THE ART OF CAUSAL CONJECTURE

Glenn Shafer

Preface

In this book, I present a new mathematical and philosophical foundation for probability and show how this foundation can help us understand causality. In doing so, I touch on many fields of scholarly endeavor, including statistics, mathematical and applied probability, artificial intelligence, and philosophy.

The various disciplines that use causal reasoning differ in the relative weight they put on security and precision of knowledge as opposed to timeliness of action. In the natural and social sciences, we often seek high levels of certainty in the identification of causes and high levels of precision in the measurement of their effects. In the practical sciences—engineering, medicine, and business—we must often act on causal conjectures that are based on more limited evidence. In artificial intelligence, action must be taken with even less deliberation. Yet these disciplines all work with the same concept of causality. This common concept is the topic of this book.

I have chosen to publish the book in MIT Press's Artificial Intelligence series because I believe that the discipline of artificial intelligence is currently in the best position to put the book's ideas to use. Because the older sciences—pure and practical—all demand a relatively high level of precision and security, they can successfully study only causal structures that are sufficiently uniform to permit relatively precise statistical estimation, and these structures can often be described adequately in the existing language of stochastic processes, without any explicit use of causal ideas. The newer discipline of artificial intelligence, which has been left on its own to deal with problems in which statistical estimation can play only a more limited role, has a more widely and keenly felt need for a conceptual framework for causal reasoning. This book has been inspired by the debate about causal reasoning in artificial intelligence, and it provides a foundation that will facilitate the further development of probabilistic expert systems based on causal models.

In the longer run, I believe that the ideas developed here will also become fundamental to the discipline of statistics. Although causal ideas permeate the use of statistics in all branches of industry, commerce, government, and science, they have not been incorporated into the basic theory of the subject. This book shows how causal ideas can be as central in statistical theory as they are in statistical practice. It does not challenge the maxim that causation cannot be proven from statistics alone, but by bringing causal ideas into the foundations of probability, it creates a framework within which causal conjectures can be more clearly quantified, debated, and confronted by statistical evidence.

The book may also contribute, in a different way, to probability as a branch of mathematics. By introducing dynamic and causal ideas into probability at the most elementary level, it liberates some basic ideas—especially the idea of a martingale—from the ponderous framework of measure theory. I believe that in the long run this will lead to simpler and hence more widely accessible expositions of some of the most important and useful branches of mathematical probability.

Finally, the book contributes to the philosophy of probability and causality. It contributes to the centuries-old debate about the meaning of probability by integrating frequency and belief into a single story in which an observer's knowledge develops step by step. It contributes to the even older debate about the meaning of causality by showing how nature can be thought of as such an observer, so that causes are steps in the development of nature's knowledge. This reveals the initmate relationship between causality and probability, and it provides a starting point for wider understanding and use of ideas on probabilistic causality that have been explored in English-speaking philosophy since Hans Reichenbach's posthumous book, *The Direction of Time* (1954).

Readers not trained in mathematical probability or mathematical statistics may fear that knowledge of existing theory in these disciplines is a prerequisite for understanding the book. This is not the case. The book's topics are quite different from those usually studied by probabilists and statisticians, and the basic ideas in the book are developed in detail from first principles. Topics from the standard theory of probability and statistics are treated in a series of appendices, but the book is designed to be read without a mastery of these appendices. I include the appendices primarily to help readers who already do know something of the standard theory. These readers will inevitably try to understand the book's new ideas in terms of the standard theory, and this process will provide many opportunities for confusion and misunderstanding. In order to minimize this confusion and misunderstanding, I have tried to explain the relation between the new ideas and standard ideas as fully as possible, and the appendices have helped me to do this without, I hope, putting obstacles in the way of readers not so interested in the standard ideas.

I have organized the book so as to emphasize the simplicity of its basic idea: that causes can be represented as steps in a probability tree. This idea leads to a great variety of causal concepts, and even to a fair amount of mathematics. But I have tried to develop the complexities step by step, so that the overall simplicity will not be obscured. In "Wide Road," Piet Hein pokes fun at scholars who make easy things hard.

> To make a name for learning
> when other roads are barred,
> take something very easy
> and make it very hard.[1]

My aim has been the opposite: to take causality, something thought very mysterious, and make it very simple.

The careful and sometimes leisurely tone of the book may give some readers the impression that it is an exposition of ideas previously published in journals or more austere monographs. In fact, virtually all the theory in the body of the book (as opposed to the appendices) is published here for the first time. I had planned earlier to publish the main ideas of Chapters 5 through 10 as an article, but I gradually became convinced that an adequate account of these ideas and their implications required a book-length presentation

In the course of writing the book, I myself have been surprised by the breadth and potential of its ideas. I am now convinced that many of topics I have touched upon, especially in the final chapters, deserve books of their own:

---

[1] Quoted with permission from page 33 of Hein's *Grooks II*, published by Borgens Forlag, Copenhagen, in 1973.

- The abstract theory of event trees and martingale trees, sketched in Chapters 11 and 12, should be extended to serve as a foundation for and generalization of standard advanced probability theory.
- The account of refining and grounding in Chapter 13 should be expanded to a more general account of the relations among probability trees for different observers, and hence to a foundation for game theory.
- The ideas of Chapter 14 should be expanded to deal with practical issues that arise in the design of experiments and observational studies.
- The ideas of Chapter 15 should be applied to examples of causal models in a variety of disciplines, from agronomy to sociology.
- Chapter 16 should be expanded to an account of computation in probability and decision trees.

I hope to contribute to some of these enterprises, and I hope many others will do so as well.

As I complete the book, I am keenly aware that even as an introduction to its topic, it will fail to meet the needs of many in the very broad audience to which it is addressed. Each of the many disciplines that stand to gain from the book would gain even more from a book tailored to its own preoccupations.

- Mathematical statisticians, once they have grasped the basic ideas, will want more concise mathematical treatments, which go on to deal with statistical testing and estimation. They will also want to see the probability-tree account of causality compared in detail with the Rubin-Holland account.
- Econometricians will want to relate the concepts of independence tracking, and sign to concepts that have been studied in their literature, including exogeneity and Granger causation.
- Other social scientists will ask whether the precise interpretations the probability-tree account provides for causal models should encourage or discourage the use of these models in their domains of study.
- Philosophers will want a thorough comparison of the probability-tree account with philosophical accounts based on counterfactuals, and they will want to see how the probability-tree account can deal with puzzles that have been studied in the philosophical literature.
- Computer scientists will want to hear less about interpretation and more about implementation.

I hope that in due time each discipline will have its own introduction to causality in probability trees. I also hope that the ideas in this book will survive to serve as common ground for all who must deal with causality in an uncertain world.

Acknowledgements

This book grew out of a seminar on causality in the Department of Philosophy of Princeton University, where I spent part of the 1992-93 academic year. I would like to thank the participants in the seminar, Dick DeVeaux, Adam Grove, Gil Harman, Paul Holland, Dick Jeffrey, and Bas van Fraassen, for their encouragement, criticism, and suggestions. Each of them made a distinctive and significant contribution to the shape of the book. Paul Holland's insistence on the need for further explication of relevance diagrams was especially seminal.

At later stages, a long list of other colleagues have provided crucial advice and encouragement. Steffen Lauritzen and Alex Kogan have been constant sources of ideas, encouragement, and gentle criticism. Judea Pearl has provided both encouragement and challenge; together with Paul Holland, he convinced me I should address experimentation and intervention as well as conditional independence. Phil Dawid and Volodya Vovk, whose own work inspired many of the ideas in the book, have been equally encouraging. David Krantz and John Tukey gave me detailed advice based on close readings of intermediate drafts. Mike Mahoney helped me improve my translation of the Huygens's manuscript. Jan von Plato introduced me to Martin-Löf's work on type theory. In its final stages, the manuscript profited greatly from close readings by Volodya Vovk and Wayne Oldford.

The book has also profited from suggestions by a number of students in classes where I have had the opportunity to use some of its ideas, especially Mindy Callen, Chingfu Chang, Jill Giles, Richard Jones, Chunyan Li, Susan Sly, Miranda Wai, and T.J. Wang at Rutgers and Peter Barsoom, Julian Dierkes, Brian Gibbs, and Becky Pettit at Princeton.

I have also received advice, encouragement, and support from many other colleagues from around the world. These include Nuel Belnap, Nancy Cartwright, Erhan Cinlar, Paul Cohen, Art Dempster, Saurav Dutta, Peter Fishburn, Doug Fisher, Clark Glymour, Russ Greiner, Allan Gut, David Heckerman, Paul Humphreys, Harsharanjeet Jagpal, Radim Jirousek, Richard Jones, Stephen Judd, Sanjeev Kulkarni, Dennis Lindley, Francesco Malvestuto, Per Martin-Löf, Stephen Omohundro, Suresh Radhadkrishnan, Eric Ristad, James Robins, Judy Roitman, Paul Rosenbaum, Uri Ronnen, Richard Scheines, Prakash Shenoy, Peter Spirtes, Wolfgang Spohn, Howard Wainer, Bruce Western, Peter Williams, and Richard Wolpert. In addition to making the book better, my interaction with these colleagues has made writing the book much more enjoyable.

I have also enjoyed my interactions with my editors at MIT Press, Fiona Stevens and Bob Prior. I appreciated their encouragement when the manuscript seemed too short and their patience as it began to seem too long.

Finally, I want to thank Nell Irvin Painter. As my wife, she has encouraged and supported the personal effort that has been required to complete this book. As an extraordinarily capable and experienced scholar and writer, she has been a source of invaluable advice.

The foundation of probability on which this book relies is the result of research stretching over two decades. This research has been generously supported by a series of grants from the National Science Foundation, as well as a Guggenheim fellowship in 1983-84, a fellowship at the Center for Advanced Study in the Behavioral Sciences in 1988-89, the Ronald G. Harper chair in the School of Business at the University of Kansas, which I held from 1988 to 1992, and research funds for 1992 to 1995 from the Graduate School of Management at Rutgers University.

The research on the probabilistic aspects of causality reported on in the book was carried out during from 1992 through 1995, with support from Rutgers and from grant SBE9213674 from the National Science Foundation. Of the many National Science Foundation grants I have held during the past twenty years, only this one, with its emphasis on a long-term project, made this kind of work possible. I want to thank Ron

Overman, director of the Science, Technology, and Society Program, for his dedicated service there.

The book would also not have been possible without the generous support provided by Rutgers University and Princeton University. I would like to thank the leadership of both institutions, especially Phipps Arabie, George Benson, George Farris, Peter Hammer, Dan Palmon and Miklos Vasarhelyi at Rutgers and Paul Benacerraf and Ruth Simmons at Princeton, for their support and their contributions to the collegial intellectual environments from which this book has benefited.