

## CAUSALITY

---

### Models, Reasoning, and Inference

Written by one of the preeminent researchers in the field, this book provides a comprehensive exposition of modern analysis of causation. It shows how causality has grown from a nebulous concept into a mathematical theory with significant applications in the fields of statistics, artificial intelligence, philosophy, cognitive science, and the health and social sciences.

The author presents and unifies the probabilistic, manipulative, counterfactual, and structural approaches to causation, and he devises simple mathematical tools for studying the relationships between causal connections and statistical associations. This book will open the way for including causal analysis in the standard curricula of statistics, artificial intelligence, epidemiology, business, social science, and economics. Students in these areas will find natural models, simple inferential procedures, and precise mathematical definitions of causal concepts that traditional texts have tended to evade or make unduly complicated.

*Causality* will be of interest to students and professionals in a wide variety of fields. Anyone who wishes to elucidate meaningful relationships from data, predict effects of actions and policies, assess explanations of reported events, or form theories of causal understanding and causal speech will find this book stimulating and invaluable.

Judea Pearl is Professor of Computer Science and Statistics and Director of the Cognitive Systems Laboratory at the University of California, Los Angeles. He is the author of *Heuristics* (1984) and *Probabilistic Reasoning in Intelligent Systems* (1988), and he has published close to 200 articles on various aspects of automated reasoning, learning, and inference. A Member of the National Academy of Engineering and a Fellow of the IEEE and the AAAI, Pearl is the recipient of the IJCAI Research Excellence Award in Artificial Intelligence (1999) “for his fundamental work on heuristic search, reasoning under uncertainty, and causality.”

### Additional Commendation for *Causality*

“Judea Pearl’s previous book, *Probabilistic Reasoning in Intelligent Systems*, was arguably the most influential book in Artificial Intelligence in the past decade, setting the stage for much of the current activity in probabilistic reasoning. In this book, Pearl turns his attention to causality, boldly arguing for the primacy of a notion long ignored in statistics and misunderstood and mistrusted in other disciplines, from physics to economics. He demystifies the notion, clarifies the basic concepts in terms of graphical models, and explains the source of many misunderstandings. This book should prove invaluable to researchers in artificial intelligence, statistics, economics, epidemiology, and philosophy, and, indeed, all those interested in the fundamental notion of causality. It may well prove to be one of the most influential books of the next decade.”

—Joseph Halpern, Computer Science Department, Cornell University

“This lucidly written book is full of inspiration and novel ideas that bring clarity to areas where confusion has prevailed, in particular concerning causal interpretation of structural equation systems, but also on concepts such as counterfactual reasoning and the general relation between causal thinking and graphical models. Finally the world can get a coherent exposition of these ideas that Judea Pearl has developed over a number of years and presented in a flurry of controversial yet illuminating articles.”

—Steffen L. Lauritzen, Department of Mathematics, Aalborg University

“Judea Pearl’s new book, *Causality: Models, Reasoning, and Inference*, is an outstanding contribution to the causality literature. It will be especially useful to students and practitioners of economics interested in policy analysis.”

— Halbert White, Professor of Economics, University of California, San Diego

“This book fulfills a long-standing need for a rigorous yet accessible treatise on the mathematics of causal inference. Judea Pearl has done a masterful job of describing the most important approaches and displaying their underlying logical unity. The book deserves to be read by all statisticians and scientists who use nonexperimental data to study causation, and would serve well as a graduate or advanced undergraduate course text.”

— Sander Greenland, School of Public Health, University of California, Los Angeles

“Judea Pearl has written an account of recent advances in the modeling of probability and cause, substantial parts of which are due to him and his co-workers. This is essential reading for anyone interested in causality.”

— Brian Skryms, Department of Philosophy, University of California, Irvine

Cambridge University Press  
978-0-521-77362-1 - Causality: Models, Reasoning, and Inference  
Judea Pearl  
Frontmatter  
[More information](#)

---

# CAUSALITY

---

## Models, Reasoning, and Inference

**Judea Pearl**

*University of California, Los Angeles*



**CAMBRIDGE**  
UNIVERSITY PRESS

Cambridge University Press  
978-0-521-77362-1 - Causality: Models, Reasoning, and Inference  
Judea Pearl  
Frontmatter  
[More information](#)

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE  
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS  
The Edinburgh Building, Cambridge CB2 2RU, UK  
40 West 20th Street, New York, NY 10011-4211, USA  
10 Stamford Road, Oakleigh, VIC 3166, Australia  
Ruiz de Alarcón 13, 28014 Madrid, Spain  
Dock House, The Waterfront, Cape Town 8001, South Africa  
<http://www.cambridge.org>

© Judea Pearl 2000

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

First published 2000  
Reprinted with corrections 2001

Printed in the United States of America

*Typeface* Times 10.25/13 *System* AMS-TEX [FH]

*A catalog record for this book is available from the British Library*

*Library of Congress Cataloging in Publication Data*

Pearl, Judea.

Causality : models, reasoning, and inference / Judea Pearl.

p. cm.

ISBN 0-521-77362-8 (hardback)

I. Causation. 2. Probabilities. I. Title.

BD541.P43 2000

122 – dc21

99-042108

CIP

ISBN 0 521 77362 8 hardback

Cambridge University Press  
978-0-521-77362-1 - Causality: Models, Reasoning, and Inference  
Judea Pearl  
Frontmatter  
[More information](#)

---

*Development of Western science is based on two great achievements: the invention of the formal logical system (in Euclidean geometry) by the Greek philosophers, and the discovery of the possibility to find out causal relationships by systematic experiment (during the Renaissance).*

Albert Einstein (1953)

*TO RUTH*

## Contents

<i>Preface</i>	<i>page</i> xiii
<b>1 Introduction to Probabilities, Graphs, and Causal Models</b>	<b>1</b>
1.1 Introduction to Probability Theory	1
1.1.1 Why Probabilities?	1
1.1.2 Basic Concepts in Probability Theory	2
1.1.3 Combining Predictive and Diagnostic Supports	6
1.1.4 Random Variables and Expectations	8
1.1.5 Conditional Independence and Graphoids	11
1.2 Graphs and Probabilities	12
1.2.1 Graphical Notation and Terminology	12
1.2.2 Bayesian Networks	13
1.2.3 The $d$ -Separation Criterion	16
1.2.4 Inference with Bayesian Networks	20
1.3 Causal Bayesian Networks	21
1.3.1 Causal Networks as Oracles for Interventions	22
1.3.2 Causal Relationships and Their Stability	24
1.4 Functional Causal Models	26
1.4.1 Structural Equations	27
1.4.2 Probabilistic Predictions in Causal Models	30
1.4.3 Interventions and Causal Effects in Functional Models	32
1.4.4 Counterfactuals in Functional Models	33
1.5 Causal versus Statistical Terminology	38
<b>2 A Theory of Inferred Causation</b>	<b>41</b>
2.1 Introduction	42
2.2 The Causal Modeling Framework	43
2.3 Model Preference (Occam's Razor)	45
2.4 Stable Distributions	48
2.5 Recovering DAG Structures	49
2.6 Recovering Latent Structures	51
2.7 Local Criteria for Causal Relations	54

2.8	Nontemporal Causation and Statistical Time	57
2.9	Conclusions	59
2.9.1	On Minimality, Markov, and Stability	61
<b>3</b>	<b>Causal Diagrams and the Identification of Causal Effects</b>	<b>65</b>
3.1	Introduction	66
3.2	Intervention in Markovian Models	68
3.2.1	Graphs as Models of Interventions	68
3.2.2	Interventions as Variables	70
3.2.3	Computing the Effect of Interventions	72
3.2.4	Identification of Causal Quantities	77
3.3	Controlling Confounding Bias	78
3.3.1	The Back-Door Criterion	79
3.3.2	The Front-Door Criterion	81
3.3.3	Example: Smoking and the Genotype Theory	83
3.4	A Calculus of Intervention	85
3.4.1	Preliminary Notation	85
3.4.2	Inference Rules	85
3.4.3	Symbolic Derivation of Causal Effects: An Example	86
3.4.4	Causal Inference by Surrogate Experiments	88
3.5	Graphical Tests of Identifiability	89
3.5.1	Identifying Models	91
3.5.2	Nonidentifying Models	93
3.6	Discussion	94
3.6.1	Qualifications and Extensions	94
3.6.2	Diagrams as a Mathematical Language	96
3.6.3	Translation from Graphs to Potential Outcomes	98
3.6.4	Relations to Robins's <i>G</i> -Estimation	102
<b>4</b>	<b>Actions, Plans, and Direct Effects</b>	<b>107</b>
4.1	Introduction	108
4.1.1	Actions, Acts, and Probabilities	108
4.1.2	Actions in Decision Analysis	110
4.1.3	Actions and Counterfactuals	112
4.2	Conditional Actions and Stochastic Policies	113
4.3	When Is the Effect of an Action Identifiable?	114
4.3.1	Graphical Conditions for Identification	114
4.3.2	Remarks on Efficiency	116
4.3.3	Deriving a Closed-Form Expression for Control Queries	117
4.3.4	Summary	118
4.4	The Identification of Plans	118
4.4.1	Motivation	118
4.4.2	Plan Identification: Notation and Assumptions	120
4.4.3	Plan Identification: A General Criterion	121
4.4.4	Plan Identification: A Procedure	124

<b>Contents</b>	ix
4.5 Direct Effects and Their Identification	126
4.5.1 Direct versus Total Effects	126
4.5.2 Direct Effects, Definition, and Identification	127
4.5.3 Example: Sex Discrimination in College Admission	128
4.5.4 Average Direct Effects	130
<b>5 Causality and Structural Models in Social Science and Economics</b>	133
5.1 Introduction	134
5.1.1 Causality in Search of a Language	134
5.1.2 SEM: How its Meaning Became Obscured	135
5.1.3 Graphs as a Mathematical Language	138
5.2 Graphs and Model Testing	140
5.2.1 The Testable Implications of Structural Models	140
5.2.2 Testing the Testable	144
5.2.3 Model Equivalence	145
5.3 Graphs and Identifiability	149
5.3.1 Parameter Identification in Linear Models	149
5.3.2 Comparison to Nonparametric Identification	154
5.3.3 Causal Effects: The Interventional Interpretation of Structural Equation Models	157
5.4 Some Conceptual Underpinnings	159
5.4.1 What Do Structural Parameters Really Mean?	159
5.4.2 Interpretation of Effect Decomposition	163
5.4.3 Exogeneity, Superexogeneity, and Other Frills	165
5.5 Conclusion	170
<b>6 Simpson's Paradox, Confounding, and Collapsibility</b>	173
6.1 Simpson's Paradox: An Anatomy	174
6.1.1 A Tale of a Non-Paradox	174
6.1.2 A Tale of Statistical Agony	175
6.1.3 Causality versus Exchangeability	177
6.1.4 A Paradox Resolved (Or: What Kind of Machine Is Man?)	180
6.2 Why There Is No Statistical Test for Confounding, Why Many Think There Is, and Why They Are Almost Right	182
6.2.1 Introduction	182
6.2.2 Causal and Associational Definitions	184
6.3 How the Associational Criterion Fails	185
6.3.1 Failing Sufficiency via Marginality	185
6.3.2 Failing Sufficiency via Closed-World Assumptions	186
6.3.3 Failing Necessity via Barren Proxies	186
6.3.4 Failing Necessity via Incidental Cancellations	188
6.4 Stable versus Incidental Unbiasedness	189
6.4.1 Motivation	189
6.4.2 Formal Definitions	191
6.4.3 Operational Test for Stable No-Confounding	192



6.5	Confounding, Collapsibility, and Exchangeability	193
6.5.1	Confounding and Collapsibility	193
6.5.2	Counfounding versus Confounders	194
6.5.3	Exchangeability versus Structural Analysis of Confounding	196
6.6	Conclusions	199
<b>7</b>	<b>The Logic of Structure-Based Counterfactuals</b>	<b>201</b>
7.1	Structural Model Semantics	202
7.1.1	Definitions: Causal Models, Actions, and Counterfactuals	202
7.1.2	Evaluating Counterfactuals: Deterministic Analysis	207
7.1.3	Evaluating Counterfactuals: Probabilistic Analysis	212
7.1.4	The Twin Network Method	213
7.2	Applications and Interpretation of Structural Models	215
7.2.1	Policy Analysis in Linear Econometric Models: An Example	215
7.2.2	The Empirical Content of Counterfactuals	217
7.2.3	Causal Explanations, Utterances, and Their Interpretation	221
7.2.4	From Mechanisms to Actions to Causation	223
7.2.5	Simon's Causal Ordering	226
7.3	Axiomatic Characterization	228
7.3.1	The Axioms of Structural Counterfactuals	228
7.3.2	Causal Effects from Counterfactual Logic: An Example	231
7.3.3	Axioms of Causal Relevance	234
7.4	Structural and Similarity-Based Counterfactuals	238
7.4.1	Relations to Lewis's Counterfactuals	238
7.4.2	Axiomatic Comparison	240
7.4.3	Imaging versus Conditioning	242
7.4.4	Relations to the Neyman–Rubin Framework	243
7.4.5	Exogeneity Revisited: Counterfactual and Graphical Definitions	245
7.5	Structural versus Probabilistic Causality	249
7.5.1	The Reliance on Temporal Ordering	249
7.5.2	The Perils of Circularity	250
7.5.3	The Closed-World Assumption	252
7.5.4	Singular versus General Causes	253
7.5.5	Summary	256
<b>8</b>	<b>Imperfect Experiments: Bounding Effects and Counterfactuals</b>	<b>259</b>
8.1	Introduction	259
8.1.1	Imperfect and Indirect Experiments	259
8.1.2	Noncompliance and Intent to Treat	261
8.2	Bounding Causal Effects	262
8.2.1	Problem Formulation	262
8.2.2	The Evolution of Potential-Response Variables	263
8.2.3	Linear Programming Formulation	266

<b>Contents</b>	xi
8.2.4 The Natural Bounds	268
8.2.5 Effect of Treatment on the Treated	269
8.2.6 Example: The Effect of Cholestyramine	270
8.3 Counterfactuals and Legal Responsibility	271
8.4 A Test for Instruments	274
8.5 Causal Inference from Finite Samples	275
8.5.1 Gibbs Sampling	275
8.5.2 The Effects of Sample Size and Prior Distribution	277
8.5.3 Causal Effects from Clinical Data with Imperfect Compliance	277
8.5.4 Bayesian Estimate of Single-Event Causation	280
8.6 Conclusion	281
<b>9 Probability of Causation: Interpretation and Identification</b>	<b>283</b>
9.1 Introduction	283
9.2 Necessary and Sufficient Causes: Conditions of Identification	286
9.2.1 Definitions, Notation, and Basic Relationships	286
9.2.2 Bounds and Basic Relationships under Exogeneity	289
9.2.3 Identifiability under Monotonicity and Exogeneity	291
9.2.4 Identifiability under Monotonicity and Nonexogeneity	293
9.3 Examples and Applications	296
9.3.1 Example 1: Betting against a Fair Coin	297
9.3.2 Example 2: The Firing Squad	297
9.3.3 Example 3: The Effect of Radiation on Leukemia	299
9.3.4 Example 4: Legal Responsibility from Experimental and Nonexperimental Data	302
9.3.5 Summary of Results	303
9.4 Identification in Nonmonotonic Models	304
9.5 Conclusions	307
<b>10 The Actual Cause</b>	<b>309</b>
10.1 Introduction: The Insufficiency of Necessary Causation	309
10.1.1 Singular Causes Revisited	309
10.1.2 Preemption and the Role of Structural Information	311
10.1.3 Overdetermination and Quasi-Dependence	313
10.1.4 Mackie's INUS Condition	313
10.2 Production, Dependence, and Sustenance	316
10.3 Causal Beams and Sustenance-Based Causation	318
10.3.1 Causal Beams: Definitions and Implications	318
10.3.2 Examples: From Disjunction to General Formulas	320
10.3.3 Beams, Preemption, and the Probability of Single-Event Causation	322
10.3.4 Path-Switching Causation	324
10.3.5 Temporal Preemption	325
10.4 Conclusions	327

Cambridge University Press  
978-0-521-77362-1 - Causality: Models, Reasoning, and Inference  
Judea Pearl  
Frontmatter  
[More information](#)

---

xii

**Contents****Epilogue The Art and Science of Cause and Effect**

A public lecture delivered November 1996 as part of the UCLA Faculty Research Lectureship Program	331
<i>Bibliography</i>	359
<i>Name Index</i>	375
<i>Subject Index</i>	379

## Preface

The central aim of many studies in the physical, behavioral, social, and biological sciences is the elucidation of cause–effect relationships among variables or events. However, the appropriate methodology for extracting such relationships from data – or even from theories – has been fiercely debated.

The two fundamental questions of causality are: (1) What empirical evidence is required for legitimate inference of cause–effect relationships? (2) Given that we are willing to accept causal information about a phenomenon, what inferences can we draw from such information, and how? These questions have been without satisfactory answers in part because we have not had a clear semantics for causal claims and in part because we have not had effective mathematical tools for casting causal questions or deriving causal answers.

In the last decade, owing partly to advances in graphical models, causality has undergone a major transformation: from a concept shrouded in mystery into a mathematical object with well-defined semantics and well-founded logic. Paradoxes and controversies have been resolved, slippery concepts have been explicated, and practical problems relying on causal information that long were regarded as either metaphysical or unmanageable can now be solved using elementary mathematics. Put simply, causality has been mathematized.

This book provides a systematic account of this causal transformation, addressed primarily to readers in the fields of statistics, artificial intelligence, philosophy, cognitive science, and the health and social sciences. Following a description of the conceptual and mathematical advances in causal inference, the book emphasizes practical methods for elucidating potentially causal relationships from data, deriving causal relationships from combinations of knowledge and data, predicting the effects of actions and policies, evaluating explanations for observed events and scenarios, and – more generally – identifying and explicating the assumptions needed for substantiating causal claims.

Ten years ago, when I began writing *Probabilistic Reasoning in Intelligent Systems* (1988), I was working within the empiricist tradition. In this tradition, probabilistic relationships constitute the foundations of human knowledge, whereas causality simply provides useful ways of abbreviating and organizing intricate patterns of probabilistic relationships. Today, my view is quite different. I now take causal relationships to be the

fundamental building blocks both of physical reality and of human understanding of that reality, and I regard probabilistic relationships as but the surface phenomena of the causal machinery that underlies and propels our understanding of the world.

Accordingly, I see no greater impediment to scientific progress than the prevailing practice of focusing all of our mathematical resources on probabilistic and statistical inferences while leaving causal considerations to the mercy of intuition and good judgment. Thus I have tried in this book to present mathematical tools that handle causal relationships side by side with probabilistic relationships. The prerequisites are startlingly simple, the results embarrassingly straightforward. No more than basic skills in probability theory and some familiarity with graphs are needed for the reader to begin solving causal problems that are too complex for the unaided intellect. Using simple extensions of probability calculus, the reader will be able to determine mathematically what effects an intervention might have, what measurements are appropriate for control of confounding, how to exploit measurements that lie on the causal pathways, how to trade one set of measurements for another, and how to estimate the probability that one event was the actual cause of another.

Expert knowledge of logic and probability is nowhere assumed in this book, but some general knowledge in these areas is beneficial. Thus, Chapter 1 includes a summary of the elementary background in probability theory and graph notation needed for the understanding of this book, together with an outline of the developments of the last decade in graphical models and causal diagrams. This chapter describes the basic paradigms, defines the major problems, and points readers to the chapters that provide solutions to those problems.

Subsequent chapters include introductions that serve both to orient the reader and to facilitate skipping; they indicate safe detours around mathematically advanced topics, specific applications, and other explorations of interest primarily to the specialist.

The sequence of discussion follows more or less the chronological order by which our team at UCLA has tackled these topics, thus re-creating for the reader some of our excitement that accompanied these developments. Following the introductory chapter (Chapter 1), we start with the hardest questions of how one can go about discovering cause–effect relationships in raw data (Chapter 2) and what guarantees one can give to ensure the validity of the relationships thus discovered. We then proceed to questions of identifiability – namely, predicting the direct and indirect effects of actions and policies from a combination of data and fragmentary knowledge of where causal relationships might operate (Chapters 3 and 4). The implications of these findings for the social and health sciences are then discussed in Chapters 5 and 6 (respectively), where we examine the concepts of structural equations and confounding. Chapter 7 offers a formal theory of counterfactuals and structural models, followed by a discussion and a unification of related approaches in philosophy, statistics, and economics. The applications of counterfactual analysis are then pursued in Chapters 8–10, where we develop methods of bounding causal relationships and illustrate applications to imperfect experiments, legal responsibility, and the probability of necessary, sufficient, and single-event causation. We end this book (Epilogue) with a transcript of a public lecture that I presented at UCLA, which provides a gentle introduction of the historical and conceptual aspects of causation.

## Preface

xv

Readers who wish to be first introduced to the nonmathematical aspects of causation are advised to start with the Epilogue and then to sweep through the other historical/conceptual parts of the book: Sections 1.1.1, 3.3.3, 4.5.3, 5.1, 5.4.1, 6.1, 7.2, 7.4, 7.5, 8.3, 9.1, 9.3, and 10.1. More formally driven readers, who may be anxious to delve directly into the mathematical aspects and computational tools, are advised to start with Section 7.1 and then to proceed as follows for tool building: Section 1.2, Chapter 3, Sections 4.2–4.4, Sections 5.2–5.3, Sections 6.2–6.3, Section 7.3, and Chapters 8–10.

I owe a great debt to many people who assisted me with this work. First, I would like to thank the members of the Cognitive Systems Laboratory at UCLA, whose work and ideas formed the basis of many of these sections: Alex Balke, Blai Bonet, David Chickering, Adnan Darwiche, Rina Dechter, David Galles, Hector Geffner, Dan Geiger, Moisés Goldszmidt, Jin Kim, Jin Tian, and Thomas Verma. Tom and Dan have proven some of the most basic theorems in causal graphs; Hector, Adnan, and Moisés were responsible for keeping me in line with the logicist approach to actions and change; and Alex and David have taught me that counterfactuals are simpler than the name may imply.

My academic and professional colleagues have been very generous with their time and ideas as I began ploughing the peaceful territories of statistics, economics, epidemiology, philosophy, and the social sciences. My mentors—listeners in statistics have been Phil Dawid, Steffen Lauritzen, Don Rubin, Art Dempster, David Freedman, and David Cox. In economics, I have benefited from many discussions with John Aldrich, Kevin Hoover, James Heckman, Ed Leamer, and Herbert Simon. My forays into epidemiology resulted in a most fortunate and productive collaboration with Sander Greenland and James Robins. Philosophical debates with James Woodward, Nancy Cartwright, Brian Skyrms, Clark Glymour, and Peter Spirtes have sharpened my thinking of causality in and outside philosophy. Finally, in artificial intelligence, I have benefited from discussions with and the encouragement of Nils Nilsson, Ray Reiter, Don Michie, Joe Halpern, and David Heckerman.

The National Science Foundation deserves acknowledgment for consistently and faithfully sponsoring the research that led to these results, with special thanks to H. Moraff, Y. T. Chien, and Larry Reeker. Other sponsors include Abraham Waksman of the Air Force Office of Scientific Research, Michael Shneier of the Office and Naval Research, the California MICRO Program, Northrop Corporation, Rockwell International, Hewlett-Packard, and Microsoft.

I would like to thank Academic Press and Morgan Kaufmann Publishers for their kind permission to reprint selected portions of previously published material. Chapter 3 includes material reprinted from *Biometrika*, vol. 82, Judea Pearl, “Causal Diagrams for Empirical Research,” pp. 669–710, Copyright 1995, with permission from Oxford University Press. Chapter 5 includes material reprinted from *Sociological Methods and Research*, vol. 27, Judea Pearl, “Graphs, Causality, and Structural Equation Models,” pp. 226–84, Copyright 1998, with permission from Sage Publications, Inc. Chapter 7 includes material reprinted from *Foundations of Science*, vol. 1, David Galles and Judea Pearl, “An Axiomatic Characterization of Causal Counterfactuals,” pp. 151–82, Copyright 1998, with permission from Kluwer Academic Publishers. Chapter 7 also includes material reprinted from *Artificial Intelligence*, vol. 97, David Galles and Judea Pearl, “Axioms

of Causal Relevance,” pp. 9–43, Copyright 1997, with permission from Elsevier Science. Chapter 8 includes material modified from *Journal of the American Statistical Association*, vol. 92, Alexander Balke and Judea Pearl, “Bounds on Treatment Effects from Studies with Imperfect Compliance,” pp. 1171–6, Copyright 1997, with permission from the American Statistical Association.

The manuscript was most diligently typed, processed, and illustrated by Kaoru Mulvihill. Jin Tian and Blai Bonet helped in proofing selected chapters. Matt Darnell did a masterful job of copyediting these pages. Alan Harvey has been my consoling ombudsman and virtual editor throughout the production process.

Finally, my humor and endurance through the writing of this book owe a great debt to my family – to Tammy, Danny, Michelle, and Leora for filling my heart with their smiles, and to my wife Ruth for surrounding me with so much love, support, and meaning.

J. P.  
Los Angeles  
August 1999