , and let \underline{Y} be a family of random variables independent of $\{N_i\}_{i \in I}$. Let $\{\lambda_i\}_{i \in I}$ be a family of positive real numbers. Suppose that

$$E \left[\int_0^T Z(t) dN_j(t) \right] = E \left[\int_0^T Z(t) \lambda_j dt \right] \quad (2.13)$$

holds true for all $j \in I$, all $T > 0$, and for all nonnegative real-valued stochastic processes $\{Z(t)\}_{t \geq 0}$ with left-continuous trajectories, and such that for all $t \geq 0$, the random variable $Z(t)$ is a functional of \underline{Y} , $\{N_i\}_{i \in I - \{j\}}$ and the restriction of N_j to $(0, t]$.

Then $\{N_i\}_{i \in I}$ is a family of independent homogeneous Poisson processes with the respective intensities $\{\lambda_i\}_{i \in I}$, and \underline{Y} and $\{N_i\}_{i \in I}$ are independent.

Proof. (Take $j = 1$ without loss of generality.) We show that it is enough to prove that

$$E \left[1_A e^{iuN_1(a,b)} \right] = P(A) \exp\{(e^{iu} - 1)\lambda_1(b - a)\}, \quad (2.14)$$

where $(a, b) \subset \mathbb{R}_+$, $u \in \mathbb{R}$ are arbitrary, and A is an arbitrary event expressible in terms of \underline{Y} , $\{N_i\}_{i \in I - \{1\}}$ and the restriction of N_1 to $(0, a]$. Indeed, with $A = \Omega$ in (2.14),

$$E[e^{iuN_1(a,b)}] = \exp\{(e^{iu} - 1)\lambda_1(b - a)\},$$

which is the characteristic function of a Poisson random variable of mean $\lambda_1(b-a)$. Equality (2.14) reads

$$E[1_A e^{iuN_1(a,b)}] = P(A)E[e^{iuN_1(a,b)}].$$

By Theorem 5.3 of Chapter 1, it follows that $N_1(a, b]$ and A are independent. But A is an arbitrary event expressible in terms of \underline{Y} , $\{N_i\}_{i \in I - \{1\}}$, and the restriction of N_1 to $(0, a]$. Therefore, $N_1(a, b]$ is independent of \underline{Y} , $\{N_i\}_{i \in I - \{1\}}$, and the restriction of N_1 to $(0, a]$. Since $(a, b]$ is arbitrary, we conclude that N_1 has independent increments and that it is independent of \underline{Y} and $\{N_i\}_{i \in I - \{1\}}$.

We now proceed to the proof of (2.14). For this we consider the process

$$X(t) = 1_A e^{iuN_1(a,t)},$$

and observe that it is piecewise constant with discontinuity times located at the events of N_1 , which we denote by $\{T_n^1\}_{n \geq 1}$. Therefore,

$$X(t) = X(0) + \sum_{n \geq 1} \{X(T_n^1) - X(T_{n-}^1)\} 1_{(a,b)}(T_n^1).$$

But $X(0) = 1_A$, and for any $T_n^1 \in (a, b]$,

$$X(T_n^1) = e^{iuN_1(a, T_n^1)} = e^{iu(N_1(a, T_n^1) + 1)} = e^{iuN_1(a, T_n^1)} e^{iu} = X(T_{n-}^1) e^{iu}.$$

Therefore,

$$X(t) = 1_A + \sum_{n \geq 1} X(T_{n-}^1)(e^{iu} - 1) 1_{(a,b)}(T_n^1)$$

or, in other notation,

$$X(t) = 1_A + \int_a^b X(s-)(e^{iu} - 1) dN_1(s). \quad (2.15)$$

Since $Z(t) = X(t-)(e^{iu} - 1)$ defines a bounded left-continuous complex-valued stochastic process $\{Z(t)\}_{t \geq 0}$ such that for all $t \geq 0$, the random variable $Z(t)$ is expressible in terms of \underline{Y} , $\{N_i\}_{i \in I - \{1\}}$ and the restriction of N_1 to $(0, t]$, equality (2.13) holds true by hypothesis. Therefore,

$$E \left[\int_a^b X(s-)(e^{iu} - 1) dN_1(s) \right] = E \left[\int_a^b X(s-)(e^{iu} - 1) \lambda_1 ds \right].$$

Now, for each $\omega \in \Omega$, $X(t-, \omega) = X(t, \omega)$ except on a countable set, and therefore in the right-hand side of the above equality one may replace $X(t-)$ by $X(t)$. Taking expectations in (2.15) therefore yields

$$E[X(b)] = P(A) + \int_a^b E[X(s)](e^{iu} - 1) \lambda_1 ds.$$

This being true for all $b \geq a$, (2.14) follows. \square

The following result can be obtained as an application of Watanabe's theorem.

Theorem 2.4. Aggregation

Consider a regular jump HMC $\{X(t)\}_{t \geq 0}$ with state space E and infinitesimal generator \mathbf{A} . Let $\tilde{E} = \{\alpha, \beta, \dots\}$ be a partition of E , and define the process $\{\tilde{X}(t)\}_{t \geq 0}$ taking its values in \tilde{E} by

$$\tilde{X}(t) = \alpha \Leftrightarrow X(t) \in \alpha. \tag{2.16}$$

Suppose that for all $\alpha \in \tilde{E}, i \in \alpha, \beta \in \tilde{E}$ with $\alpha \neq \beta$,

$$\sum_{j \in \beta} q_{ij} = \tilde{q}_{\alpha\beta} \tag{2.17}$$

(with the meaning that the left-hand side is independent of $i \in \alpha$). Then $\{\tilde{X}(t)\}_{t \geq 0}$ is a regular jump HMC with state space \tilde{E} and infinitesimal generator $\tilde{\mathbf{A}}$, with off-diagonal terms $\tilde{q}_{\alpha\beta}$ given by (2.17).

Proof. This is a statement concerning the distribution of process $\{X(t)\}_{t \geq 0}$, and therefore we may suppose that $\{X(t)\}_{t \geq 0}$ is generated as in Section 1 as a Poisson system, using independent HPPs $\{N_{ij}\}_{i, j \in E: i \neq j}$ with respective intensities $\{q_{ij}\}_{i, j \in E: i \neq j}$ and an initial state $X(0)$ independent of the above HPPs. Then, for $f : \tilde{E} \rightarrow \mathbb{R}, s \in [0, t]$,

$$\begin{aligned} f(\tilde{X}(t)) &= f(\tilde{X}(s)) + \sum_{\substack{i, j \neq i \\ i \neq j}} \int_{\{s, t\}} \{f(\tilde{X}(u)) - f(\tilde{X}(u-))\} 1_{\{X(u-)=i\}} dN_{ij}(u) \\ &= f(\tilde{X}(s)) + \sum_{\substack{\alpha, \beta \in \tilde{E} \\ \alpha \neq \beta}} \int_{\{s, t\}} \{f(\beta) - f(\alpha)\} \sum_{i \in \alpha} \left(1_{\{X(u-)=i\}} \left(\sum_{j \in \beta} dN_{ij}(u) \right) \right). \end{aligned}$$

Define for all $\alpha, \beta \in \tilde{E}, \alpha \neq \beta$, the point process $\tilde{N}_{\alpha\beta}$ by

$$\tilde{N}_{\alpha\beta}(0, t] = \int_{\{0, t\}} \sum_{i \in \alpha} \left(1_{\{X(s-)=i\}} \left(\sum_{j \in \beta} dN_{ij}(s) \right) \right) + \int_{\{s, t\}} 1_{\{\tilde{X}(s-)\neq\alpha\}} d\hat{N}_{\alpha, \beta}(s),$$

where the (dummy) point processes $\{\hat{N}_{\alpha\beta}\}_{\substack{\alpha, \beta \in \tilde{E} \\ \alpha \neq \beta}}$ form an independent family of HPPs with intensities $\{\tilde{q}_{\alpha\beta}\}_{\substack{\alpha, \beta \in \tilde{E} \\ \alpha \neq \beta}}$, respectively, and are independent of $X(s)$ and $\{N_{ij}\}_{i, j \in E: i \neq j}$. Then

$$f(\tilde{X}(t)) = f(\tilde{X}(s)) + \sum_{\substack{\alpha, \beta \in \tilde{E} \\ \alpha \neq \beta}} (f(\beta) - f(\alpha)) \int_{\{s, t\}} 1_{\{\tilde{X}(u-)=\alpha\}} d\tilde{N}_{\alpha\beta}(u). \tag{2.18}$$

Suppose (proof comes later) that $\{\tilde{N}_{\alpha, \beta}\}_{\substack{\alpha, \beta \in \tilde{E} \\ \alpha \neq \beta}}$ is a family of independent HPPs with respective intensities $\{\tilde{q}_{\alpha\beta}\}_{\substack{\alpha, \beta \in \tilde{E} \\ \alpha \neq \beta}}$. It follows from (2.18) that $\{\tilde{X}(t)\}_{t \geq 0}$ is a homogeneous Markov chain, with transition matrix $\tilde{\mathbf{A}}$.

To prove the assumed result concerning $\{\tilde{N}_{\alpha\beta}\}$, we apply Watanabe's theorem with the family \underline{Y} reduced to the single random variable $X(0)$. Take $Z(t)$ to be as specified in Theorem

2.1, where the family $\{N_i\}$ is now $\{\tilde{N}_{\alpha\beta}\}$. We obtain

$$E \left[\int_{(0,T]} Z(t) d\tilde{N}_{\alpha\beta}(t) \right] = \sum_{i \in \alpha} \sum_{j \in \beta} E \left[\int_{(0,T]} Z(t) 1_{\{X(t-)=i\}} dN_{ij}(t) \right] + E \left[\int_{(0,T]} Z(t) 1_{\{\tilde{X}(t-) \neq \alpha\}} d\hat{N}_{\alpha\beta} \right],$$

and this quantity is equal, by the smoothing formula, to

$$\begin{aligned} & \sum_{i \in \alpha} \sum_{j \in \beta} E \left[\int_0^T Z(t) 1_{\{X(t-)=i\}} q_{ij} dt \right] + E \left[\int_0^T Z(t) 1_{\{\tilde{X}(t-) \neq \alpha\}} \tilde{q}_{\alpha\beta} dt \right] \\ &= E \left[\int_0^T Z(t) \left[\left(\sum_{j \in \beta} q_{ij} \right) 1_{\{\tilde{X}(t-)=\alpha\}} + \tilde{q}_{\alpha\beta} 1_{\{\tilde{X}(t-) \neq \alpha\}} \right] dt \right] \\ &= E \left[\int_0^T Z(t) \tilde{q}_{\alpha\beta} dt \right], \end{aligned}$$

where we used (2.17). □

3 Poisson Systems

3.1 The Purely Poissonian Description

Most continuous-time Markov chains arising in operations research receive a natural description in terms of a mapping transforming an *input* into an output called the *state process*. In *Poisson systems*, the input consists of HPPs, also called the *driving HPPs*.

The construction of the state process from the basic driving HPPs is similar to the one used in Section 1 for jump HMCs, only with more details corresponding to a “natural” description of the system considered.

The rationale behind this construction is that there are sources of events that can be either active or inactive depending on the state of the process. Active sources then compete to produce an event that will be the cause for a transition.

We shall now proceed with the precise definitions.

Let S and E be two countable sets, the set of *event sources* (*sources* for short) and the set of *states*, respectively. With each source–state pair (s, i) is associated a number $c(s, i) \geq 0$ measuring the *intensity of activity* of source s when the state is i . An event source s is said to be *active in state* i if $c(s, i) > 0$. We denote by $A(i)$ the collection of sources s active in state i . With each source–state pair (s, i) is associated a probability distribution $p(s, i, \cdot)$ on E .

The four objects $E, S, p(\cdot, \cdot, \cdot)$ and $c(\cdot, \cdot)$ constitute the *transition mechanism*.

The collection T of triples (s, i, j) such that $c(s, i)p(s, i, j) > 0$ is called the set of *allowed transitions*. If $(s, i, j) \in T$, one says that a transition $i \rightarrow j$ can be triggered by an event produced by source s .

With each source $s \in S$ is associated an intensity, or *rate*, $\lambda_s > 0$. Let $\{N_{s,i,j}\}_{(s,i,j) \in T}$ be a family of independent HPPs with the respective intensities $\{a_{s,i,j}\}_{(s,i,j) \in T}$, where

$$a_{s,i,j} = \lambda_s c(s, i)p(s, i, j),$$

and define for all pairs of states (i, j) ,

$$a_{ij} = \sum_{s \in S} \lambda_s c(s, i)p(s, i, j), \quad (3.1)$$

and

$$a_i = \sum_{j \in E} a_{ij}. \quad (3.2)$$

We introduce the *stability assumption*: For all $i \in E$,

$$a_i < \infty. \quad (3.3)$$

We first give the purely Poissonian description of the state of the system, which is in general not directly readable from the natural description of the system one wishes to model, but which has the advantage of immediately leading to the infinitesimal generator.

The *state process* $\{X(t)\}_{t \geq 0}$ with values in $E_\Delta = E \cup \{\Delta\}$, where Δ is an element outside E , is constructed as follows.

The initial state $X(0)$ is a random variable independent of the HPPs $\{N_{s,i,j}\}_{(s,i,j) \in T}$. Suppose the state at time t is $X(t) = i \in E$. The active sources are then all the sources $s \in A(i)$. The HPPs competing for the determination of the next transition after t are $\{N_{s,i,j}\}_{s \in A(i), p(s,i,j) > 0}$. If N_{s_0,i,j_0} is the HPP producing the first event strictly after t , say at time $t + \tilde{\tau}$, then a transition from state i to state j_0 occurs at time $t + \tilde{\tau}$. The stability condition $a_i < \infty$ ensures that $\tilde{\tau} > 0$, since a_i is the sum of the intensities of the $N_{s,i,j}$ in competition when the state is i (Theorem 1.3, Chapter 8).

This recursive procedure started from time 0 gives a sequence of transition times $\tilde{\tau}_1, \tilde{\tau}_2, \dots$. Possibly, there is a random index N such that $\tilde{\tau}_N = \infty$, in which case we let $\tilde{\tau}_k = \infty$ for $k \geq N$. The embedded process is $\tilde{X}_1, \tilde{X}_2, \dots$, where $\tilde{X}_n = X(\tilde{\tau}_n)$, with the convention $X(\infty) = \Delta$. This procedure defines $X(t)$ for all $t \in [0, \tilde{\tau}_\infty)$, where $\tilde{\tau}_\infty = \lim_{n \uparrow \infty} \tilde{\tau}_n$ is called the *explosion time*. For $t \geq \tilde{\tau}_\infty$, we set $X(t) = \Delta$.

The same arguments as in Section 1 show that $\{X(t)\}_{t \geq 0}$ is a strong Markov jump HMC with state space E_Δ . From now on, we shall assume that $\tilde{\tau}_\infty = \infty$, and therefore that $\{X(t)\}_{t \geq 0}$ is a regular jump HMC on E . (In all the queuing networks that we shall consider, this assumption holds true, with an easy proof.)

Infinitesimal Generator of a Poisson System

We shall now look in detail at the first transition at time $\tilde{\tau}_1$, which is typical of all other transitions. Suppose that $X(0) = i$. The competing processes are the $N_{s,i,j}$ where $s \in A(i)$ and j satisfies $p(s, i, j) > 0$. If one is not interested in the particular source triggering the transition, one can lump all the $N_{s,i,j}$, $s \in A(i)$, to form a single HPP N_{ij} of intensity a_{ij} . We are then in the same situation as in Section 1, with one exception: a_{ii} may be strictly positive. This motivates the notations $\tilde{\tau}_n$ instead of τ_n , \tilde{X}_n instead of X_n , a_{ij} instead of q_{ij} , etc., which are there to warn us that a *pseudo-transition* can occur, i.e., a transition from i to i . Otherwise, the situation is similar to that of Section 1, and we can state the following result, defining $\tilde{X}_0 = X(0)$:

(α) $\{\tilde{X}_n\}_{n \geq 0}$ is a discrete-time HMC with values in $E_\Delta = E \cup \{\Delta\}$ and transition matrix $\{\tilde{p}_{ij}\}_{i,j \in E_\Delta}$ given by $\tilde{p}_{\Delta\Delta} = 1$, $\tilde{p}_{i\Delta} = 1$ if $a_i = 0$, and if $i \in E$, $a_i > 0$, $j \in E$,

$$\tilde{p}_{ij} = \frac{a_{ij}}{a_i}. \tag{3.4}$$

(β) Given $\{\tilde{X}_n\}_{n \geq 0}$, the sequence $\{\tilde{\tau}_{n+1} - \tilde{\tau}_n\}_{n \geq 0}$ is independent, and for all $n \geq 0$, $i, i_0, i_1, \dots, i_{n-1} \in E$, $x \in \mathbb{R}_+$,

$$P(\tilde{\tau}_{n+1} - \tilde{\tau}_n \geq x \mid \tilde{X}_n = i, \tilde{X}_{n-1} = i_{n-1}, \dots, \tilde{X}_0 = i_0) = e^{-a_i x}. \tag{3.5}$$

So far there is nothing new with respect to the construction of Section 1 except that we allow pseudo-transitions which can play a role in some situations, in feedback queues, for instance (see Example 4.4 below). Let us call $\{\tau_n\}_{n \geq 0}$ and $\{X_n\}_{n \geq 0}$ the sequence of *true* transitions of the state process (that is, from some state to a different state), and $\{X_n\}_{n \geq 0}$ the embedded process at the true transition times. If we are interested only in true transitions, it suffices to get rid of all the HPPs $N_{s,i,i}$ because they do not affect the state process. We therefore retrieve the situation of Section 1 with

$$q_{ij} = \sum_{s \in S} \lambda_s c(s, i) p(s, i, j), \tag{3.6}$$

where $j \neq i$, and

$$q_i = \sum_{j \in E, j \neq i} \sum_{s \in S} \lambda_s c(s, i) p(s, i, j). \tag{3.7}$$

3.2 The GSMP construction

The abbreviation GSMP stands for generalized semi-Markov process. A GSMP is a generalization of a Poisson system, and uses the same mechanism

$$(E, S, c(\cdot, \cdot), p(\cdot, \cdot, \cdot)).$$

The basic feature is that each source $s \in S$ can generate as many independent random variables as needed with the same cumulative distribution function $G_s(x)$. In addition, all the random variables generated by the various sources are independent.

At each transition time $\tilde{\tau}_n$, the situation of the system is described by the state

$$X(\tilde{\tau}_n) = \tilde{X}_n$$

and a collection of *consumable* random variables

$$(Y_s(n), s \in S).$$

The transition time $\tilde{\tau}_{n+1}$ is given by

$$\tilde{\tau}_{n+1} - \tilde{\tau}_n = \inf_{s \in A(\tilde{X}_n)} \frac{Y_s(n)}{c(s, \tilde{X}_n)}. \tag{3.8}$$

Denoting by \tilde{s}_n the source $s \in A(\tilde{X}_n)$ realizing the infimum, the next state \tilde{X}_{n+1} is chosen according to

$$P(\tilde{X}_{n+1} = j \mid \tilde{X}_n = i, \tilde{s}_n = s) = p(s, i, j), \tag{3.9}$$

and \tilde{X}_{n+1} is independent of $\tilde{X}_0, \tilde{s}_0, \dots, \tilde{X}_{n-1}, \tilde{s}_{n-1}$ given \tilde{X}_n and \tilde{s}_n .

At the origin of times $\tilde{\tau}_0 = 0$, the set $(Y_s(0), s \in S)$ is an independent collection of random variables with the respective cumulative distribution function $(G_s(x), s \in S)$. The set of consumable random variables $(Y_s(n+1), s \in S)$ available at transition time $\tilde{\tau}_{n+1}$ is given by

$$\begin{aligned} Y_s(n+1) &= Y_s(n) \text{ if } s \notin A(\tilde{X}_n), \\ Y_s(n+1) &= Y_s(n) - c(s, \tilde{X}_n)(\tilde{\tau}_{n+1} - \tilde{\tau}_n) \text{ if } s \in A(\tilde{X}_n), s \neq \tilde{s}_n, \end{aligned}$$

and for $s = \tilde{s}_n$, the consumed random variables $Y_s(n)$ is replaced by a new random variable $Y_s(n+1)$, with the c.d.f $G_s(x)$ and independent of $(Y_s(k), k \leq n, s \in S)$.

This constructive algorithm is easily implemented in software. An alternative way to visualize the construction uses *markers*. With each source $s \in S$ is associated a point process $N_s = \{T_n^s\}_{n \geq 0}$, where $T_0^s = 0$, and for $n \geq 0, T_{n+1}^s - T_n^s = S_{n+1}^s$, where $\{S_n^s\}_{n \geq 1}$ is an i.i.d sequence with cumulative distribution function $G_s(x)$. The processes $\{N_s\}_{s \in S}$ are mutually independent. Each point process has a marker, which is located at time t at position

$$M_s(t) = \int_0^t c(s, X(u)) du.$$

In particular, the marker on N_s moves at speed $c(s, i)$ when the GSMP is in state i . At time $\tilde{\tau}_n$, the set of consumable $(Y_s(n+1), s \in S)$ is given by the forward recurrence times (see Figure 9.4.1)

$$Y_s(n+1) = \inf \{T_k^s; T_k^s > M_s(\tau_n)\} - M_s(\tau_n). \tag{3.10}$$




Figure 9.4.1. Markers and consumable variables

Only the markers on the point processes N_s such that $s \in A(\tilde{X}_n)$ are moving, at the respective speeds $c(s, \tilde{X}_n)$, and the first to meet the next event of its point process provokes a transition. The time of transition $\tilde{\tau}_{n+1}$ and the new state \tilde{X}_{n+1} are given by (3.8) and (3.9), respectively, where \tilde{s}_n is the source corresponding to the marker that first reached an event of its point process.

Poisson Systems as GSMPS

When $G_s(x) = 1 - e^{-\lambda_s x}$ for all $s \in S$, the GSMPS construction gives the same state process as the purely Poissonian description. Indeed, in this construction, at time $\tilde{\tau}_n$, there is a competition among the

$$\{S_{\tilde{\tau}_n} N_{(s,i,j)}\}_{(s,j) \in T(i)},$$

where $T(i) = \{(s, j) : (s, i, j) \in T\}$. By the strong Markov property for HPPs, these HPPs are independent and distributed as

$$\{N_{(s,i,j)}\}_{(s,j) \in T(i)}, \tag{3.11}$$

and moreover, they are independent of what happened before time $\tilde{\tau}_n$. We can therefore consider that $\tilde{\tau}_n = \tilde{\tau}_0 = 0$ without loss of generality.

The first transition time $\tilde{\tau}_1$ is obtained by competition among the HPPs in (3.11), where

we suppose that $\tilde{X}_0 = i$. This competition can be organized in two stages. First there is a competition among the

$$\{N_{(s,i)}\}_{s \in A(i)},$$

where

$$N_{(s,i)} = \sum_{j; (s,i,j) \in T} N_{(s,i,j)},$$

which produces the transition time

$$\tilde{\tau}_1 = \inf_{s \in A(i)} Z_{(s,i)}(0),$$

where $Z_{(s,i)}(0)$ is the first point of $N_{(s,i)}$, an exponential random variable with mean $(\lambda_s c(s,i))^{-1}$. One can write

$$Z_{(s,i)}(0) = \frac{Y_s(0)}{c(s,i)},$$

where $Y_s(0)$ is an exponential random variable with mean λ_s^{-1} , and therefore

$$\tilde{\tau}_1 - \tilde{\tau}_0 = \inf_{s \in A(\tilde{X}_0)} \frac{Y_s(0)}{c(s, \tilde{X}_0)},$$

and we retrieve (3.8). Let \tilde{s}_0 be the source giving the infimum. The new state \tilde{X}_1 is chosen by competition among the $\{N_{\tilde{s}_0,i,j}\}_{j \in E, p(\tilde{s}_0,i,j) > 0}$. Therefore,

$$P(\tilde{X}_1 = j \mid \tilde{X}_0 = i, \tilde{s}_0 = s) = p(s, i, j),$$

and we retrieve (3.9).

3.3 Markovian Queues as Poisson Systems

Example 3.1. *The M/M/1/∞/FIFO Queue*

One of the simplest examples of a Poisson system is the M/M/1/∞ queue, a model of a waiting line. Consider a ticket booth with a single attendant, or *server*, in it. Customers wait in line in front of the booth, and the facility is so large that no bound is imposed on the number of customers waiting for service. In other words, the *waiting room has infinite capacity*. Such a system will be called a *1/∞ service system*, where 1 is for the number of servers, and ∞ is for the capacity of the waiting room.

Customer arrivals are modeled by a homogeneous Poisson process $\{T_n\}_{n \geq 1}$ of intensity $\lambda > 0$. Customer n arriving at time T_n brings a service request σ_n , which means that the server will need σ_n units of time for processing the request of customer n . The sequence $\{\sigma_n\}_{n \geq 1}$ is assumed i.i.d, with exponential distribution of mean μ^{-1} . Also, the arrival sequence $\{T_n\}_{n \geq 1}$ and the service sequence $\{\sigma_n\}_{n \geq 1}$ are supposed independent. Such a pattern of arrivals is called an *M/M input*. In this notation, introduced by D. Kendall, M means “Markovian”. Indeed, the Poisson process is Markovian, and exponential distributions are intimately connected with the Markov property.

The server attends one customer at a time and does not remain idle as long as there is at least one customer in the *system* (ticket booth plus waiting room). Once the service of a customer is started it cannot be interrupted before completion.

The above system is called an M/M/1/∞ queue. Its description could be complemented by an indication of the *service discipline* used: for instance, FIFO (first in first out), where the server, after completion of a service, chooses his next customer at the head of the line. However, we shall see that the service discipline turns out to be irrelevant if one is interested in the congestion process counting the number of customers present in the system.

We shall now describe the M/M/1/∞/FIFO queue as a Poisson system.

For this, we take $S = \{\alpha, \delta\}$, where α stands for *arrival* and δ for *departure*. The state space is $E = \mathbb{N} = \{0, 1, 2, \dots\}$, $i \in E$ representing the number of customers present in the system (in the waiting line or being attended by the server). If $i > 0$, $A(i) = \{\alpha, \delta\}$, meaning that when the number of customers in the system is strictly positive, an arrival or a departure could occur. If $i = 0$, $A(0) = \{\alpha\}$, meaning that if the system is empty, one cannot expect a departure, only an arrival.

All the activity intensities that are not null are taken equal to 1. Also, $\lambda_\alpha = \lambda$ and $\lambda_\delta = \mu$, because we want the interarrival times to be exponential random variables with mean λ^{-1} and the service requests to be exponential with mean μ^{-1} .

If in state $i \geq 0$ a transition is triggered by an event of source $s_0 = \alpha \in A(i)$, the next state is $j = i + 1$ (an arrival increases the number of customers by 1); whereas in state $i > 0$, if $s_0 = \delta \in A(i)$, the next state is $j = i - 1$. Thus, $p(\alpha, i, i + 1) = 1$ for all $i \geq 0$, and $p(\delta, i, i - 1) = 1$ for all $i > 0$, and all other probabilities $p(s, i, j)$ are null.

Figure 9.4.2 depicts a typical evolution of the state process, starting from state $i_0 = 0$.




Figure 9.4.2. GSMP construction of M/M/1/∞/FIFO

Formula (3.6) makes the infinitesimal generator readable from the natural description in terms of the rates λ_s , the speeds $c(s, i)$, and the transition probabilities $p(s, i, j)$. For the above M/M/1/ ∞ /FIFO queue, it can be readily checked that

$$q_{i,i+1} = \lambda, \quad q_{i,i-1} = \mu \mathbf{1}_{\{i \geq 0\}}. \quad \diamond$$

Remark 3.1.

One should point out that even for the simplest example—the M/M/1/ ∞ queue—the Poisson system description differs from the usual regenerative description of the corresponding Markov chain. The Poisson system description gives all the details about the mechanism of generation of transitions, and its collapse into the usual regenerative description is accompanied by a loss of information concerning the fine details.

Indeed, one cannot distinguish statistically between the state process of an M/M/1/ ∞ /FIFO queue that has an infinitesimal generator of the form

$$\mathbf{A} = (\lambda + \mu)(\mathbf{K} - I),$$

where

$$\mathbf{K} = \begin{pmatrix} 0 & 1 & 0 & 0 & \cdots \\ \frac{\mu}{\lambda + \mu} & 0 & \frac{\lambda}{\lambda + \mu} & 0 & \cdots \\ 0 & \frac{\mu}{\lambda + \mu} & 0 & \frac{\lambda}{\lambda + \mu} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix},$$

from the uniform Markov chain constructed from a Poisson process of intensity $\lambda + \mu$ and a random walk with reflecting barrier 0 with the transition matrix \mathbf{K} . It is actually not easy to discover from the uniform description that one is dealing with an M/M/1/ ∞ queue. This difficulty increases with the size and the complexity of the system; see Section 5 on queuing networks. \diamond

Example 3.2. *M/M/K/0 (The Erlang Queue)*

This queueing system receives the same description as the M/M/1/ ∞ /FIFO queue except for the following features. There are $K \geq 1$ servers, and the capacity of the waiting room is 0. In particular, an arriving customer finding the K servers busy is not accepted in the system, and therefore the number of customers present in the system at any given time is less than or equal to K .

A customer entering the system (i.e., arriving and seeing one or more servers idle) will select a free server at random.

One possible transition mechanism is the following:

$$\begin{aligned} S &= \{0, 1, \dots, K\}, \\ E &= \{\text{all subsets of } S \text{ containing } 0\}, \\ A(i) &= \{\text{all event sources in } i\}, \\ c(s, i) &= 1 \text{ for all } s \in A(i). \end{aligned}$$

The transition probabilities $p(\cdot, \cdot, \cdot)$ will be described in a few lines. The Poisson processes corresponding to event sources $1, 2, \dots, K$ all have the same intensity $\mu > 0$, whereas the Poisson process of event source 0 has the intensity $\lambda > 0$. Source 0 corresponds, as in the previous example, to the arrivals.

For any state $i \in E$, $i - \{0\}$ is a subset of $\{1, \dots, K\}$, representing the servers that are busy when the system is in that state. If $X(t) = i$, the number of customers in the system is $|i| - 1$, where $|i|$ is the cardinal of the set i . An event source $s \in [1, K]$ corresponds to server s , and the sequence $S_n^s = T_n^s - T_{n-1}^s (n \geq 1)$ is the sequence of successive service times provided by this server.

We shall now give the transition probabilities $p(\cdot, \cdot, \cdot)$. For this we need to distinguish two states: $i_0 = \{0\}$ and $i_K = \{0, 1, \dots, K\}$ corresponding respectively to an empty and a full system.

If $i \neq i_0$ and a transition is triggered on event source $s \in i$ where $s > 0$, this means that server s releases a customer, and the next state is then $i - \{s\}$.

If $i \neq i_K$ and a transition is triggered on event source 0 (which means that a new customer arrives), the next state is $i + \{r\}$, where r is chosen at random in $S - i$, the set of idle servers in state i .

If $i = i_K$ and a transition is triggered on event source 0 (a new customer arrives), the state does not change (the new customer is not accepted, all servers being busy).

The M/M/K/0 model is also called by the name of its inventor, the Danish telephone engineer K. Erlang, who used it to study quantitatively a *telephone switch* with K lines, or *channels*. Each customer finding a free line is connected and uses it for the time of a conversation. A customer finding all channels busy is rejected, or in the best case routed to another switch. Erlang was able to obtain his famous *blocking formula* giving the probability in stationary regime that a given customer finds all lines busy. This is the first formula of queueing theory; see Example 5.2 below.

The infinitesimal generator can be computed from formula (3.6) and (exercise) this gives for the nonnull terms

$$q_{i,i+\{r\}} = \lambda \frac{1}{K + 1 - |i|} \quad \text{if } i \in E, r \in \{1, \dots, K\}, r \notin i,$$

$$q_{i,i-\{r\}} = \mu \quad \text{if } i \in E, r \in \{1, \dots, K\}, r \in i.$$

Defining $Q(t) = |X(t)| - 1$ (the cardinal of $X(t)$ minus 1, that is, the number of busy servers), we can apply Theorem 3.3 to prove that this is a regular jump HMC with state space $\tilde{E} = \{0, 1, \dots, K\}$ and infinitesimal generator \tilde{A} given by

$$\tilde{q}_{n,n+1} = \lambda \text{ if } n \in [0, K - 1]$$

$$\tilde{q}_{n,n-1} = n\mu \text{ if } n \in [1, K].$$

A given state $n \in \tilde{E}$ is obtained by grouping the states $i \in E$ such that $|i| - 1 = n$.

To show that $\tilde{q}_{n,n-1} = \mu n$ for $n \in [1, K]$, we must consider the transitions of $\{X(t)\}$ from i such that $|i| - 1 = n$ to j such that $|j| - 1 = n - 1$. Fixing i such that $|i| - 1 = n$, there are exactly n states j such that $q_{ij} > 0$ and $|j| - 1 = n - 1$, namely all states j of the form $j = i - \{r\}$, where $r \in \{1, \dots, K\}$ and $r \in i$. The corresponding sum $\sum_j q_{ij} = n\mu$ and is independent of i such that $|i| - 1 = n$. Thus condition (2.17) of Theorem 3.3 is satisfied for $\alpha = n, \beta = n - 1$. The proof of $\tilde{q}_{n,n+1} = \lambda$ is similar. \diamond

Example 3.3. *M/M/1/∞/LIFO Preemptive Resume.*

This queuing system receives a description similar to that of a M/M/1/∞/FIFO, except for the FIFO discipline which now becomes LIFO *preemptive resume* (the abbreviation LIFO stands for last in first out). A customer upon arrival goes right to the ticket booth, and the customer who was receiving service is sent back to the waiting room, where he will stand in front of the line (at least until the time when another rude customer shows up, sending the first rude customer to the front of the queue, and so on). This type of discipline is called *preemptive*. The phrase *preemptive resume* means that a preempted customer does not have to start from scratch: When the server sees him next time, he will resume work where it was left.

This type of discipline is not as unfair as it may appear. First of all, all customers being equally rude, they endure as much as they hurt. A customer with a large service request spends a larger time at the ticket booth and is therefore more exposed than a customer with a modest request, and it is precisely in this sense that the discipline is fair. It makes longer requests who are responsible for congestion wait longer in the system. The precise result is that the expected sojourn time of a customer in the system, given that its service request is x , is equal to $x/(1 - \rho)$, where $\rho = \lambda/\mu$; see (Wolff, 1989).

The construction of M/M/1/∞/FIFO *preemptive resume* as a Poisson system is as follows

$$\begin{aligned} S &= \mathbb{N}, \\ E &= \{\text{all finite subsets of } S \text{ containing } 0\}, \\ A(i) &= \{0, i_{n(i)}\}, \text{ where } i = \{0, i_1, i_2, \dots, i_{n(i)}\} \in E, \\ c(s, i) &= 1 \text{ if } s \in A(i), \\ p(0, i, i + \{i_{n(i)} + 1\}) &= 1, \quad p(i_{n(i)}, i, i - \{i_{n(i)}\}) = 1. \end{aligned}$$

Source 0 is the arrival source. The last prescription is that when a new customer shows up (that is, a transition is triggered on source 0), the state being i , a new source is added to i to form the next state, the one just above the set i , and this source becomes immediately active. When source $i_{n(i)}$ triggers an event, this means a departure from the queue, and the source $i_{n(i)}$ disappears from state i , and the source $i_{n(i)-1}$ is reactivated. The HPP $\{T_n^0\}_{n \geq 1}$ has intensity $\lambda > 0$, and for all $k \geq 1, \{T_n^k\}_{n \geq 1}$ has intensity $\mu > 0$. When $k \geq 1$, the interevent times in $\{T_n^k\}_{n \geq 1}$ correspond to service times.

We see that in state i , there are $n(i)$ customers in the system. Thus the congestion process at time t is $|X(t)| - 1$, where $|X(t)|$ is the cardinal of the set $X(t)$. \diamond

Example 4.1 was chosen for its simplicity. Example 4.2 shows that the state space E should sometimes be chosen larger than the space in which a process of interest, here the congestion process, takes its values; however, the latter should be a function of the state process, and this is why $i \in E$ is sometimes called a *macrostate*. Example 4.3 contains the same teachings as Example 4.2, plus the feature that all sources in the (macro) state i need not be active.

In the queuing literature, especially when applications to communications networks or computer networks are considered, a queuing system is represented by the pictogram of Figure 9.4.3, where the input arrow represents the arrival stream of *jobs* (customers), the output arrow represents the stream of *completed jobs* (served customers), the circle is a *processor* (the service system), and the stack is a *buffer* (a waiting room, where customers wait for a server to be free).




Figure 9.4.3. The holy pictogram of queuing theory!

This basic pictogram can be richly adorned. For instance, in Figure 9.4.3 we have a system with K servers, a waiting room of capacity c , λ is the arrival rate of customers, and the derivation shows that the customers finding all K servers busy and a full waiting room are rejected. Also, there is an indication of the average service time, $1/\mu$, and of the fact that the incoming stream is Poisson. In this pictogram, one would take it as implicit that the service times sequence is i.i.d and independent of the arrival process, unless other assumptions are explicitly mentioned. One sometimes also gives the service discipline (LIFO, FIFO, etc.).

Example 3.4. *M/M/1/∞/FIFO Queue with Instantaneous Feedback*

Although one can add and suppress pseudo-transitions at will without altering the state process, pseudo-transitions are not always meaningless. For instance, consider an $M/M/1/\infty$ queue with instantaneous feedback, where a customer finishing service either leaves the system with probability $1 - p$ or is immediately recycled with probability p at the end of the waiting line or at the service booth if there is an empty waiting line, with a new independent exponential service request (see Figure 9.4.4). Then the times of service completion of recycled customers do not correspond to a genuine transition of $\{X(t)\}_{t \geq 0}$, where $X(t)$ is the number of customers present in the system at time t . ◇




Figure 9.4.4. A queue with instantaneous feedback

4 Markovian Queuing Theory

4.1 Isolated Markovian Queues

This section constitutes a brief introduction to queuing theory, which finds applications in operations research and in the performance analysis of communications networks.

We have seen in Section 4 the M/M/1/∞ queue and the M/M/K/0 queue, for which the congestion processes $\{X(t)\}_{t \geq 0}$, where $X(t)$ is now the number of customers in the system at time t , are special cases of birth-and-death processes with an infinitesimal generator of the form

$$\mathbf{A} = \begin{pmatrix} -\lambda_0 & \lambda_0 & 0 & 0 & \cdots \\ \mu_1 & -(\lambda_1 + \mu_1) & \lambda_1 & 0 & \cdots \\ 0 & \mu_2 & -(\lambda_2 + \mu_2) & \lambda_2 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \quad (4.1)$$

where the state space is $E = \mathbb{N}$, or $E = \{0, 1, \dots, N\}$ for finite N . In any case, for all the queuing processes we are going to consider, $\lambda_i > 0$ for all $i \in E$ except $i = N$ when $E = \{0, 1, \dots, N\}$, and $\mu_i > 0$ for all $i \in E$ except $i = 0$. These conditions guarantee irreducibility.

The global balance equations are

$$\lambda_0 \pi(0) = \mu_1 \pi(1)$$

and

$$(\lambda_i + \mu_i) \pi(i) = \lambda_{i-1} \pi(i-1) + \mu_{i+1} \pi(i+1)$$

for $i \geq 1$, with the convention $\mu_{N+1} = 0$ if $E = \{0, 1, \dots, N\}$. In Example 3.4 of Chapter 8, a stationary distribution π was shown to exist if and only if

$$\sum_{i \in E} \prod_{n=1}^i \frac{\lambda_{n-1}}{\mu_n} < \infty, \quad (4.2)$$

in which case, for $i \geq 1$,

$$\pi(i) = \pi(0) \prod_{n=1}^i \frac{\lambda_{n-1}}{\mu_n} \quad (4.3)$$

and

$$\pi(0) = \left(1 + \sum_{i \in E} \prod_{n=1}^i \frac{\lambda_{n-1}}{\mu_n} \right)^{-1}. \quad (4.4)$$

The ergodicity condition (4.2) is, of course, automatically satisfied when the state space is finite.

Example 4.1. *M/M/1/∞*

In this case $E = \mathbb{N}$, $\lambda_i \equiv \lambda > 0$, and $\mu_i = \mu > 0$ for all $i \geq 1$. The ergodicity condition reads $\sum_{i \geq 1} \left(\frac{\lambda}{\mu}\right)^i < \infty$, that is,

$$\rho \stackrel{\text{def}}{=} \frac{\lambda}{\mu} < 1.$$

This condition is natural, since $\rho = \lambda E[\sigma_1]$, the *traffic intensity*, is the average rate of work entering the system per unit time, and should not exceed the maximal speed of service, equal to 1. The solution of the balance equation is

$$\pi(i) = (1 - \rho)\rho^i. \quad \diamond$$

Example 4.2. *M/M/K/0 (Erlang Loss System)*

Here $E = \{0, \dots, K\}$. The solution of the balance equations is

$$\pi(i) = \frac{\rho^i / i!}{\sum_{n=0}^K \rho^n / n!}$$

for $i \in [0, K]$. In particular,

$$\pi(K) = \frac{\rho^K / K!}{\sum_{n=0}^K \rho^n / n!}$$

is the *blocking probability*, the probability of finding the K channels busy. The corresponding formula is due to Erlang. The distribution π is called a Poisson distribution truncated at K , since

$$\pi(i) = P(Z = i \mid Z \leq K)$$

where Z is a Poisson random variable with mean ρ . ◇

Example 4.3. *M/M/∞/∞ or Pure Delay*

This is not really a queuing system, but a *pure delay* system. The arrival process, the service times sequence, and the waiting room are as in the M/M/1/∞ model, but now there is an infinity of servers, and therefore no queuing, since anyone entering the system finds an idle server. The state space is $E = \mathbb{N}$, and the parameters are $\lambda_i \equiv \lambda > 0$ and $\mu_i = i\mu$. The form of the birth parameter is due to the fact that when $i > 0$ customers are present in the system, they are all being served, and therefore there are i independent exponential random variables (the service times) of mean μ^{-1} , being consumed at speed 1. A transition

$i \rightarrow i - 1$ will take place as soon as one of them is consumed. If $X(t) = i$, the transition $i \rightarrow i - 1$ will therefore occur in the time interval $(t, t + h]$ with the probability $i\mu h + o(h)$.

The ergodicity condition (4.2) is $\sum_{i=1}^{\infty} \frac{\rho^i}{i!} < \infty$ and is always satisfied. The solution of the balance equations is the Poisson distribution

$$\pi(i) = e^{-\rho} \frac{\rho^i}{i!}. \quad \diamond$$

Example 4.4. *M/M/K/∞*

This is the M/M/1/∞ queue, except that there are now K servers. Here $\lambda_i \equiv \lambda > 0$ and $\mu_i = \inf(i, K)\mu$. Indeed, if there are $X(t) = i \leq K$ customers, there are i independent exponentials of mean μ^{-1} active in provoking a downward transition. Therefore, the probability that a downward transition occurs in $(t, t + h]$ is $i\mu h + o(h)$. If there are $X(t) = i > K$ customers, only K exponentials are active, since there are only K servers, and therefore a downward transition occurs in the interval $(t, t + h]$ with probability $K\mu h + o(h)$.

The ergodicity condition (4.2) is satisfied only if

$$\rho \stackrel{\text{def}}{=} \frac{\lambda}{\mu} < K.$$

Again, this just says that the average incoming work per unit time cannot exceed the maximal service speed, which is K when all servers are busy.

The stationary distribution is then, for $i \in [1, K]$,

$$\pi(i) = \pi(0) \frac{\rho^i}{i!},$$

and for $i \geq K$,

$$\pi(i) = \pi(0) \frac{\rho^K}{K!} \left(\frac{\rho}{K}\right)^{i-K},$$

where

$$\pi(0)^{-1} = \sum_{i=0}^{K-1} \frac{\rho^i}{i!} + \frac{\rho^K}{K!} \frac{1}{1 - \rho/K}.$$

In this system, the probability of waiting is the probability of entering the system when the K servers are busy, that is $\pi(\geq K) \stackrel{\text{def}}{=} \sum_{i \geq K} \pi(i)$. One obtains *Erlang's waiting formula*

$$\pi(\geq K) = \frac{\frac{\rho^K}{K!} \frac{1}{1 - \rho/K}}{\sum_{i=0}^{K-1} \frac{\rho^i}{i!} + \frac{\rho^K}{K!} \frac{1}{1 - \rho/K}}. \quad \diamond$$

Example 4.5. *M/M/1/∞ Processor Sharing*

The arrival process is of the M/M type, as in the M/M/1/∞ queue, and the waiting-room capacity is infinite. Just as a LIFO nonpreemptive discipline or a FIFO discipline does not

make a difference in the infinitesimal generator, the processor-sharing discipline does not either. In this discipline the server is equally shared among the customers in the system. More precisely, if at time t there are $X(t) = i$ customers in the system, each one is served at speed $1/i$. For a transition $i \rightarrow i - 1$, there are i independent exponentials of mean $1/\mu$, consumed at speed $1/i$. Therefore, the probability of such a transition in the time interval $(t, t + h]$ is $i \frac{\mu}{i} h + o(h) = \mu h + o(h)$. \diamond

The last example suggests that the statistics of the congestion process $\{X(t)\}_{t \geq 0}$ are independent of the service discipline in an M/M/1/ ∞ queue, and this is indeed the case if we consider only service disciplines such that the server works at full speed, equal to 1, whenever there is at least one customer in the system.

We shall now see a nice application of reversibility.

Example 4.6. *Burke–Reich Output Theorem*

This result is due to Burke, and the proof using a reversibility argument is due to Reich; see (Kelly, 1979) for bibliographical details. Let $\{X(t)\}_{t \geq 0}$ be a birth-and-death process on \mathbb{N} with birth parameters of the form

$$\lambda_i \equiv \lambda, \quad (4.5)$$

and suppose that $\mu_i > 0$ for all $i \geq 1$. The corresponding chain is irreducible, and we shall assume that it is ergodic, that is,

$$\sum_{i \geq 1} \frac{\lambda^i}{\mu_1 \cdots \mu_i} < \infty.$$

In this case, for $i \geq 1$,

$$\pi(i) = \pi(0) \frac{\lambda^i}{\mu_1 \cdots \mu_i},$$

where

$$\pi(0) = \left(1 + \sum_{i \geq 1} \frac{\lambda^i}{\mu_1 \cdots \mu_i} \right)^{-1}.$$

This chain is reversible, since $\pi(i + 1)q_{i+1,i} = \pi(i)q_{i,i+1}$, that is, $\pi(i + 1)\mu_{i+1} = \pi(i)\lambda$, for all $i \geq 0$. Note that the same would be true of any irreducible ergodic birth-and-death process, but the restriction (4.5) on the birth parameters has a nice consequence in queuing theory. Indeed, if we interpret the upward transitions as due to arriving customers and the downward transitions as due to departing customers, condition (4.5) says that the arrival process is an HPP (Problem 9.3.1).

Suppose that the queue is in steady state, and therefore the reversed process has the same distribution as the direct process. When time is reversed, the point process of departures becomes the point process of arrivals, and therefore, in view of the reversibility property, the reversed process of departures is an HPP. Now, the probabilistic nature of a HPP does not change when time is reversed. Therefore, the departure process is a Poisson process.

Also, since in direct time, for any time $t \geq 0$, the state $X(t)$ is independent of the future at time t of the arrival process (Markov property of Poisson processes), it follows from reversibility that $X(t)$ is independent of the past at time t of the departure process. \diamond




Figure 9.5.1. Reich's proof of Burke's theorem

4.2 The M/GI/1/∞/FIFO Queue

Markovian models are easy to handle, but they sometimes are not adequate. The Poissonian assumption for the input is very often justified, in particular when the law of rare events applies. However, the exponential assumption is usually unrealistic. Note that there are situations where the exponential assumption leads to the correct result even in the case where the service time distribution is arbitrary. The typical case is Erlang's system, in which the substitution of an arbitrary service distribution with the same mean as the original exponential distribution leaves unaltered the stationary distribution of the number of busy lines at an arbitrary time. In particular, Erlang's blocking formula remains valid, and this robustness explains the success of this formula, which was widely applied in the design of early telephone switches. This lucky phenomenon is called *insensitivity*, and is not so rare. For instance, the M/M/∞ queue and the M/M/1/∞/LIFO *preemptive* queue are insensitive. However, since insensitivity is not a general phenomenon, one must work a little more, and we shall consider in this subsection the M/GI/1/∞/FIFO queue, a non-Markovian queue for which the analysis is possible. This queue is exactly like an M/M/1/∞/FIFO queue, only with a general distribution for the i.i.d service times sequence,

$$P(\sigma_1 \leq x) = G(x). \quad (4.6)$$

Embedded Congestion Process

The corresponding congestion process $\{X(t)\}_{t \geq 0}$ is no longer Markovian. Fortunately, the process $\{X_n\}_{n \geq 0}$, where

$$X_n = X(\tau_n) \tag{4.7}$$

and τ_n is the n th departure time (see Figure 9.5.2), $\tau_0 = 0$, is a discrete-time HMC. This is a direct consequence of the extended strong Markov property of HPPs. The detailed proof is left for the reader (Problem 9.5.3). This discrete-time chain is called the *chain embedded at departure times*. Since the congestion process is taken as right-continuous, we see that X_n is the number of customers that customer n leaves behind him when he has completed service.




Figure 9.5.2. The embedded process at departure times

A glance at Figure 9.5.2 reveals that

$$X_{n+1} = (X_n - 1)^+ + Z_{n+1}, \tag{4.8}$$

where Z_{n+1} is the number of customers arriving in the interval $(\alpha_n, \alpha_n + \sigma_{n+1}]$, where $\alpha_n = \tau_n$ if $X_n > 0$ and $\alpha_n = T'_n$ if $X_n = 0$ (Fig. 9.5.2).

The time α_n is a stopping time with respect to the arrival process $N = N_\alpha$ and the first n service requests $(\sigma_1, \dots, \sigma_n)$, and the $(n + 1)$ st service request σ_{n+1} is independent of the arrival process before α_n and of the previous requests $(\sigma_1, \dots, \sigma_n)$. In particular, in view of Theorem 1.1,

$$Z_{n+1} = N(\alpha_n, \alpha_n + \sigma_{n+1}]$$

is independent of X_0, \dots, X_n , and $\{Z_n\}_{n \geq 1}$ is an i.i.d sequence, with each Z_n distributed as

$$\tilde{Z} = \tilde{N}(0, \tilde{\sigma}],$$

where \tilde{N} is some HPP with intensity λ , and $\tilde{\sigma}$ is a random variable independent of \tilde{N} , with cumulative distribution function $G(x)$.

We recognize an avatar of the fetish example “The Repair Shop” for which we have




Figure 9.5.3. Arrival process, FIFO workload, and congestion processes

already proven that $\{X_n\}_{n \geq 0}$ is indeed an HMC with transition matrix

$$\begin{pmatrix} a_0 & a_1 & a_2 & a_3 & \dots \\ a_0 & a_1 & a_2 & a_3 & \dots \\ 0 & a_0 & a_1 & a_2 & \dots \\ 0 & 0 & a_0 & a_1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix},$$

where $a_i = P(Z_{n+1} = i) = P(\tilde{Z} = i)$, that is,

$$a_i = \int_0^\infty e^{-\lambda t} \frac{(\lambda t)^i}{i!} dG(t). \tag{4.9}$$

For this particular instance of the repair shop HMC, the necessary and sufficient condition of irreducibility $a_0 > 0, a_0 + a_1 < 1$ is verified.

Recurrence of this chain depends on the value

$$\rho = E[Z_1] = E[\tilde{Z}] = \sum_{i=1}^\infty i \int_0^\infty e^{-\lambda t} \frac{(\lambda t)^i}{i!} dG(t).$$

A straightforward calculation gives

$$\rho = \lambda \int_0^{\infty} t dG(t) = \lambda E[\sigma]. \quad (4.10)$$

If $\rho < 1$, the chain is positive recurrent; if $\rho = 1$, it is null recurrent; and if $\rho > 1$, the chain is transient (the proof of this was completed in Example 3.1 of Chapter 5).

Again, this result is natural if we observe that ρ is the average work entering the system per unit of time and remember that the server works at a maximum speed of one unit of work completed per unit of time.

In Example 5.5 of Chapter 2 we found that in the positive recurrent case $\rho < 1$, the generating function of the stationary distribution π is given by the formula

$$\sum_{i=0}^{\infty} \pi(i) z^i = \pi(0) \frac{(z-1)g_Z(z)}{z - g_Z(z)}, \quad (4.11)$$

where

$$\pi(0) = 1 - \rho \quad (4.12)$$

and $g_Z(z)$ is the generating function of Z_n . The latter can be explicitly computed using expression (4.9):

$$g_Z(z) = \int_0^{\infty} e^{-\lambda t} \left(\sum_{i=0}^{\infty} \int_0^{\infty} \frac{(\lambda t)^i}{i!} z^i \right) dG(t),$$

that is

$$g_Z(z) = \int_0^{\infty} e^{-\lambda t(1-z)} dG(t). \quad (4.13)$$

Embedded Sojourn Process

In the $M/GI/1/\infty/\text{FIFO}$ queuing system, the number of customers $X_n = X(\tau_n)$ left behind by the n th customer when he leaves the system is exactly the number of customers arriving during the time interval $(T_n, \tau_n]$, that is,

$$X_n = N(T_n, T_n + V_n], \quad (4.14)$$

where V_n is the sojourn time of the n th customer in the system. Invoking the extended strong Markov property of HPPs and noting that V_n depends only on $(\sigma_1, \dots, \sigma_n)$ and the past of N at time T_n , it follows from (4.14) that at equilibrium (for $\{X_n\}_{n \geq 0}$),

$$\sum_{i=0}^{\infty} \pi(i) z^i = E \left[z^{\tilde{N}(V_n)} \right],$$

where \tilde{N} is a Poisson process of intensity λ independent of V_n . Now,

$$\begin{aligned} E \left[z^{\tilde{N}(V_n)} \right] &= \int_0^{\infty} E \left[z^{\tilde{N}(v)} \right] dF_{V_n}(v) \\ &= \int_0^{\infty} e^{\lambda v(z-1)} dF_{V_n}(v) = \Phi_{V_n}(\lambda(z-1)), \end{aligned}$$

where $F_{V_n}(v)$ is the c.d.f of V_n , and Φ_{V_n} is the Laplace transform of V_n . We therefore see that the distribution of V_n is independent of n , and that calling V any random variable with the same distribution as V_n , we have

$$\sum_{i=0}^{\infty} \pi(i)z^i = \Phi_V(\lambda(z-1)), \quad (4.15)$$

where

$$\Phi_V(s) = \int e^{\lambda s} dF_V(s). \quad (4.16)$$

Note that this expresses the fact that at equilibrium, X_n is distributed as

$$X = N(V),$$

where N is an HPP of intensity λ , and V is a random variable independent of N with the distribution of the stationary sojourn time.

4.3 The GI/M/1/∞/FIFO Queue

The GI/M/1/∞/FIFO queue is of the same nature as the M/M/1/∞/FIFO queue, except that now the arrival process is not Poissonian, but renewal. The interarrival times form an i.i.d sequence with cumulative distribution function $F(x)$ and mean λ^{-1} . The service times are exponential with mean μ^{-1} .

As in the M/GI/1/∞/FIFO queue, the congestion process $\{X(t)\}_{t \geq 0}$ is *not* Markovian, but the system is amenable to Markovian analysis. Indeed, if we define

$$X_n = X(T_n-),$$

the number of customers in the system seen upon arrival by the n th customer, the process $\{X_n\}_{n \geq 1}$ is an HMC. To see this, we can take the GSMP description of the queue analogous to the one given in Example 4.1, except that now the input process $\{T_n\}_{n \geq 0}$ corresponding to the source δ is a renewal process. In Figure 9.5.4, not all the crossed events in source δ are used as departure times, but only $X_n + 1$ of them at most. More precisely,

$$X_{n+1} = (X_n + 1 - N_\delta(T_n, T_{n+1}))^+.$$

Defining

$$Z_{n+1} = N_\delta(T_n, T_{n+1}],$$

the sequence $\{Z_n\}_{n \geq 1}$ is i.i.d. (use, for instance, the extended strong Markov property for HPPs), and therefore $\{X_n\}_{n \geq 0}$ is an HMC by Theorem 2.1 of Chapter 2. Moreover, writing

$$P(Z_{n+1} = k) = \int_0^\infty e^{-\mu t} \frac{(\mu t)^k}{k!} dF(t) \stackrel{\text{def}}{=} b_k,$$

the i th row of the transition matrix \mathbf{P} is

$$(1 - \sum_{k=0}^i b_k, b_i, b_{i-1}, \dots, b_0, 0, 0, \dots).$$




Figure 9.5.4. GSMP construction of GI/M/1/∞/FIFO

Since $b_k > 0$ for all $k \geq 0$, \mathbf{P} is irreducible and aperiodic. We first observe that

$$\sum_{k=0}^{\infty} kb_k = \rho^{-1}, \tag{4.17}$$

where

$$\rho = \frac{\lambda}{\mu} \tag{4.18}$$

is the traffic intensity. Indeed,

$$\begin{aligned} \sum_{k=0}^{\infty} kb_k &= E[Z_{n+1}] = E[N_{\delta}(T_n, T_{n+1})] \\ &= E\left[\int_0^{\infty} 1_{(T_n, T_{n+1})}(t) dN_{\delta}(t)\right] = E\left[\int_0^{\infty} 1_{(T_n, T_{n+1})}(t) \mu dt\right] \\ &= \mu E[T_{n+1} - T_n] = \frac{\lambda}{\mu}, \end{aligned}$$

where we have used the smoothing formula of Poisson calculus. We want to show that the embedded HMC $\{X_n\}_{n \geq 0}$ is

- positive recurrent if $\rho < 1$,

- null recurrent if $\rho = 1$,
- transient if $\rho > 1$.

A. First we show that if $\rho < 1$, the chain is positive recurrent. For this it suffices to prove the existence of a stationary distribution. We make the educated guess that π has the form

$$\pi(i) = \xi^i(1 - \xi) \quad (4.19)$$

for some $\xi \in (0, 1)$. To verify this guess, we must find $\xi \in (0, 1)$ such that for all $i \geq 1$,

$$\sum_{j=i-1}^{\infty} \xi^{j-i+1} b_{j-i+1} = \xi \quad (4.20)$$

and

$$\sum_{j=0}^{\infty} \left(\sum_{k=j+1}^{\infty} b_k \right) \xi^j = 1, \quad (4.21)$$

since these equations are the balance equations when π is given by (4.19).

Equations (4.20) all reduce to

$$\xi = g_Z(\xi), \quad (4.22)$$

where

$$g_Z(\xi) = \sum_{k=0}^{\infty} b_k \xi^k \quad (4.23)$$

is the generating function of Z_1 . Since all the b_k 's are positive, there is a unique solution ξ_0 of (4.22) in $(0, 1)$ if and only if

$$g'_Z(1) = \sum_{k=1}^{\infty} k b_k > 1,$$

that is, $\rho < 1$. We must verify (4.21). The left-hand side equals

$$\sum_{k=1}^{\infty} \sum_{j=0}^{k-1} b_k \xi^j = \sum_{k=1}^{\infty} b_k \left(\frac{1 - \xi^k}{1 - \xi} \right) = \frac{1}{1 - \xi} \left(1 - b_0 - \sum_{k=1}^{\infty} b_k \xi^k \right),$$

and in view of (4.22), this equals

$$\frac{1}{1 - \xi} (1 - b_0 - (\xi - b_0)) = 1.$$

B. Consider the matrix

$$\tilde{\mathbf{P}} = \begin{pmatrix} b_0 & b_1 & b_2 & \dots \\ b_0 & b_1 & b_2 & \dots \\ & b_0 & b_1 & \dots \\ & & b_0 & \dots \end{pmatrix}, \quad (4.24)$$

This is the transition matrix of the chain imbedded at the arrival times of an $M/GI/1/\infty$ system with traffic intensity

$$\tilde{\rho} = \frac{\mu}{\lambda} = \frac{1}{\rho}. \tag{4.25}$$

We know from Theorem 3.4, Chapter 5, that the irreducible HMC with transition matrix \mathbf{P} is transient if and only if there exists a nontrivial bounded solution $h : E \rightarrow \mathbb{R}$ of

$$h(j) = \sum_{k \neq 0} p_{jk} h(k), \quad j \neq 0, \tag{4.26}$$

where 0 is an arbitrary state. Equation (4.26) is

$$h(j) = \sum_{i=1}^{j+1} b_{j-i+1} h(i), \quad j \neq 0. \tag{4.27}$$

We can assume without loss of generality that $h(1) \geq 0$ (otherwise replace function h by $-h$). Defining

$$\tilde{\pi}(0) = h(1)b_0; \quad \tilde{\pi}(1) = h(1)(1 - b_0); \quad \tilde{\pi}(j) = h(j) - h(j - 1), \quad (j \geq 2), \tag{4.28}$$

it is easy to verify that

$$\tilde{\pi} \tilde{\mathbf{P}} = \tilde{\pi}. \tag{4.29}$$

We show that $\tilde{\pi} \geq 0$. Indeed, (4.29) is, for $j \geq 0$,

$$\tilde{\pi}(j) = \tilde{\pi}(0)b_j + \sum_{i=1}^{j+1} \tilde{\pi}(i)b_{j-i+1}.$$

Add these equations for $j = 0$ to n , and solve to obtain, for $n \geq 0$,

$$\tilde{\pi}(n + 1)b_0 = \tilde{\pi}(0)c_n + \sum_{i=1}^n \tilde{\pi}(i)c_{n-i+1},$$

where

$$c_n = 1 - b_0 - \dots - b_n > 0.$$

Therefore, if $\tilde{\pi}(i) \geq 0$ for $i = 0$ to n , then $\tilde{\pi}(n + 1) \geq 0$. But $\tilde{\pi}(0) \geq 0$. Therefore, by induction, $\tilde{\pi} \geq 0$.

Equation (4.29) has a nontrivial bounded solution with bounded sum $\sum_{i=0}^{\infty} \tilde{\pi}(i)$ if and only if $\tilde{\rho} < 1$. Equivalently, as we have just shown, (4.27) has a nontrivial bounded solution if and only if (4.29) has a nontrivial bounded solution with bounded sum. Therefore, by Theorem 3.4 of Chapter 5, \mathbf{P} is transient if and only if $\tilde{\mathbf{P}}$ is positive recurrent, i.e., if and only if $\rho > 1$.

C. It remains to show that \mathbf{P} is recurrent null if $\rho = 1$. From the previous discussion, \mathbf{P} cannot be transient. It suffices therefore to show that \mathbf{P} cannot be recurrent positive. Indeed, if it were, there would be a stationary distribution π for \mathbf{P} . Writing

$$\tilde{h}(j) = \pi(0) + \dots + \pi(j - 1), \quad j \geq 1, \tag{4.30}$$

and using $\sum_{k=0}^{\infty} b_k = 1$, we can verify that

$$\tilde{h}(j) = \sum_{k \neq 0} \tilde{\rho}_{jk} \tilde{h}(k), \quad j \neq 0. \quad (4.31)$$

Therefore, (4.31) has a nontrivial bounded solution, and by Theorem 3.4 of Chapter 5, this is equivalent to transience of $\tilde{\mathbf{P}}$. But we know that in this case necessarily $\tilde{\rho} > 1$, and therefore $\rho < 1$, a contradiction with our assumption $\rho = 1$. Therefore, for $\rho = 1$, \mathbf{P} is recurrent null.

Embedded Waiting Time Process

An arriving customer finds $X_n = X(T_n^-)$ customers in front of him, and therefore, in the FIFO discipline, his waiting time W_n (before he starts to be served) is the time needed for the server to take care of the X_n customers present at time T_n^- .

By the extended strong Markov property applied to the HPP N_δ in the GSMP construction (see Figure 9.5.4), the point process N_δ after T_n is an HPP of intensity μ and is independent of X_n . In particular,

$$W_n = \sum_{j=1}^{X_n} Y_j,$$

where the Y_j are i.i.d exponentials of mean μ^{-1} and common characteristic function

$$E[e^{iuY}] = \frac{\mu}{\mu - iu},$$

and are independent of X_n . Also,

$$P(X_n = k) = (1 - \xi_0)\xi_0^k.$$

Therefore, in steady state,

$$\begin{aligned} E[e^{iuW_n}] &= E\left[\sum_{k=0}^{\infty} e^{iu\sum_{j=1}^k Y_j} 1_{\{X_n=k\}}\right] \\ &= \sum_{k=0}^{\infty} E[e^{iuY}]^k P(X_n = k) \\ &= \sum_{k=0}^{\infty} \left(\frac{\mu}{\mu - iu}\right)^k (1 - \xi_0)\xi_0^k. \end{aligned}$$

We find for the characteristic function of the stationary waiting time

$$(1 - \xi_0) + \xi_0 \left(\frac{\mu(1 - \xi_0)}{\mu(1 - \xi_0) + iu} \right).$$

This is the characteristic function of a random variable that is null with probability $1 - \xi_0$ and exponential of mean $[\mu(1 - \xi_0)]^{-1}$ with probability ξ_0 .

For more information in classical Queuing Theory, the reader is directed, for instance, to Gross and Harris (1985).

4.4 Markovian Queuing Networks

Jackson Network

A Jackson network (Jackson, 1957) is an open network of interconnected queues. There are K stations, and each station has a $1/\infty$ service system, that is, a unique server working at unit speed and an infinite waiting room. There are two types of customers queuing at a given station, (1) those which are fed-back, that is, who have received service in another or the same station and are rerouted to the given station for more service, and (2) those who enter the network for the first time (see Figure 9.6.1).




Figure 9.6.1. Jackson network

The *exogenous arrivals* into station i form an HPP, denoted by \bar{N}_i , with intensity $\bar{\lambda}_i \in [0, \infty)$. The sequence of service times at station i are exponential random variables of mean $1/\mu_i \in (0, \infty)$.

The service times in the same and in different stations are independent and independent of the exogenous input HPPs \bar{N}_i , and the latter HPPs are independent of one another.

The routing is of the Bernoulli type. Each customer just completing service in station i tosses a $(K + 1)$ -faced die with probabilities $r_{i,1}, \dots, r_{i,K}, r_i$ with the effect that the customer is sent to station j with probability r_{ij} or leaves the system with probability $r_i = 1 - \sum_{j=1}^K r_{ij}$. The matrix

$$\mathbf{R} = \{r_{ij}\}_{1 \leq i, j \leq K}$$

is the *routing matrix*. The successive tosses of the routing dice of all stations are independent, and independent of the exogeneous arrival processes and of all the service times.

This is the original Jackson model, which can be enriched by the introduction of service speeds. If there are n_i customers in station i , the server works at speed $\phi_i(n_i)$, where $\phi_i(0) = 0$ and $\phi_i(n_i) > 0$ for all $n_i \geq 1$.

Let $X_i(t)$ be the number of customers in station i at time t , and define

$$X(t) = (X_1(t), \dots, X_K(t)).$$

It can be verified that the process $\{X(t)\}_{t \geq 0}$ is a regular jump HMC with state space $E = \mathbb{N}^K$ and with infinitesimal generator $\mathbf{A} = \{q_{n,n'}\}$, where all the nonzero off-diagonal terms are

$$\begin{aligned} q_{n,n+e_i} &= \bar{\lambda}_i, \\ q_{n,n-e_i} &= \mu_i \phi_i(n_i) r_i \mathbf{1}_{\{n_i > 0\}}, \\ q_{n,n-e_i+e_j} &= \mu_i \phi_i(n_i) r_{ij} \mathbf{1}_{\{n_i > 0\}}. \end{aligned}$$

where $n = (n_1, \dots, n_K) \in E = \mathbb{N}^K$, and e_i is the i th vector of the canonical basis of \mathbb{R}^K .

This generator is independent of the service strategy—LIFO, FIFO, or processor-sharing.

The form of the infinitesimal generator can be obtained from a full description of the network as a Poisson system. However, we shall be content with heuristic arguments. For instance, the expression for $q_{n,n-e_i+e_j}$ follows from the intuitive considerations below. If the state at time t is n , a transfer from station i to station j requires that the exponential random variable with mean $1/\mu_i$ representing the required service of the customer being served at station i at time t be terminated between times t and $t+h$ (probability $\mu_i \phi_i(n_i)h$ up to the first order in h), and that the corresponding customer be routed to station j (probability r_{ij}).

We shall assume that the chain $\{X(t)\}_{t \geq 0}$ is irreducible. This is the case when

(α) for all $j \in [1, K]$, there exist $i, i_1, \dots, i_m \in [1, K]$ such that

$$\bar{\lambda}_i r_{ii_1} r_{i_1 i_2} \cdots r_{i_m j} > 0,$$

and

(β) for all $j \in [1, K]$, there exist $j_1, j_2, \dots, j_\ell, k \in [1, K]$ such that

$$r_{jj_1} r_{j_1 j_2} \cdots r_{j_\ell k} r_k > 0$$

(the proof is left for the reader).

Recalling that the $\mu_i > 0$ for all stations, condition (α) tells us that any station is *exogenously supplied*, and (β) tells that any station has an *outlet*.

Consider now the $(K+1) \times (K+1)$ matrix

$$\tilde{\mathbf{R}} = \begin{pmatrix} r_{11} & r_{12} & \cdots & r_{1K} & r_1 \\ r_{21} & r_{22} & \cdots & r_{2K} & r_2 \\ \vdots & \vdots & & \vdots & \vdots \\ r_{K1} & r_{K2} & \cdots & r_{KK} & r_K \\ 0 & 0 & \cdots & 0 & 1 \end{pmatrix}.$$

It can be interpreted as the transition matrix of an HMC on the finite state space $\{1, \dots, K, K+1\}$. Condition (β) implies that state $K+1$ is absorbing and the states $1, \dots, K$ are transient.

It follows that $(1 - \mathbf{R})^{-1}$ exists and therefore the solution of the system of equations

$$\lambda_i = \bar{\lambda}_i + \sum_{j=1}^K \lambda_j r_{ji}, \quad (4.32)$$

that is, with obvious notations, $\lambda = \bar{\lambda} + \mathbf{R}\lambda$, has a unique solution

$$\lambda = (1 - \mathbf{R})^{-1} \bar{\lambda} = \left(\sum_{n=0}^{\infty} \mathbf{R}^n \right) \bar{\lambda},$$

where the latter expression shows that this solution is nonnegative. Equations (4.32) are called the *traffic equations*, because they give a necessary relation between the average numbers of customers λ_i entering station i in steady state if the network is ergodic. Indeed, λ_i is equal to the exogeneous rate of arrivals $\bar{\lambda}_i$ plus the sum of all average rates of transfer from other stations. From station j , the corresponding rate is $\alpha_j r_{ji}$, where α_j is the average rate of customers finishing service in station j . But at equilibrium $\alpha_j = \lambda_j$, since the average number of customers in station j remains constant, whence the traffic equations (4.32).

Of course, this heuristic argument is not needed, and in particular, one need not attempt to identify λ_i in (4.32) as the average incoming arrival rate in station i . This identification is possible if equilibrium is guaranteed. However, even in the nonergodic cases, the traffic equations have a solution, and a unique solution.

The existence of a probability distribution π on $E = \mathbb{N}^K$ satisfying Kolmogorov's balance equations is a necessary and sufficient condition of ergodicity of the network. We do the case where the service speeds of all the servers are equal to 1 and where $r_{ii} = 0$, and refer the reader to Problem 9.6.4 for the general case.

The global balance equations for the Jackson network are

$$\begin{aligned} \pi(n) \left\{ \sum_{i=1}^K (\lambda_i + \mu_i r_i 1_{\{n_i > 0\}}) \right\} &= \sum_{i=1}^K \pi(n - e_i) \bar{\lambda}_i 1_{\{n_i > 0\}} + \sum_{i=1}^K \pi(n + e_i) \mu_i r_i \\ &+ \sum_{i=1}^K \sum_{j=1}^K \pi(n + e_i - e_j) \mu_i r_{ij} 1_{\{n_j > 0\}}. \end{aligned}$$

It turns out that if the solution of the traffic equation satisfies

$$\rho_i = \frac{\lambda_i}{\mu_i} < 1 \quad (4.33)$$

for all $i \in [1, K]$, then the network is ergodic, and its stationary distribution is given by

$$\pi(n) = \prod_{i=1}^K \pi_i(n_i), \quad (4.34)$$

where π_i is the stationary distribution of an M/M/1/ ∞ queue with traffic intensity ρ_i ,

$$\pi_i(n_i) = \rho_i^{n_i} (1 - \rho_i). \quad (4.35)$$

To prove this, we shall apply the reversal test (Theorem 5.5 of Chapter 8). We define the generator $\tilde{\mathbf{A}}$ on $E = \mathbb{N}^K$ by

$$\pi(n)\tilde{q}_{n,n'} = \pi(n')q_{n',n}$$

and check that

$$\sum \tilde{q}_{n,n'} = q_n.$$

Here

$$q_n = \sum_{i=1}^K \lambda_i + \mu_i 1_{\{n_i > 0\}}.$$

The generator $\tilde{\mathbf{A}}$ is given by

$$\pi(n)\tilde{q}_{n,n+e_i} = \pi(n+e_i)q_{n+e_i,n} = \pi(n+e_i)\mu_i r_i,$$

$$\pi(n)\tilde{q}_{n,n-e_i} = \pi(n-e_i)q_{n-e_i,n} = \pi(n-e_i)\bar{\lambda}_i 1_{\{n_i > 0\}},$$

$$\pi(n)\tilde{q}_{n,n+e_i-e_j} = \pi(n+e_i-e_j)q_{n+e_i-e_j,n} = \pi(n+e_i-e_j)\mu_i r_{ij} 1_{\{n_j > 0\}},$$

and therefore, taking into account the specific form of $\pi(n)$ given by (4.34) and (4.35),

$$\tilde{q}_{n,n+e_i} = \rho_i \mu_i r_i = \lambda_i r_i,$$

$$\tilde{q}_{n,n-e_i} = \frac{\bar{\lambda}_i}{\rho_i} 1_{\{n_i > 0\}},$$

$$\tilde{q}_{n,n+e_i-e_j} = \frac{\rho_i}{\rho_j} \mu_i r_{ij} 1_{\{n_j > 0\}} = \frac{1}{\rho_j} \lambda_i r_{ij} 1_{\{n_j > 0\}}.$$

We have to verify that

$$\sum_{i=1}^K (\bar{\lambda}_i + \mu_i 1_{\{r_i > 0\}}) = \sum_{i=1}^K \left(\lambda_i r_i + \frac{\bar{\lambda}_i}{\rho_i} 1_{\{n_i > 0\}} + \frac{1}{\rho_i} \left(\sum_{j=1}^K \lambda_j r_{ji} \right) 1_{\{n_i > 0\}} \right).$$

By the traffic equation, $\sum_{j=1}^K \lambda_j r_{ji} = \lambda_i - \bar{\lambda}_i$, and therefore the right-hand side of the previous equality is

$$\sum_{i=1}^K \left(\lambda_i r_i + \frac{1}{\rho_i} \lambda_i 1_{\{n_i > 0\}} \right) = \sum_{i=1}^K (\lambda_i r_i + \mu_i 1_{\{n_i > 0\}}),$$

and it remains to check that

$$\sum_{i=1}^K \bar{\lambda}_i = \sum_{i=1}^K \lambda_i r_i.$$

For this we need only sum the traffic equations.

Gordon–Newell Network

A closed Jackson network, also called a *Gordon–Newell network* (Gordon and Newell, 1967), is one for which

$$\bar{\lambda}_i = 0, \quad r_i = 0 \tag{4.36}$$

for all $i, 1 \leq i \leq K$. In other words, there is no inlet and no outlet, and therefore the number of customers in the networks remains constant, and we shall call it N . The state space is

$$E = \{(n_1, \dots, n_K) \in \mathbb{N}^K, \sum_{i=1}^K n_i = N\}.$$

The traffic equations are now

$$\lambda = \mathbf{R}\lambda, \tag{4.37}$$

and since \mathbf{R} is a stochastic matrix in this case, which we shall assume irreducible, it has an infinity of solutions, all multiples of the same vector, which is the stationary distribution of \mathbf{R} .

It is true that the vector of average traffics through the stations is a solution of (4.37). We do not know which one, in contrast with the open network, where the solution of the traffic equation is unique. Nevertheless, if we take any positive solution, we see by inspection that for all $n \in E$,

$$\pi(n) = \frac{1}{G(N, K)} \prod_{i=1}^K \rho_i^{n_i}, \tag{4.38}$$

where $\rho_i = \lambda_i/\mu_i$ is a stationary solution. Here $G(N, K)$ is the normalizing factor

$$G(N, K) = \sum_{\substack{n \in \mathbb{N}^K \\ n_1 + \dots + n_K = N}} \prod_{i=1}^K \rho_i^{n_i}. \tag{4.39}$$

Note that under the irreducibility assumption for the routing matrix \mathbf{R} , the chain $\{X(t)\}_{t \geq 0}$ is itself irreducible, and since the state space is finite, it is positive recurrent, with a unique stationary distribution. In particular, for closed networks, there is no ergodicity condition.

The problem with closed Jackson networks resides in the computation of the normalizing constant $G(N, K)$. Brute-force summation via (4.39) is practically infeasible for large populations and/or large networks. Instead, we can use the following algorithm.

Define $G(j, \ell)$ to be the coefficient of z^j in the power series development of

$$g_\ell(z) = \prod_{i=1}^{\ell} \frac{1}{1 - \rho_i z} = \prod_{i=1}^{\ell} \left(\sum_{n_i=0}^{\infty} \rho_i^{n_i} z^{n_i} \right).$$

The normalizing factor is indeed equal to $G(N, K)$. Since

$$g_\ell(z) = g_{\ell-1}(z) + \rho_\ell z g_\ell(z),$$

we find the recurrence equation

$$G(j, \ell) = G(j, \ell - 1) + \rho_\ell G(j - 1, \ell) \tag{4.40}$$

with the initial conditions

$$G(j, 1) = \rho_1^j \quad (j \geq 0), \tag{4.41}$$

$$G(0, \ell) = 1 \quad (\ell \geq 1).$$

It is also of interest to be able to compute the *utilization* of server i , defined by

$$U_i(N, K) = P(X_i(t) > 0), \quad (4.42)$$

which gives the average *throughput* from station i to station j ,

$$d_{ij}(N, K) = \mu_i U_i(N, K) r_{ij}.$$

This is the average number of customers transferred from station i to station j in one unit of time. Since

$$U_i(N, K) = \sum_{\substack{n_1 + \dots + n_K = N \\ n_i > 0}} \pi(n) = \frac{1}{G(N, K)} \sum_{\substack{n_1 + \dots + n_K = N \\ n_i > 0}} \prod_{j=1}^K \rho_j^{n_j},$$

we see that $G(N, K)U_i(N, K)$ is the coefficient of z^N in

$$\tilde{g}_{N,i}(z) = \left(\prod_{\substack{j=1 \\ j \neq i}}^K \left(\sum_{n_j=0}^{\infty} \rho_j^{n_j} z^{n_j} \right) \right) \left(\sum_{n_i=1}^{\infty} \rho_i^{n_i} z^{n_i} \right).$$

Now,

$$\tilde{g}_{N,i}(z) = g_N(z) \rho_i z,$$

and therefore $G(N, K)U_i(N, K) = \rho_i G(N-1, K)$, i.e.,

$$U_i(N, K) = \rho_i \frac{G(N-1, K)}{G(N, K)}. \quad (4.43)$$

For additional information on algorithmic computation of closed networks characteristics, see, for instance, Robertazzi (1990).

Closed Jackson network models arise in the following situation, where an open network is operated with a blocking admission policy: If there are already N customers in the network, the newcomers wait at a gate (in a queue) until one customer is released from the network, at which time one among the blocked customers, if there are any, is admitted. In the network, there are at most N customers. At “saturation,” there is always one customer ready to replace a departing customer, since the gate queue is infinite, by definition of saturation. Therefore, at saturation, or for all practical purposes near saturation, everything looks as if a departing customer was being immediately recycled (see Figure 9.6.2).

It is important in practice to be able to compute the average number of customers $d(N)$ passing through the entrance point A per unit of time: This is the maximum throughput, and therefore it is related to the efficiency of the system, from the point of view of the operator (who makes money with customers). From the point of view of the quality of service, an important parameter is $W(N)$, the average time spent by a customer between A and B . As a matter of fact, N , $W(N)$, and $d(N)$ are related by

$$d(N)W(N) = N, \quad (4.44)$$

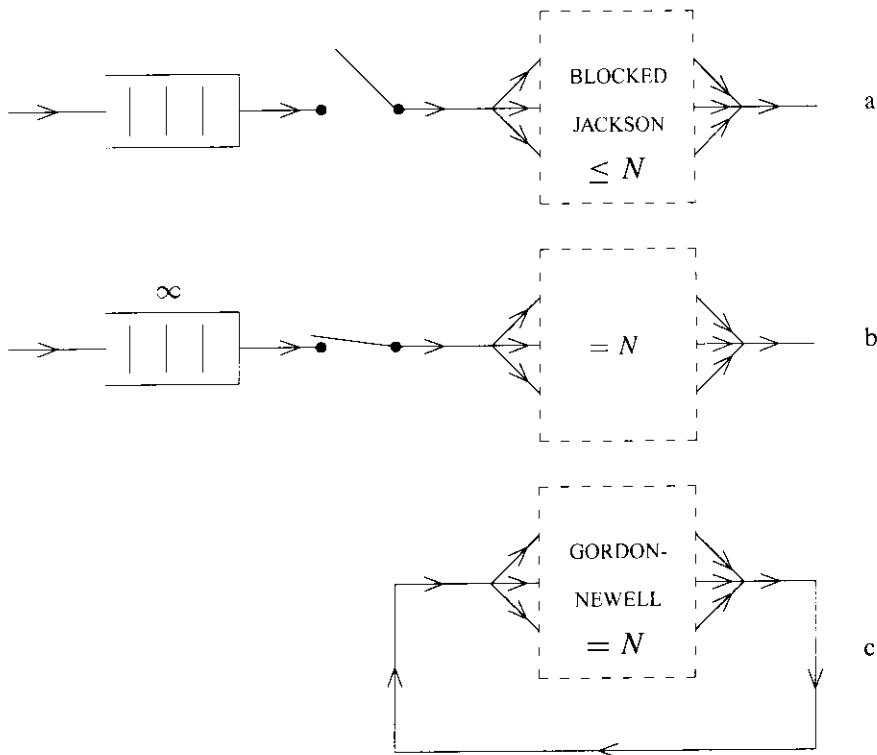


Figure 9.6.2. Gordon–Newell network as a saturated Jackson network

a particular case of Little’s formula; see (Wolff, 1989). One can therefore compute $d(N)$ via formulas such as (4.43) and the algorithm (4.41), and then $W(N)$ via (4.44), and then choose an operating point N that provides the required balance between the operator’s profit and the customer’s comfort.

Problems

9.1.1. The regenerative structure of a regular jump HMC suggests that for simulating such a chain, one needs to generate *two* random variables per transition: One fixes the time of occurrence of the event, and the other one tells what the next state is. Show that for a finite-state regular jump HMC, one really needs to generate one single exponential random variable per transition *in the long run*. Discuss the feasibility when the state space is large in terms of memory requirements.

9.3.1 Let $\{X(t)\}_{t \geq 0}$ be an irreducible ergodic regular jump HMC, and take for initial distribution $\mu = \pi$, the stationary distribution. Suppose, moreover, that with each state $i \in E$ is associated a subset $T(i) \subset E$ such that $\sum_{j \in T(i)} q_{ij} = \lambda$, independent of $i \in E$. Let N be

the counting process recording all transition times τ_n such that $X(\tau_n) \in T(X(\tau_n-))$. Show that N is a Poisson process with intensity λ .

9.3.2 Let N be an HPP with intensity $\lambda > 0$ defined on the probability space (Ω, \mathcal{F}, P) . Define a set function P' on \mathcal{F} by

$$P'(A) = E_P[L_T 1_A],$$

where E_P refers to expectation with respect to P , and

$$L_T = \left(\frac{\lambda'}{\lambda}\right)^{N(T)} \exp\{-(\lambda' - \lambda)T\},$$

where $\lambda' > 0$. Show that P' is a probability, and that under P' , N is an HPP with intensity λ' .

9.4.1 Consider the following model of a population in terms of birth and death. The lifetimes of the individuals appearing in the population are independent and exponentially distributed with the same mean μ^{-1} , where $\mu \in (0, \infty)$. An individual introduced into the population (by birth or by immigration) gives birth to children according to an HPP of intensity $\lambda > 0$ (the reader will interpret this statement) as long as she is alive. All the birth point processes attached to individuals are independent, and independent of all the lifetimes. There is also an immigration process, an HPP of intensity a , $a \in (0, \infty)$, independent of all the rest. Give a Poisson system description of this population process, and show that the process counting the number of individuals in the population is a birth-and-death process with parameters $\lambda_n = n\lambda + a$, $\mu_n = n\mu$.

9.4.2 Give a purely Poissonian description of the M/M/1/ ∞ /FIFO queue with instantaneous feedback, with a first-come-first-served discipline (head-of-the-line customer is served first); see Example 4.4 for the definition. Show that the congestion process counting the number of customers in the system at time t is a regular jump HMC, and give its infinitesimal generator.

9.4.3 Consider the general Poisson system of Section 4. Call $\tilde{\tau}_n$ the n th transition (maybe a pseudo-transition) time, where $\tilde{\tau}_0 \equiv 0$. Call \tilde{s}_n the source responsible for the transition at time $\tilde{\tau}_n$, with $\tilde{s}_0 \in S$ arbitrary and independent of the driving HPPs. Show that $\{(\tilde{s}_n, X(\tilde{\tau}_n))\}_{n \geq 0}$ forms a discrete-time HMC, and give its transition matrix.

9.4.4 Consider an M/M/1/ ∞ /LIFO preemptive resume queue. See Example 4.3 for the definition and the notation. Show that the congestion process (counting the customers in the system: ticket booth plus waiting room) is a regular jump HMC with the same infinitesimal generator as the corresponding process in an M/M/1/ ∞ /FIFO queueing system.

9.5.1 Consider the M/M/1/ ∞ /FIFO queue with instantaneous Bernoulli feedback of Example 4.4. Define the process $\{Z(t)\}_{t \geq 0}$ with values in $\{-1, +1\}$ as follows: $Z(0)$ is chosen independently of everything else, and the process $\{Z(t)\}_{t \geq 0}$ switches from one value to the other whenever a customer is fed-back. Show that $Y(t) = (X(t), Z(t))$ defines a regular jump HMC $\{Y(t)\}_{t \geq 0}$, and give its infinitesimal generator. Give a necessary and sufficient condition of ergodicity and compute the stationary distribution in the ergodic case. Give,


at equilibrium, the distribution of the times between transitions of $\{Z(t)\}_{t \geq 0}$. Show that the times of departure of customers definitively leaving the system form a Poisson process. Do the same for the times of end of service.

9.5.2 Prove that the congestion process in an $M/GI/1/\infty/FIFO$ queue observed at the departure times is a discrete-time HMC.

9.5.3 Let $\{X(t)\}_{t \geq 0}$ be an irreducible regular jump HMC with state space E and infinitesimal generator A . Suppose that it is ergodic, with stationary probability π . Show that at equilibrium ($P(X(t) = i) = \pi(i)$ for all $t \geq 0$, all $i \in E$), the average number of transitions from i to j in a unit time interval is $\pi(i)q_{ij}$. Use this to interpret the detailed balance equations for a reversible chain.

9.5.4 Consider the closed Jackson network of the theory, with a pure ring structure, i.e., $r_{i,i+1} = 1, i \in [1, K - 1]$, and $r_{K,1} = 1$. Compute the stationary distribution.

9.5.5 Consider the closed Jackson network below where all service times at different queues have the same (exponential) distribution of mean 3. Compute for $N = 1, \dots, 10$ the average time spent by a customer to go from the leftmost point A to the rightmost point B , and the average number of customers passing by A per unit time.



9.5.6 The following modification of the basic Jackson network is considered. For all $i \in [1, K]$, the server at station i has a speed of service $\phi_i(n_i)$ when there are n_i customers present in station i , where $\phi_i(k) > 0$ for all $k \geq 1$ and $\phi_i(0) = 0$. The new infinitesimal generator is obtained from the standard one by replacing μ_i by $\mu_i \phi_i(n_i)$. Check this and show that if for all $i \in [1, K]$,

$$A_i \stackrel{\text{def}}{=} 1 + \sum_{n_i=1}^{\infty} \left(\frac{\rho_i^{n_i}}{\prod_{k=1}^{n_i} \phi(k)} \right) < \infty,$$

where $\rho_i = \lambda_i / \mu_i$ and λ_i is the solution of the traffic equation, then the network is ergodic with stationary distribution

$$\pi(n) = \prod_{i=1}^K \pi_i(n_i),$$

where

$$\pi_i(n_i) = \frac{1}{A_i} \frac{\rho_i^{n_i}}{\prod_{k=1}^{n_i} \phi_i(k)}.$$

9.5.7 The basic Jackson network of the theory is now modified by assigning K_i servers with unit speed of service to each station $i \in [1, K]$. Show that the infinitesimal generator

of this network is the same as that of the network of the previous problem, with $\phi_i(n_i) = \inf(K_i, n_i)$. Give the ergodicity condition corresponding to that of the previous problem.

9.5.8 Consider the closed Jackson network of the theory, with $N = 1$ customer. Let $\{Y(t)\}_{t \geq 0}$ be the process giving the position of this customer, i.e., $Y(t) = i$ if he/she is in station i at time t . Show that $\{Y(t)\}_{t \geq 0}$ is a regular jump HMC, irreducible, and give its stationary distribution. Observe that $\lambda_i = \alpha_i \mu_i$, $i \in [1, K]$, is a solution of the traffic equation.

9.5.9 Consider the Jackson network of the theory with $K = 2$, $\bar{\lambda}_2 = 0$, $r_{12} = 1$, $r_2 = 1$ (two queues in series). Show that the point process counting the customers passing from station 1 to station 2 is a Poisson process, and give its intensity. Show that for all $t \geq 0$, $X_2(t)$ is independent of $(X_1(s), 0 \leq s \leq t)$.

Appendix

1 Number Theory and Calculus

1.1 Greatest Common Divisor

Let $a_1, \dots, a_k \in \mathbb{N}$ be such that $\max(a_1, \dots, a_k) > 0$. Their greatest common divisor (g.c.d) is the largest positive integer dividing all of them. It is denoted by $\text{g.c.d}(a_1, \dots, a_k)$. Clearly, removing all zero elements does not change the g.c.d, so that we may assume without loss of generality that all the a_k 's are positive.

Let $\{a_n\}_{n \geq 1}$ be a sequence of positive integers. The sequence $\{d_k\}_{k \geq 1}$ defined by $d_k = \text{g.c.d}(a_1, \dots, a_k)$ is bounded below by 1 and is nonincreasing, and it therefore has a limit $d \geq 1$, a positive integer called the g.c.d of the sequence $\{a_n\}_{n \geq 1}$. Since the d_k 's are integers, the limit is attained after a finite number of steps, and therefore there exists a positive integer k_0 such that $d = \text{g.c.d}(a_1, \dots, a_k)$ for all $k \geq k_0$.

Lemma 1.1.

Let $S \subset \mathbb{Z}$ contain at least one nonzero element and be closed under addition and subtraction. Then S contains a least positive element a , and $S = \{ka ; k \in \mathbb{Z}\}$.

Proof. Let $c \in S$, $c \neq 0$. Then $c - c = 0 \in S$. Also $0 - c = -c \in S$. Therefore, S contains at least one positive element. Denote by a the smallest positive element of S . Since S is closed under addition and subtraction, S contains a , $a + a = 2a, \dots$ and $0 - a = -a, 0 - 2a = -2a, \dots$, that is, $\{ka ; k \in \mathbb{Z}\} \subset S$.

Let $c \in S$. Then $c = ka + r$, where $k \in \mathbb{Z}$ and $0 \leq r < a$. Since $r = c - ka \in S$, we cannot have $r > 0$, because this would contradict the definition of a as the smallest positive integer in S . Therefore, $r = 0$, i.e., $c = ka$. Therefore, $S \subset \{ka ; k \in \mathbb{Z}\}$. \square

Lemma 1.2.

Let a_1, \dots, a_k be positive integers with greatest common divisor d . There exist $n_1, \dots, n_k \in \mathbb{Z}$ such that $d = \sum_{i=1}^k n_i a_i$.

Proof. The set $S = \left\{ \sum_{i=1}^k n_i a_i ; n_1, \dots, n_k \in \mathbb{Z} \right\}$ is closed under addition and subtraction, and therefore, by Lemma 1.1, $S = \{ka ; k \in \mathbb{Z}\}$, where $a = \sum_{i=1}^k n_i a_i$ is the smallest positive integer in S .

Since d divides all the a_i 's, d divides a , and therefore $0 < d \leq a$. Also, each a_i is in S and is therefore a multiple of a , which implies that $a \leq \text{g.c.d.}(a_1, \dots, a_k) = d$. Therefore, $d = a$. \square

Theorem 1.1.

Let d be the g.c.d of $A = \{a_n ; n \geq 1\}$, a set of positive integers that is closed under addition. Then A contains all but a finite number of the positive multiples of d .

Proof. We may assume without loss of generality that $d = 1$ (otherwise, divide all the a_n by d). For some k , $d = 1 = \text{g.c.d.}(a_1, \dots, a_k)$, and therefore by Lemma 1.2,

$$1 = \sum_{i=1}^k n_i a_i$$

for some $n_1, \dots, n_k \in \mathbb{Z}$. Separating the positive from the negative terms in the latter equality, we have $1 = M - P$, where M and P are in A .

Let $n \in \mathbb{N}$, $n \geq P(P-1)$. We have $n = aP + r$, where $r \in [0, P-1]$. Necessarily, $a \geq P-1$, otherwise, if $a \leq P-2$, then $n = aP + r < P(P-1)$. Using $1 = M - P$, we have that $n = aP + r(M - P) = (a - r)P + rM$. But $a - r \geq 0$. Therefore, n is in A . We have thus shown that any $n \in \mathbb{N}$ sufficiently large—say $n \geq P(P-1)$ —is in A . \square

1.2 Abel's Theorem

Lemma 1.3. *Abel's Lemma*

Let $\{b_n\}_{n \geq 1}$ and $\{a_n\}_{n \geq 1}$ be two sequences of real numbers such that

$$b_1 \geq b_2 \geq \dots \geq b_n \geq 0,$$

and such that for some real numbers m and M , and all $n \geq 1$,

$$m \leq a_1 + \dots + a_n \leq M.$$

Then, for all $n \geq 1$,

$$b_1 m \leq a_1 b_1 + \dots + a_n b_n \leq b_1 M. \quad (1.1)$$

Proof. Let $s_n = a_1 + \dots + a_n$, and use Abel's summation technique to obtain

$$\begin{aligned} a_1 b_1 + \dots + a_n b_n &= b_1 s_1 + b_2 (s_2 - s_1) + \dots + b_n (s_n - s_{n-1}) \\ &= s_1 [b_1 - b_2] + \dots + s_{n-1} [b_{n-1} - b_n] + s_n [b_n]. \end{aligned}$$

The bracketed terms are all nonnegative, and therefore replacing each s_i by its lower bound or upper bound yields the result. \square

We recall without proof a standard result of calculus.

Lemma 1.4.

The sum of a uniformly convergent series of continuous functions is a continuous function.

Theorem 1.2. Abel's Theorem

Let $\{a_n\}_{n \geq 1}$ be a sequence of real numbers such that the radius of convergence of the power series $\sum_{n=0}^{\infty} a_n z^n$ is 1. Suppose that the sum $\sum_{n=0}^{\infty} a_n$ is convergent. Then the power series $\sum_{n=0}^{\infty} a_n x^n$ is uniformly convergent in $[0, 1]$ and

$$\lim_{x \uparrow 1} \sum_{n=0}^{\infty} a_n x^n = \sum_{n=0}^{\infty} a_n, \quad (1.2)$$

where $x \uparrow 1$ means that x tends to 1 from below.

Proof. It suffices to prove that $\sum_{n=0}^{\infty} a_n x^n$ is uniformly convergent in $[0, 1]$, since (1.2) then follows by Lemma 1.4. Write $A_n^p = a_n + \dots + a_p$. By convergence of $\sum_{n=0}^{\infty} a_n$, for all $\epsilon > 0$, there exists $n_0 \geq 1$ such that $p \geq n \geq n_0$ implies $|A_n^p| \leq \epsilon$, and therefore, since for $x \in [0, 1]$, the sequence $\{x^n\}_{n \geq 0}$ is nonincreasing, Abel's lemma gives, for all $x \in [0, 1]$,

$$|a_n x^n + \dots + a_p x^p| \leq \epsilon x^n \leq \epsilon,$$

from which uniform convergence follows. \square

Theorem 1.3.

Let $\{a_n\}_{n \geq 0}$ be a sequence of *nonnegative* real numbers such that the power series $\sum_{n=0}^{\infty} a_n z^n$ has a radius of convergence 1. If

$$\lim_{x \uparrow 1} \sum_{n=0}^{\infty} a_n x^n = a \leq \infty, \quad (1.3)$$

then

$$\sum_{n=0}^{\infty} a_n = a. \quad (1.4)$$

Proof. For $x \in [0, 1)$, $\sum_{n=0}^{\infty} a_n x^n \leq \sum_{n=0}^{\infty} a_n$ (the a_n are nonnegative), and therefore by (1.3), $a \leq \sum_{n=0}^{\infty} a_n$. This proves the result when $a = \infty$.

We now suppose that $a < \infty$. From $\sum_{n=0}^p a_n = \lim_{x \uparrow 1} \sum_{n=0}^p a_n x^n$, we have that $\sum_{n=0}^p a_n \leq a < \infty$. Thus, $\sum_{n=1}^p a_n$ is a nondecreasing sequence, converging to some α , $\alpha \leq a < \infty$. By Abel's theorem, $\lim_{x \uparrow 1} \sum_{n=0}^{\infty} a_n x^n = \alpha$, and therefore $\alpha = a$ and $\sum_{n=0}^{\infty} a_n = a$. \square

Theorem 1.3 is a partial converse of Abel's theorem, requiring the additional condition that $\{a_n\}_{n \geq 0}$ be a nonnegative sequence. In general, the partial converses of Abel's theorem are called Tauberian theorems. For instance:

Theorem 1.4. *Tauber's theorem*

Let $\{a_n\}_{n \geq 0}$ be a sequence of real numbers such that the radius of convergence of $\sum_{n=0}^{\infty} a_n z^n$ is 1. If

$$\lim_{x \uparrow 1} \sum_{n=0}^{\infty} a_n x^n = a < \infty \quad (1.5)$$

and if moreover

$$a_n = o\left(\frac{1}{n}\right), \quad (1.6)$$

then $\sum_{n=0}^{\infty} a_n$ converges to a .

The proof is omitted, since the result will not be used in this book; see, however, E.C. Titchmarsh, *The Theory of Functions*, n th edition, Oxford University Press, 1986, §1.23.

Theorem 1.5. *Cesaro's lemma*

Let $\{b_n\}_{n \geq 0}$ be a sequence of real numbers such that

$$\lim_{n \uparrow \infty} b_n = 0.$$

Then

$$\lim_{n \uparrow \infty} \frac{b_1 + \cdots + b_n}{n} = 0.$$

Proof. The sequence $\{b_n\}_{n \geq 0}$ is bounded in absolute value, say by K . For fixed arbitrary $\epsilon > 0$, there exists n_0 such that $n > n_0$ implies $|b_n| \leq \epsilon$, and therefore

$$\begin{aligned} \left| \frac{b_1 + \cdots + b_n}{n} \right| &\leq \left| \frac{b_1 + \cdots + b_{n_0}}{n} \right| + \left| \frac{b_{n_0} + \cdots + b_n}{n} \right| \\ &\leq \frac{n_0 K}{n} + \frac{n - n_0}{n} \epsilon \leq 2\epsilon, \end{aligned}$$

if n is sufficiently large. □

1.3 Lebesgue's Theorems for Series

The results in this subsection concern the validity of exchanging limits and infinite sums.

Theorem 1.6. *Dominated Convergence*

Let $\{a_{nk}\}_{n \geq 1, k \geq 1}$ be an array of real numbers such that for some sequence $\{b_k\}_{k \geq 1}$ of non-negative numbers satisfying

$$\sum_{k=1}^{\infty} b_k < \infty, \quad (1.7)$$

it holds that for all $n \geq 1, k \geq 1$,

$$|a_{nk}| \leq b_k. \quad (1.8)$$

If moreover for all $k \geq 1$,

$$\lim_{n \uparrow \infty} a_{nk} = a_k, \tag{1.9}$$

then

$$\lim_{n \uparrow \infty} \sum_{k=1}^{\infty} a_{nk} = \sum_{k=1}^{\infty} a_k. \tag{1.10}$$

Proof. Let $\epsilon > 0$ be fixed. Since $\sum_{k=1}^{\infty} b_k$ is a convergent series, one can find $M = M(\epsilon)$ such that $\sum_{k=M+1}^{\infty} b_k < \frac{\epsilon}{3}$. Since $|a_{nk}| \leq b_k$ and therefore $|a_k| \leq b_k$, we have

$$\sum_{k=M+1}^{\infty} |a_{nk}| + \sum_{k=M+1}^{\infty} |a_k| \leq \frac{2\epsilon}{3}.$$

Now, for sufficiently large n ,

$$\sum_{k=1}^M |a_{nk} - a_k| \leq \frac{\epsilon}{3}.$$

Therefore, for sufficiently large n ,

$$\left| \sum_{k=1}^{\infty} a_{nk} - \sum_{k=1}^{\infty} a_k \right| \leq \sum_{k=1}^M |a_{nk} - a_k| + \sum_{k=M+1}^{\infty} |a_{nk}| + \sum_{k=M+1}^{\infty} |a_k| \leq \frac{\epsilon}{3} + \frac{2\epsilon}{3} = \epsilon. \quad \square$$

Theorem 1.7. Monotone Convergence

Let $\{a_{nk}\}_{n \geq 1, k \geq 1}$ be an array of nonnegative real numbers such that for all $k \geq 1$, the sequence $\{a_{nk}\}_{n \geq 1}$ is nondecreasing and such that

$$\lim_{n \uparrow \infty} a_{nk} = a_k \leq \infty. \tag{1.11}$$

Then

$$\lim_{n \uparrow \infty} \sum_{k=1}^{\infty} a_{nk} = \sum_{k=1}^{\infty} a_k. \tag{1.12}$$

Proof. If $\sum_{k=1}^{\infty} a_k < \infty$, the result is a direct application of the dominated convergence theorem.

For the case $\sum_{k=1}^{\infty} a_k = \infty$, let $A > 0$ be fixed, and choose $M = M(A)$ such that $\sum_{k=1}^M a_k \geq 2A$. For sufficiently large n , $\sum_{k=1}^M (a_k - a_{nk}) \leq A$. Therefore, for sufficiently large n ,

$$\sum_{k=1}^{\infty} a_{nk} \geq \sum_{k=1}^M a_{nk} + \sum_{k=1}^M (a_{nk} - a_k) \geq 2A - A = A. \quad \square$$

Theorem 1.8. Fatou's Lemma

Let $\{a_{nk}\}_{n \geq 1, k \geq 1}$ be an array of nonnegative real numbers. Then

$$\sum_{k=1}^{\infty} \liminf_{n \uparrow \infty} a_{nk} \leq \liminf_{n \uparrow \infty} \sum_{k=1}^{\infty} a_{nk}. \tag{1.13}$$

Proof. By definition of \liminf , for fixed k ,

$$z_{nk} \stackrel{\text{def}}{=} \inf(a_{nk}, a_{n+1,k}, \dots)$$

increases, as $n \uparrow \infty$, to $\liminf_{n \uparrow \infty} a_{nk}$. Therefore, by monotone convergence,

$$\sum_{k=1}^{\infty} \liminf_{n \uparrow \infty} a_{nk} = \lim_{n \uparrow \infty} \uparrow \sum_{k=1}^{\infty} z_{nk}.$$

But since $z_{nk} \leq a_{nk}$,

$$\sum_{k=1}^{\infty} z_{nk} \leq \sum_{k=1}^{\infty} a_{nk},$$

and therefore

$$\lim_{n \uparrow \infty} \sum_{k=1}^{\infty} z_{nk} \leq \lim_{n \uparrow \infty} \inf \sum_{k=1}^{\infty} a_{nk}. \quad \square$$

1.4 Infinite Products

Theorem 1.9. Infinite Products

Let $\{a_n\}_{n \geq 1}$ be a sequence of numbers of the interval $[0, 1)$.

(a) If $\sum_{n=1}^{\infty} a_n < \infty$, then

$$\lim_{n \uparrow \infty} \prod_{k=1}^n (1 - a_k) > 0.$$

(b) If $\sum_{n=1}^{\infty} a_n = \infty$, then

$$\lim_{n \uparrow \infty} \prod_{k=1}^n (1 - a_k) = 0.$$

Proof. (a): For any numbers c_1, \dots, c_n in $[0, 1)$, it holds that $(1 - c_1)(1 - c_2) \cdots (1 - c_n) \geq 1 - c_1 - c_2 - \cdots - c_n$ (proof by induction). Since $\sum_{n=1}^{\infty} a_n$ converges, there exists N such that for all $n \geq N$,

$$a_N + \cdots + a_n < \frac{1}{2}.$$

Therefore, defining $\pi(n) = \prod_{k=1}^n (1 - a_k)$, we have that for all $n \geq N$,

$$\frac{\pi(n)}{\pi(N-1)} = (1 - a_N) \cdots (1 - a_n) \geq 1 - (a_N + \cdots + a_n) \geq \frac{1}{2}.$$

Therefore, the sequence $\{\pi(n)\}_{n \geq N}$ is a nonincreasing sequence bounded from below by $\frac{1}{2}\pi(N-1) > 0$, so that $\lim_{n \uparrow \infty} \pi(n) > 0$.

(b): Using the inequality $1 - a \leq e^{-a}$ when $a \in [0, 1)$, we have that $\pi(n) \leq e^{-a_1 - a_2 - \cdots - a_n}$, and therefore, if $\sum_{n=1}^{\infty} a_n = \infty$, $\lim_{n \uparrow \infty} \pi(n) = 0$. \square

1.5 Tychonov's Theorem

Theorem 1.10. *Tychonov's Theorem*

Let $\{x_n\}_{n \geq 0}$ be a sequence of elements of $[0, 1]^{\mathbb{N}}$, that is

$$x_n = (x_n(0), x_n(1), \dots),$$

such that $x_n(k) \in [0, 1]$ for all $k, n \in \mathbb{N}$. There exists a strictly increasing sequence of integers $\{n_l\}_{l \geq 0}$ and an element $x \in \{0, 1\}^{\mathbb{N}}$ such that

$$\lim_{l \uparrow \infty} x_{n_l}(k) = x(k) \tag{1.14}$$

for all $k \in \mathbb{N}$.

Proof. Since the sequence $\{x_n(0)\}_{n \geq 0}$ is contained in the closed interval $[0, 1]$, by the Boltzano–Weierstrass theorem, one can extract a subsequence $\{x_{n_0(l)}(0)\}_{l \geq 0}$ such that

$$\lim_{l \uparrow \infty} x_{n_0(l)}(0) = x(0)$$

for some $x(0) \in [0, 1]$. In turn, one can extract from $\{x_{n_0(l)}(1)\}_{l \geq 0}$ a subsequence $\{x_{n_1(l)}(1)\}_{l \geq 0}$ such that

$$\lim_{l \uparrow \infty} x_{n_1(l)}(1) = x(1)$$

for some $x(1) \in [0, 1]$. Note that

$$\lim_{l \uparrow \infty} x_{n_1(l)}(0) = x(0).$$

Iterating this process, we obtain for all $j \in \mathbb{N}$ a sequence $\{x_{n_j(l)}\}_{l \geq 0}$ that is a subsequence of each sequence $\{x_{n_0(l)}(1)\}_{l \geq 0}, \dots, \{x_{n_{j-1}(l)}(1)\}_{l \geq 0}$ and such that

$$\lim_{l \uparrow \infty} x_{n_j(l)}(k) = x(k)$$

for all $k \leq j$, where $x(1), \dots, x(j) \in [0, 1]$. The diagonal sequence $n_l = n_l(l)$ then establishes (1.14). □

1.6 Subadditive Functions

Theorem 1.11. *Subadditive Function Theorem*

Let $f : (0, \infty) \rightarrow \mathbb{R}$ be a nonnegative function such that $\lim_{t \downarrow 0} f(t) = 0$, and assume that f is subadditive, that is,

$$f(t + s) \leq f(t) + f(s)$$

for all $s, t \in (0, \infty)$. Define the (possibly infinite) nonnegative real number

$$q = \sup_{t > 0} \frac{f(t)}{t}.$$

Then

$$\lim_{h \downarrow 0} \frac{f(h)}{h} = q. \tag{1.15}$$

Proof. By definition of q , there exists for each $a \in [0, q)$ a real number $t_0 > 0$ such that $\frac{f(t_0)}{t_0} \geq a$.

To each $t > 0$, there corresponds an $n \in \mathbb{N}$ such that $t_0 = nt + \delta$ with $0 \leq \delta < t$. Since f is subadditive, $f(t_0) \leq nf(t) + f(\delta)$, and therefore

$$a \leq \frac{nt}{t_0} \frac{f(t)}{t} + \frac{f(\delta)}{t_0},$$

which implies

$$a \leq \liminf_{t \downarrow 0} \frac{f(t)}{t}.$$

($\delta \rightarrow 0$, and therefore $f(\delta) \rightarrow 0$, as $t \rightarrow 0$; also $\frac{nt}{t_0} \rightarrow 1$ as $t \rightarrow 0$.) Therefore,

$$a \leq \liminf_{t \downarrow 0} \frac{f(t)}{t} \leq \limsup_{t \downarrow 0} \frac{f(t)}{t} \leq q,$$

from which (1.15) follows, since a can be chosen arbitrarily close to q (arbitrarily large, when $q = \infty$). \square

2 Linear Algebra

2.1 Eigenvalues and Eigenvectors

The basic results of the theory of matrices relative to eigenvalues and eigenvectors will now be reviewed, and the reader is referred to the classic text of Gantmacher (1959), or to the more recent (Horn and Johnson, 1985) for the proofs.

Let A be a square matrix of dimension $r \times r$, with complex coefficients. If there exists a scalar $\lambda \in \mathbb{C}$ and a column vector $v \in \mathbb{C}^r$, $v \neq 0$, such that

$$Av = \lambda v \text{ (resp., } v^T A = \lambda v^T), \quad (2.1)$$

then v is called a right eigenvector (resp., a left eigenvector) of A associated with the eigenvalue λ . There is no need to distinguish between right and left eigenvalues because if there exists a left eigenvector associated with the eigenvalue λ , then there exists a right eigenvector associated with the same eigenvalue λ . This follows from the facts that the set of eigenvalues of A is exactly the set of roots of the *characteristic equation*

$$\det(\lambda I - A) = 0 \quad (2.2)$$

where I is the $r \times r$ identity matrix, and that

$$\det(\lambda I - A) = \det(\lambda I - A^T).$$

The *algebraic multiplicity* of λ is its multiplicity as a root of the *characteristic polynomial* $\det(\lambda I - A)$.

If $\lambda_1, \dots, \lambda_k$ are *distinct* eigenvalues corresponding to the right eigenvectors v_1, \dots, v_k and the left eigenvectors u_1, \dots, u_k , then v_1, \dots, v_k are independent, and so are u_1, \dots, u_k .

Call R_λ (resp. L_λ) the set of right eigenvectors (resp., left eigenvectors) associated with the eigenvalue λ , plus the null vector. Both L_λ and R_λ are vector subspaces of \mathbb{C}^r , and they have the same dimension, called the *geometric multiplicity* of λ . In particular, the largest number of independent right eigenvectors (resp., left eigenvectors) cannot exceed the sum of the geometric multiplicities of the distinct eigenvalues.

The matrix A is called *diagonalizable* if there exists a nonsingular matrix Γ of the same dimensions such that

$$\Gamma A \Gamma^{-1} = \Lambda, \quad (2.3)$$

where

$$\Lambda = \text{diag}(\lambda_1, \dots, \lambda_r)$$

for some $\lambda_1, \dots, \lambda_r \in \mathbb{C}$, not necessarily distinct. It follows from (2.3) that with $U = \Gamma^T$, $U^T A = U^T \Lambda$, and with $V = \Gamma^{-1}$, $AV = V\Lambda = \Lambda V$, and therefore $\lambda_1, \dots, \lambda_r$ are eigenvalues of A , and the i th row of $U^T = \Gamma$ (resp., the i th column of $V = \Gamma^{-1}$) is a left eigenvector (resp., right eigenvector) of A associated with the eigenvalue λ_i . Also, $A = \Gamma^{-1} \Lambda \Gamma$ and therefore $A^n = \Gamma^{-1} \Lambda^n \Gamma$, that is

$$A^n = V \Lambda^n U^T. \quad (2.4)$$

Clearly, if A is diagonalizable, the sum of the geometric multiplicities of A is exactly equal to r . It turns out that the latter is a sufficient condition of diagonalizability of A . Therefore, A is diagonalizable if and only if the sum of the geometric multiplicities of the distinct eigenvalues of A is equal to r .

Example 2.1. Distinct Eigenvalues

By the last result, if the eigenvalues of A are distinct, A is diagonalizable. In this case, the diagonalization process can be described as follows. Let $\lambda_1, \dots, \lambda_r$ be the r distinct eigenvalues and let u_1, \dots, u_r and v_1, \dots, v_r be the associated sequences of left and right eigenvectors, respectively. As mentioned above, u_1, \dots, u_r form an independent collection of vectors, and so do v_1, \dots, v_r . Define

$$U = [u_1 \cdots u_r], V = [v_1 \cdots v_r]. \quad (2.5)$$

Observe that if $i \neq j$, $u_i^T v_j = 0$. Indeed, $\lambda_i u_i^T v_j = u_i^T A v_j = \lambda_j u_i^T v_j$, which implies $(\lambda_i - \lambda_j) u_i^T v_j = 0$, and in turn $u_i^T v_j = 0$, since $\lambda_i \neq \lambda_j$ by hypothesis. Since eigenvectors are determined up to multiplication by an arbitrary nonnull scalar, one can choose them in such a way that $u_i^T v_i = 1$ for all $i \in [1, r]$. Therefore,

$$U^T V = I, \quad (2.6)$$

where I is the $r \times r$ identity matrix. Also, by definition of U and V ,

$$U^T A = \Lambda U^T, AV = \Lambda V. \quad (2.7)$$

In particular, by (2.6), $A = V \Lambda U^T$. From the last identity and (2.6) again, we obtain for all $n \geq 0$,

$$A^n = V \Lambda^n U^T, \quad (2.8)$$

that is,

$$A^n = \sum_{i=1}^r \lambda_i^n v_i u_i^T. \quad (2.9)$$

◇

2.2 Exponential of a Matrix

Theorem 2.1.

Let A be a finite $r \times r$ matrix with complex elements. For each $t > 0$, the series

$$\sum_{n=0}^{\infty} t^n \frac{A^n}{n!} \quad (2.10)$$

converges componentwise, and it is denoted by e^{tA} . For any $r \times r$ complex matrix B such that

$$AB = BA \quad (2.11)$$

we have

$$e^{t(A+B)} = e^{tA} e^{tB} \quad (2.12)$$

for all $t > 0$. Also,

$$\frac{d}{dt} e^{tA} = A e^{tA} = e^{tA} A. \quad (2.13)$$

Proof. Let $a_{ij} = a_{ij}(1)$, where $a_{ij}(n)$ is the general entry of A^n . Let $\Delta = \max |a_{ij}|$. A straightforward inductive proof shows that

$$a_{ij}(n) \leq r^{n-1} \Delta^n \quad (2.14)$$

for all i, j . Therefore, the general term of (2.10) is bounded in absolute value by the convergent series

$$1 + \sum_{n=1}^{\infty} t^n \frac{r^{n-1} \Delta^n}{n!},$$

and the componentwise convergence of (2.10) is proven.

Using the commutativity hypothesis (2.11),

$$\begin{aligned} \sum_{k=0}^n t^k \frac{(A+B)^k}{k!} &= \sum_{k=0}^n \sum_{j=0}^k t^j \frac{A^j}{j!} t^{k-j} \frac{B^{k-j}}{(k-j)!} \\ &= \left(\sum_{j=0}^n t^j \frac{A^j}{j!} \right) \left(\sum_{l=0}^n t^l \frac{B^l}{l!} \right) + C_n, \end{aligned}$$

where it is easy to bound C_n and show that $\lim_{n \rightarrow \infty} C_n = 0$, so that (2.12) follows by letting $n \rightarrow \infty$. To prove (2.13), write

$$\frac{e^{(t+h)A} - e^{tA}}{h} = e^{tA} \frac{e^{hA} - I}{h} = \frac{e^{hA} - I}{h} e^{tA},$$

and

$$\frac{e^{hA} - I}{h} = A + h \left(\sum_{n \geq 2} h^{n-2} \frac{A^n}{n!} \right).$$

In view of (2.14), each element of the latter matrix series is bounded, and the result follows. \square

Example 2.2.

If A is diagonalizable, i.e.,

$$A = V \Lambda U^T,$$

where $\Lambda = \text{diag}\{\lambda_1, \dots, \lambda_r\}$ is the diagonal eigenvalue matrix, $V = [v_1, \dots, v_r]$, $U = [u_1, \dots, u_r]$, v_i (resp., u_i) is a right (resp., left) eigenvector corresponding to λ_i , and $U^T V = I$, then

$$A^n = V \Lambda^n U^T,$$

so that

$$e^{tA} = \sum_{n=0}^{\infty} t^n \frac{A^n}{n!} = V \sum_{n=0}^{\infty} t^n \frac{\Lambda^n}{n!} U^T.$$

But

$$\sum_{n=0}^{\infty} t^n \frac{\Lambda^n}{n!} = \text{diag}\{e^{\lambda_1 t}, \dots, e^{\lambda_r t}\}$$

Finally,

$$e^{tA} = V \text{diag}\{e^{\lambda_1 t}, \dots, e^{\lambda_r t}\} U^T.$$

2.3 Gershgorin's Bound

Theorem 2.2. Gershgorin's Bound

Let A be a finite $r \times r$ matrix with complex elements. Then for any eigenvalue λ , and all $k \in [1, r]$,

$$|\lambda - a_{kk}| \leq \min(r_k, s_k), \quad (2.15)$$

where $r_k = \sum_{j=1, j \neq k}^r |a_{kj}|$ and $s_k = \sum_{j=1, j \neq k}^r |a_{jk}|$.

Proof. Let $v = (v(1), \dots, v(r))^T$ be a right eigenvector corresponding to the eigenvalue λ , and let k be such that

$$|v(k)| = \|v\|_{\infty} \stackrel{\text{def}}{=} \max_i |v(i)|.$$

The k th equality in $Av = \lambda v$ is

$$\sum_{j=1}^r a_{kj}v(j) = \lambda v(k),$$

or equivalently,

$$(\lambda - a_{kk})v(k) = \sum_{j=1, j \neq k}^r a_{kj}v(j).$$

Therefore,

$$|\lambda - a_{kk}||v(k)| \leq \sum_{j=1, j \neq k}^r |a_{kj}||v(j)| \leq r_k \|v\|_{\infty}.$$

Dividing by $\|v\|_{\infty} = |v(k)|$ yields

$$|\lambda - a_{kk}| \leq r_k.$$

By considering the left eigenvector u associated with λ , we obtain

$$|\lambda - a_{kk}| \leq s_k. \quad \square$$

Corollary 2.1.

Let \mathbf{P} be an $r \times r$ stochastic matrix. Then, for all eigenvalues λ ,

$$|\lambda - p_{kk}| \leq 1 - p_{kk}. \quad (2.16)$$

In particular, when the smallest eigenvalue λ_r is real,

$$\lambda_r \geq -1 + 2 \min_k (p_{kk}). \quad (2.17)$$

Proof. For a stochastic matrix,

$$r_k = \sum_{j=1, j \neq k}^r p_{kj} = 1 - p_{kk}. \quad \square$$

3 Probability

3.1 Expectation Revisited

There are two basic technical tools in probability theory that are indispensable, in this book and elsewhere: the monotone convergence theorem and the dominated convergence theorem. These results give sufficient conditions that allow interchange of limit and expectation.

For most purposes in this book, we need only the versions of the Lebesgue theorems given in Section 1.3. However, on a few occasions, the more general statements in terms of

expectations are necessary. To state them in their generality, we need to briefly summarize the definition of expectation for arbitrary random variables, which are neither discrete nor absolutely continuous. For proofs, see (Shiryayev, 1987) or (Williams, 1991).

Let X be a nonnegative random variable. Its expectation, denoted by $E[X]$ as usual, is defined by

$$E[X] = \lim_{n \uparrow \infty} \left\{ \sum_{k=0}^{n2^n-1} \frac{k}{2^n} P\left(\frac{k}{2^n} \leq X < \frac{k+1}{2^n}\right) + nP(X > n) \right\}. \quad (3.1)$$

The general term of the right-hand side is clearly increasing with n , and therefore the limit is well-defined, possibly infinite, however.

From (3.1), it is clear that

$$E[1_A] = P(A). \quad (3.2)$$

For a random variable X that is not nonnegative, the procedure already used to define $E[X]$ in the discrete case and in the probability density case is still applicable, i.e., $E[X] = E[X^+] - E[X^-]$ if not both $E[X^+]$ and $E[X^-]$ are infinite. If $E[X^+]$ and $E[X^-]$ are infinite, the expectation is not defined. If $E[|X|] < \infty$, X is said to be integrable, and then $E[X]$ is a finite number.

The basic properties of the expectation so defined are linearity and monotonicity: If X_1 and X_2 are random variables with expectations, then for all $\lambda_1, \lambda_2 \in \mathbb{R}$,

$$E[\lambda_1 X_1 + \lambda_2 X_2] = \lambda_1 E[X_1] + \lambda_2 E[X_2] \quad (3.3)$$

whenever the right-hand side has meaning (i.e., is not an $\infty - \infty$ form). Also, if $X_1 \leq X_2$, P-a.s., then

$$E[X_1] \leq E[X_2]. \quad (3.4)$$

It follows from this that if $E[X]$ is well-defined, then

$$|E[X]| \leq E[|X|]. \quad (3.5)$$

These properties will be accepted without proof. It is clear from (3.1) and (3.2) that if X is a simple random variable, i.e.,

$$X(\omega) = \sum_{i=1}^N \alpha_i 1_{A_i}(\omega),$$

where $\alpha_i \in \mathbb{R}$, $A_i \in \mathcal{F}$, then

$$E[X] = \sum_{i=1}^N \alpha_i P(A_i).$$

The definitions of the mean, the variance, and the characteristic function of a random variable are the same as in the discrete case and the probability density case:

$$\begin{aligned} m_X &= E[X], \\ \sigma_X^2 &= E[(X - m_X)^2] = E[X^2] - m_X^2, \\ \phi_X(u) &= E[e^{iuX}], \end{aligned}$$

where $E[e^{iuX}] = E[\cos(uX)] + iE[\sin(uX)]$. (For a complex valued random variable $X = X_1 + iX_2$, where X_1 and X_2 are real valued integrable random variables, $E[X] = E[X_1] + iE[X_2]$ defines the expectation of X .)

Now let $g : \mathbb{R} \rightarrow \mathbb{R}$ be some function such that $E[|g(X)|] < \infty$ where X is a real random variable with c.d.f. $F(x)$. By definition of the symbol on the left-hand side of (3.6),

$$\int_{\mathbb{R}} g(x)dF(x) := E[g(X)], \quad (3.6)$$

and $\int_{\mathbb{R}} g(x)dF(x)$ is called the Stieltjes–Lebesgue integral of $g(x)$ with respect to $F(x)$. It can be shown that if X admits a p.d.f. $f(x)$, i.e., $F(x) = \int_{-\infty}^x f(y)dy$, then

$$\int_{\mathbb{R}} g(x)dF(x) = \int_{\mathbb{R}} g(x)f(x)dx$$

in accordance with our previous definition of expectation for absolutely continuous random variables.

Clearly, the Stieltjes–Lebesgue integral inherits from expectation the properties of linearity and monotonicity.

At this point, we shall quote without proof a formula that is true whenever X is a *nonnegative* random variable:

$$E[X] = \int_0^{\infty} [1 - F(x)]dx. \quad (3.7)$$

We will now mention a difficulty that was concealed in the definition of expectation in the case of a random vector $X = (X_1, \dots, X_n)$ with a probability density f_X . For instance, we have defined for $Y = g(X_1, \dots, X_n)$ the expectation $E[Y]$ by

$$E[Y] = \int_{-\infty}^{+\infty} \cdots \int_{-\infty}^{+\infty} g(x_1, \dots, x_n)f_X(x_1, \dots, x_n)dx_1 \cdots dx_n.$$

But now, suppose that Y admits a probability density f_Y . The expectation $E[Y]$ could be computed by means of the formula

$$E[Y] = \int_{-\infty}^{+\infty} yf_Y(y)dy.$$

The question is this: Are the last two definitions consistent? Without the help of the abstract definition of expectation, such a question cannot be answered in general. Fortunately, the answer is yes. The theory tells us that all the special definitions of expectation given in the previous sections are consistent with (3.1).

3.2 Lebesgue's Theorems for Expectation

The most important technical results relative to expectation are the celebrated Lebesgue theorems, which give fairly general conditions under which the limit and expectation symbols

can be interchanged, i.e.,

$$E \left[\lim_{n \uparrow \infty} X_n \right] = \lim_{n \uparrow \infty} E[X_n]. \quad (3.8)$$

Theorem 3.1. Monotone Convergence Theorem.

Let $\{X_n\}_{n \geq 1}$ be a sequence of random variables such that for all $n \geq 1$,

$$0 \leq X_n \leq X_{n+1}, \text{ P-a.s.}$$

Then (3.8) holds.

Theorem 3.2. Dominated Convergence Theorem.

Let $\{X_n\}_{n \geq 1}$ be a sequence of random variables such that for all ω outside a set \mathcal{N} of null probability there exists $\lim_{n \uparrow \infty} X_n(\omega)$ and such that for all $n \geq 1$

$$|X_n| \leq Y, \text{ P-a.s.},$$

where Y is some integrable random variable. Then (3.8) holds.

Example 3.1. A Counterexample

We shall see that (3.8) is not always true when $\lim_{n \uparrow \infty} X_n$ exists. Indeed, take the following probabilistic model: $\Omega = [0, 1]$, and P is the Lebesgue measure on $[0, 1]$ (the probability of $[a, b] \subset [0, 1]$ is the length $b - a$ of this interval). Thus, ω is a real number in $[0, 1]$, and a random variable is a real function defined on $[0, 1]$. Take for X_n the function whose graph is a triangle with base $[0, \frac{2}{n}]$ and height n . Clearly, $\lim_{n \uparrow \infty} X_n(\omega) = 0$ and $E[X_n] = \int_0^1 X_n(x) dx = 1$, so that $E[\lim_{n \uparrow \infty} X_n] = 0 \neq \lim_{n \uparrow \infty} E[X_n] = 1$.

Example 3.2.

Let $\{S_n\}_{n \geq 1}$ be a sequence of nonnegative random variables. Then

$$E \left[\sum_{n=1}^{\infty} S_n \right] = \sum_{n=1}^{\infty} E[S_n]. \quad (3.9)$$

It suffices to apply the monotone convergence theorem, with $X_n = \sum_{k=1}^n S_k$. \diamond

Example 3.3.

Let $\{S_n\}_{n \geq 1}$ be a sequence of real random variables such that $\sum_{n \geq 1} E[|S_n|] < \infty$. Then (3.9) holds. It suffices to apply the dominated convergence theorem with $X_n = \sum_{k=1}^n S_k$ and $Y = \sum_{k=1}^{\infty} |S_k|$. By the result of Example 3.2, $E[Y] = \sum_{k=1}^{\infty} E[|S_k|] < \infty$. \diamond

Example 3.4. The Product Formula

As another illustration of the monotone convergence theorem we shall now prove the product formula

$$E[XY] = E[X]E[Y], \quad (3.10)$$

where X and Y are independent integrable random variables or independent nonnegative random variables. We do the proof in the nonnegative case, from which the integrable case follows easily, using the decomposition of a random variable into its nonnegative and nonpositive parts.

First notice that for any nonnegative random variable Z , if we define

$$Z_n = \sum_{k=1}^{n2^n-1} \frac{k}{2^n} \mathbf{1}_{\{\frac{k}{2^n} \leq Z < \frac{k+1}{2^n}\}} + n \mathbf{1}_{\{Z \geq n\}},$$

then $\lim_{n \uparrow \infty} Z_n = Z$, P-a.s. Thus, by the Lebesgue monotone convergence theorem, $\lim_{n \uparrow \infty} E[Z_n] = E[Z]$. With obvious notations, $\lim_{n \uparrow \infty} X_n Y_n = XY$, and therefore $\lim_{n \uparrow \infty} E[X_n Y_n] = E[XY]$. Also, $\lim_{n \uparrow \infty} E[X_n]E[Y_n] = E[X]E[Y]$. It is then clear that (3.10) will be proved if

$$E[X_n Y_n] = E[X_n]E[Y_n]$$

is shown to be true for all $n \geq 1$. But then, this verification amounts to observing that for all i, j ,

$$E[\mathbf{1}_{\{\frac{j}{2^n} < X \leq \frac{j+1}{2^n}\}} \mathbf{1}_{\{\frac{i}{2^n} < Y \leq \frac{i+1}{2^n}\}}] = E[\mathbf{1}_{\{\frac{j}{2^n} < X \leq \frac{j+1}{2^n}\}}] E[\mathbf{1}_{\{\frac{i}{2^n} < Y \leq \frac{i+1}{2^n}\}}],$$

that is

$$\begin{aligned} P\left(X \in \left(\frac{j}{2^n}, \frac{j+1}{2^n}\right], Y \in \left(\frac{i}{2^n}, \frac{i+1}{2^n}\right]\right) \\ = P\left(X \in \left(\frac{j}{2^n}, \frac{j+1}{2^n}\right]\right) P\left(Y \in \left(\frac{i}{2^n}, \frac{i+1}{2^n}\right]\right). \end{aligned}$$

The latter equality is just the independence of X and Y , and the proof is therefore completed. \diamond

Bibliography

- Aldous, D. and P. Diaconis, Shuffling cards and stopping times, *American Mathematical Monthly*, 93, 333–348, 1981.
- Aldous, D. and J.A. Fill, *Reversible Markov Chains and Random Walks on Graphs*, draft available at <http://www.stat.Berkeley.EDU/users/aldous/book.html>.
- Anderson, W.J., *Continuous Time Markov Chains: An Applications Oriented Approach*, Springer, New York, 1991.
- Anily, S. and A. Federgruen, Simulated annealing methods with general acceptance probabilities, *Journal of Applied Probability*, 24, 657–667, 1987.
- Asmussen, S., *Applied Probability and Queues*, Wiley, Chichester, 1987.
- Bartlett, M.S., On theoretical models for competitive and predatory biological systems, *Biometrika*, 44, 1957.
- Barker, A.A., Monte Carlo calculations of the radial distribution functions for a proton–electron plasma, *Austral. J. Phys.*, 18, 119–133, 1965.
- Baxter, R.J., *Exactly Solved Models in Statistical Mechanics*, Academic Press, London, 1982.
- Besag, J., Spatial interaction and the statistical analysis of lattice systems, *Journal of the Royal Statistical Society*, B-31, 192–236, 1974.
- Bhattacharya, R.N. and E.C. Waymire, *Stochastic Processes with Applications*, Wiley, New York, 1990.
- Capetanakis, J.I., Tree algorithm for packet broadcast channels, *IEEE Transactions on Information Theory*, IT-25, 505–515, 1979.
- Chung, K.L., *Markov Chains with Stationary Transition Probabilities*, Springer, Berlin, 2nd edition, 1967.
- Diaconis, P., *Group Representations in Probability and Statistics*, Institute of Mathematical Statistics, Hayward, California, 1988.

- Diaconis, P. and J.A. Fill, Strong stationary times via a new form of duality, *The Annals of Probability*, 18, 4, 1483–1522, 1991.
- Diaconis, P. and L. Saloff-Coste, What do we know about the Metropolis algorithm?, preprint, 1996.
- Diaconis, P. and D. Stroock, Geometric bounds for eigenvalues of Markov chains, *The Annals of Applied Probability*, 1, 1, 36–61, 1991.
- Doebelin, W., Sur les propriétés asymptotiques de mouvements régis par certains types de chaînes simples, *Bulletin Mathématique de la Société Roumaine des Sciences*, 39, No.1, 57–115, No.2, 3–61, 1937.
- Dobrushin, R.L., Central limit theorems for non-stationary Markov chains II, *Theory of Probability and its Applications*, 1, 329–383 (English translation), 1956.
- Dobrushin, R.L., Existence of a phase transition in two and three-dimensional Ising models, *Theory of Probability and its Applications*, 10, 193–213, 1965.
- Fayolle, G., Etude du comportement d'un canal radio partagé entre plusieurs utilisateurs, Thèse de Docteur-Ingénieur, Université Paris 6, 1975.
- Fayolle, G.; V.A. Malishev, and M.V. Menshikov, *Topics in the Constructive Theory of Countable Markov Chains*, Cambridge University Press, 1995.
- Fill, J.A., Eigenvalue bounds on convergence to stationarity for non-reversible Markov chains, with an application to the exclusion process, *The Annals of Applied Probability*, 1, 1, 62–87, 1991.
- Fill, J.A., An interruptible algorithm for perfect sampling via Markov chains, *Annals of Applied Probability*, 8, 131–162, 1998.
- Fishman, G.S., *Monte Carlo*, Springer, New York, 1996.
- Fréchet, M., *Recherches Théoriques Modernes sur le Calcul des Probabilités. Second livre*. Hermann, Paris, 1938.
- Frigessi, A., C.-R. Hwang and L. Younès, Optimal spectral properties of reversible stochastic matrices, Monte Carlo methods and the simulation of Markov random fields, *The Annals of Applied Probability*, 2, 3, 610–628, 1992.
- Gantmacher, F.R., *Applications of the Theory of Matrices*, Inter Science, New York, 1959.
- Geman, D., *Random Fields and Inverse Problems in Imaging*. In: *Lecture Notes in Mathematics*, 1427, 113–193, Springer, Berlin, 1990.
- Geman, S. and D. Geman, Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images, *IEEE Transactions of Pattern Analysis and Machine Intelligence*, 6, 721–741, 1984.
- Gibbs, W., *Elementary Principles of Statistical Mechanics*, Yale University Press, 1902.

- Glauber, R., Time-dependent statistics of the Ising model, *J. Math. Physics*, 4, 294–307.
- Gordon, W.J. and G.F. Newell, Closed queuing systems with exponential servers, *Operations Research*, 15, 224–25, 1967.
- Griffeath, D., A Maximal coupling for Markov chains, *Z. Wahrsch. Verw. Gebiete*, 31, 95–106, 1975.
- Grimmett, G.R., A Theorem on random fields, *Bulletin of the London Mathematical Society*, 81–84, 1973.
- Grimmett, G.R. and D.R. Stirzaker, *Probability and Random Processes*, Clarendon Press, Oxford, 1992.
- Grinstead, C.M. and J.L. Snell, *Introduction to Probability*, American Mathematical Society, 1997.
- Gross, D. and C.M. Harris, *Fundamentals of Queuing Theory*, Wiley, New York, 1975, 1984.
- Gut, A., *An Intermediate Course in Probability*, Springer, New York, 1995.
- Hajek, B., Cooling schedule for optimal annealing, *Mathematics of Operations Research* 13, 311–329, 1988.
- Hajnal, J., Weak ergodicity in nonhomogeneous Markov chains, *Proc. Cambridge Philos. Society*, 54, 236–246, 1958.
- Hammersley, J.M. and P. Clifford, Markov fields on finite graphs and lattices, unpublished manuscript, 1968.
- Hastings, W.K., Monte Carlo sampling methods using Markov chains and their applications, *Biometrika*, 57, 97–109, 1970.
- Horn, R.A. and C.R. Johnson, *Matrix Analysis*, Cambridge University Press, 1985.
- Iosifescu, I., *Finite Markov Chains and their Applications*, Wiley, Chichester, UK, 1980.
- Isaacson, D.L. and R.W. Madsen, *Markov Chains*, Wiley, New York, 1976.
- Ising, E., Beitrag zur Theorie des Ferromagnetismus, *Zeitschrift für Physik*, 31, 253–258, 1925.
- Jackson, J.R., Networks of waiting lines, *Operations Research*, 5, 518–521, 1957.
- Jerrum, M. and A. Sinclair, Approximating the permanent, *SIAM Journal of Computing*, 18, 1149–1178, 1989.
- Kac, M., *Random Walk and the Theory of Brownian Motion*, *American Mathematical Monthly*, 54, 369–391, 1947.
- Keilson, J., *Markov Chain Models, Rarity and Exponentiality*, Springer, New York, 1979.

- Karlin, S. and M.H. Taylor *A First Course in Stochastic Processes*, Academic Press, New York, 1975.
- Karlin, S. and M.H. Taylor *A Second Course in Stochastic Processes*, Academic Press, New York, 1981.
- Kelly, F.P., *Reversibility and Stochastic Networks*, Wiley, New York, 1979.
- Kelly, F.P. and B.D. Ripley, Markov point processes, *Journal of the London Mathematical Society*, 15, 188–192, 1977.
- Kemeny, J.G., Generalization of fundamental matrix, *Linear Algebra and its Applications*, 38, 193–206, 1991.
- Kinderman, R. and J.L. Snell, *Markov Random Fields and their Applications*, Contemporary Math., Vol. 1, Providence, Rhode Island: Amer. Math. Soc., 1980.
- Kirkpatrick, S., C.D. Gelatt, M.P. Vecchi, Optimization by simulated annealing, *Science*, 220, 671–680, 1982.
- Kolmogorov, A., Anfangsgründe der Theorie der Markoffschen Ketten mit unendlich vielen möglichen Zuständen, *Rec. Math. Moskov (Mat. Sbornik)*, 1, 43, 607–610, 1936.
- Lanford, O.E. and D. Ruelle, Observables at infinity and states with short range correlations in statistical mechanics, *Communications in Mathematical Physics*, 13, 194–215, 1969.
- Lindvall, T., Probabilistic proof of Blackwell's renewal theorem, *Annals of Probability*, 5, 482–485, 1977.
- Lindvall, T., *Lectures on the Coupling Method*, Wiley, New York, 1992.
- Liu, J., Eigenanalysis for a Metropolis sampling scheme with comparisons to rejection sampling and importance sampling, *Statistics and Computing*, to appear, 1995.
- Markov, A.A., Extension of the law of large numbers to dependent events (in Russian), *Bull. Soc. Phys. Math. Kazan*, 2, 15, 155–156, 1906.
- Metropolis, N.; M.N. Rosenbluth; A.W. Rosenbluth; A.H. Teller and E. Teller, Equations of state calculations by fast computing machines, *J. Chem. Phys.*, 21, 1087–92, 1953.
- Meyn, S.P. and R.L. Tweedie, *Markov Chains and Stochastic Stability*, Springer, New York, 1993.
- Mihaïl, M., Combinatorial aspects of expanders, Ph.D. dissertation, Department of Computer Science, Harvard University, 1989.
- Peskun, P.H., Optimum Monte Carlo sampling using Markov chains, *Biometrika*, 60, 3, 607–612, 1973.
- Peierls, R., On Ising's model of ferromagnetism, *Proceedings of the Cambridge Philosophical Society*, 32, 477–481, 1936.

- Pitman, J., Uniform rates of convergence for Markov chains transition probabilities, *Z. Wahrsch. Verw. Gebiete*, 29, 193–227, 1974.
- Revuz, D., *Markov Chains*, North-Holland, Amsterdam, 1984.
- Ripley, B.D., *Stochastic Simulation*, Wiley, New York, 1987.
- Robertazzi, T.G., *Computer Networks and Systems: Queuing Theory and Performance Evaluation*, Springer, New York, 1990.
- Rom, R. and M. Sidi, *Multiple Access Protocols, Performance and Analysis*, Springer, New York, 1990.
- Ross, S., *A Course in Simulation*, Macmillan, New York, 1987.
- Rudin, W., *Principles of Mathematical Analysis*, 3rd edition, McGraw–Hill, New York, 1976.
- Seneta, E., *Nonnegative Matrices and Markov Chains*, 2nd edition, Springer, New York, 1981.
- Shiryayev, A.N., *Probability*, Springer-Verlag, New York, 1987.
- Walrand, J., *An Introduction to Queuing Networks*, Prentice–Hall, Englewood Cliffs, New Jersey, 1988.
- Williams, D., *Probability and Martingales*, Cambridge University Press, 1991.
- Winkler, G., *Image Analysis, Random Fields and Dynamical Monte Carlo Methods*, Springer, New York, 1995.
- Wolff, R.W., *Stochastic Modeling and the Theory of Queues*, Prentice–Hall, Englewood Cliffs, New Jersey, 1989.
- Yan, D. and H. Mukai, Stochastic discrete optimization, *SIAM J. on Control and Optimization*, 30, 3, 549–612, 1992.

Author Index

- Aarts, 309
Aldous, 219, 225, 227
Anderson, 344
Anily, 244, 312
Asmussen, 365
Barker, 295
Besag, 268
Bhattacharya, 134
Burke, 397
Capetanakis, 174
Chung, 337
Clifford, 261
Cors, 309
Diaconis, 212, 219, 224, 225, 296, 298
Dobrushin, 235, 282
Erlang, 390
Fayolle, 170, 173
Federgruen, 244, 312
Fill, 211, 219, 224, 227
Fishman, 290
Foster, 167, 186
Frigessi, 302
Gantmacher, 198, 424
Gelatt, 309
Geman, 271, 311
Gordon, 410
Grimmett, 262
Grinstead, 229
Gross, 406
Hajek, 315
Hajnal, 241
Hammersley, 261
Harris, 406
Hastings, 294
Hirsch, 339
Horn, 424
Hwang, 302
Iosifescu, 195
Isaacson, 235, 247
Ising, 258
Jackson, 407
Johnson, 424
Kac, 136
Karlin, 195
Kelly, 274
Kemeny, 229
Kinderman, 285
Kirkpatrick, 309
Lanford, 282
Lindvall, 129, 136, 223
Liu, 298
Madsen, 235, 247
Malyshev, 170
Menshikov, 170
Metropolis, 294
Mihail, 211
Mihailov, 174
Mukai, 310
Newell, 410
Pakes, 169
Peierls, 283
Peskun, 299
Puterman, 184
Reich, 397
Ripley, 274, 290
Robertazzi, 412

- Rom, 178
Rosenbluth, 294
Ross, 290
Ruelle, 282
Saloff-Coste, 296, 298
Seneta, 198, 235
Shiryayev, 185, 429
Sidi, 178
Smale, 339
Snell, 229, 285
Stroock, 212
- Taylor, 195
Teller, 294
Titchmarsh, 420
Tsybahov, 174
Waymire, 134
Williams, 185, 429
Winkler, 280
Wolff, 392, 413
Yan, 310
Younès, 302

Subject Index

Abbreviations and Acronyms

a.c., 15
ALOHA, 108, 173
c.d.f., 14
c.f., 34
FIFO, 388
g.f., 30
GSMP, 385
HMC, 54
HPP, 324
i.i.d., 44
i.o., 42
LIFO, 392
M/G/1, 398
M/M/1/ ∞ , 388
M/M/K/0, 390
MAP, 275
MCMC, 293
p.d., 17
p.d.f., 15
SLE, 204
SLEM, 207
SLLN, 45

Abel's theorem, 419
absolutely continuous, 15, 25
absorbing chain, 155
absorption probability, 65, 156, 189
acceptance–rejection method, 292
almost-sure convergence, 43
aperiodic, 74

auto-Poisson model, 320
autobinomial model, 268

backward differential system, 339
backward recurrence time, 54, 147
Barker's sampler, 301
basic renewal equation, 138
Bayes's rules, 10
Bayes's sequential formula, 10
Bernoulli–Gaussian model, 278
binary tree protocol, 174
binomial random variable, 21
birth-and-death parameters, 338
birth-and-death process, 79, 338, 352
Borel's SLLN, 45
Borel–Cantelli lemma, 42
branching process, 61, 189
bus paradox, 147

Cauchy random variable, 34
cellular automaton, 190
Cesaro's lemma, 420
Chapman-Kolmogorov equation, 330
characteristic function, 34
Chebyshev's inequality, 45
chi-square contrast, 207
clique, 254
clock of a uniform Markov chain, 332
closed set, 71
collision-resolution protocol, 174
communication class, 71
competition theorem, 328
conditional expectation, 39
conditional independence, 12

- conditional probability, 9
 conductance, 215
 conservative, 336
 continuous-time HMC, 330
 contrast bound, 211
 convergence in variation, 128
 cooling, 309
 counting integral, 376
 counting process, 324
 coupling, 125
 coupling inequality, 129
 coupling method, 129
 coupling time, 129
 cumulative distribution function, 14
- data sequence, 138
 defective, 137
 defective renewal theorem, 142
 delayed regenerative theorem, 149
 detailed balance equation, 81
 Dirac sequence, 138
 Dirichlet form, 204
 Dirichlet problem, 183
 discrete random variable, 4
 distance in variation, 207
 DLR problem, 282
 Dobrushin's inequality, 236
 dominated convergence, 420, 431
 Dynkin's formula, 380
- eigenvalues, 424
 eigenvectors, 424
 embedded HMC, 348
 energy, 256
 equilibrium, 76
 ergodic, 360
 ergodic coefficient, 235
 ergodic theorem, 110
 Erlang's blocking formula, 395
 Erlang's waiting formula, 396
 essential, 349
 essential generator, 350
 event, 2
 excessive, 179
 exclusive and exhaustive causes, 10
- expectation, 17, 25, 428
 explosion, 351
 explosion time, 337
 exponential of a matrix, 426
 exponential random variable, 24
 extended stopping time, 370
- Fatou's lemma, 421
 first-step analysis, 65
 flip-flop, 331
 forward differential system, 339
 forward recurrence time, 147, 192
 Foster's theorem, 167
 fundamental matrix, 155, 226
 fundamental solution, 138
- Galton-Watson process, 61
 Gaussian random variable, 24
 generalized semi-Markov process, 385
 generating function, 30
 generator, 350
 geometric random variable, 21
 Gershgorin's bound, 427
 Gibbs distribution, 256
 Gibbs potential, 257
 Gibbs sampler, 286
 Gibbs-Markov equivalence, 261
 global balance equation, 75, 342
 Gordon-Newell network, 410
- Hammersley-Clifford theorem, 261
 harmonic, 179
 homogeneous Markov chain, 53
- independence, 7
 independent coupling, 130
 indicator function, 18
 infinite product, 422
 infinitesimal generator, 336
 intensity of a HPP, 324
 interrenewal sequence, 137
 invariant measure, 100
 inverse distribution method, 291
 irreducible, 72
- Jackson network, 407

- jump process, 337
- Kolmogorov's SLLN, 45
- lattice, 140
- lattice renewal theorem, 142
- local characteristics, 333
- local energy, 261
- local specification, 254
- Lyapunov function, 169

- Möbius formula, 262
- Markov coupling, 162
- Markov property, 12
- Markov random field, 254
- Markov's inequality, 45
- martingale, 179
- martingale convergence theorem, 185
- maximum principle, 180
- mean, 19
- mean return time, 104
- Metropolis sampler, 300
- monotone convergence, 421, 431
- Monte Carlo, 290
- Monte Carlo Markov chain, 293
- multinomial distribution, 28
- multiple-access, 108, 173, 174

- n-invariant measure, 357
- neighborhood, 254
- nonhomogeneous Markov chain, 239
- nonlattice, 140
- normalized potential, 264
- null recurrent, 96

- optimal control, 184
- optional sampling theorem, 185
- Orey's theorem, 131
- outcome, 1

- Pakes's lemma, 169
- period, 72
- permanent, 349
- Perron–Frobenius, 197
- point process, 323
- Poisson calculus, 375
- Poisson process, 324
- Poisson random variable, 22
- Poisson system, 383
- positive recurrent, 95
- positivity condition, 274
- potential matrix, 154
- potential matrix criterion, 97
- preemptive, 392
- preemptive resume, 392
- probability, 4
- probability density function, 15, 25
- probability space, 5
- processor sharing, 396
- product formula, 29, 431
- proper, 137
- pure delay, 395

- random vector, 25
- random walk, 59, 77, 88, 97, 105, 170
- randomized stopping time, 219
- rare events, 22
- real random variable, 4
- recurrent, 95
- regenerative cycle, 87
- regenerative processes, 145
- regenerative theorem, 147
- regular jump HMC, 337
- regular process, 337
- renewal distribution, 137
- renewal equation, 137
- renewal reward theorem, 117
- renewal sequence, 54, 137
- renewal theorem, 137
- renewal time, 137
- repair shop, 152, 170, 187
- retrodiction formula, 10
- return time, 83
- Reuter's criterion, 351
- reversal test, 81, 362
- reversibilization, 211
- reversible HMC, 81
- reversible transition matrix, 81, 201

- sample space, 1
- separation, 219

- sequential continuity, 6
- sigma-additivity, 5
- sigma-field, 3
- simulated annealing, 305
- Smith's formula, 147
- smoothing formula, 376
- snake chain, 90
- spectral theorem, 204
- stable, 336
- stationary coupling time, 223
- stationary distribution, 75, 342
- stationary distribution criterion, 104
- stationary regime, 76
- steady state, 76
- Stieltjes-Lebesgue integral, 15, 376, 430
- stochastic automaton, 63
- stochastic matrix, 54
- stochastic process, 53, 330
- stopping time, 83, 345
- strong law of large numbers, 45
- strong Markov property, 84, 345
- strong Markov property of HPPs, 370
- strong stationary time, 220
- strongly ergodic, 240
- sub- σ -additivity, 6
- subadditive function theorem, 423
- subharmonic, 179
- submartingale, 179
- subordinated chain, 332
- success runs, 96, 106
- successive conditioning, 41
- superharmonic, 179
- supermartingale, 179
- t-invariant measure, 357
- t-null recurrent, 357
- t-positive recurrent, 357
- telescope formula, 19
- time reversal, 81, 361
- traffic intensity, 395
- transfer matrix method, 260
- transient, 95
- transition graph, 56
- transition matrix, 54
- transition semigroup, 330
- transition time, 348
- Tychonov's theorem, 423
- uniform Markov chain, 332
- uniform random variable, 24
- urn of Ehrenfest, 64, 76, 82, 133
- variance, 19
- Wald's lemma, 20
- Watanabe's theorem, 380
- weakly ergodic, 240
- weighted path lower bound, 213
- weighted path upper bound, 213
- Zermelo's paradox, 133

Texts in Applied Mathematics

(continued from page ii)

31. *Brémaud*: Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues.
32. *Durrant*: Numerical Methods for Wave Equations in Geophysical Fluid Dynamics.
33. *Thomas*: Numerical Partial Differential Equations: Conservation Laws and Elliptic Equations.
34. *Chicone*: Ordinary Differential Equations with Applications.
35. *Kevorkian*: Partial Differential Equations: Analytical Solution Techniques, 2nd ed.
36. *Dullerud/Paganini*: A Course in Robust Control Theory: A Convex Approach.
37. *Quarteroni/Sacco/Saleri*: Numerical Mathematics.
38. *Gallier*: Geometric Methods and Applications: For Computer Science and Engineering.
39. *Atkinson/Han*: Theoretical Numerical Analysis: A Functional Analysis Framework.
40. *Brauer/Castillo-Chávez*: Mathematical Models in Population Biology and Epidemiology.
41. *Davies*: Integral Transforms and Their Applications, 3rd ed.
42. *Deufhard/Bornemann*: Scientific Computing with Ordinary Differential Equations.
43. *Deufhard/Hohmann*: Numerical Analysis in Modern Scientific Computing: An Introduction, 2nd ed.
44. *Knabner/Angermann*: Numerical Methods for Elliptic and Parabolic Partial Differential Equations.
45. *Larsson*: Partial Differential Equations with Numerical Methods.
46. *Pedregal*: Introduction to Optimization.