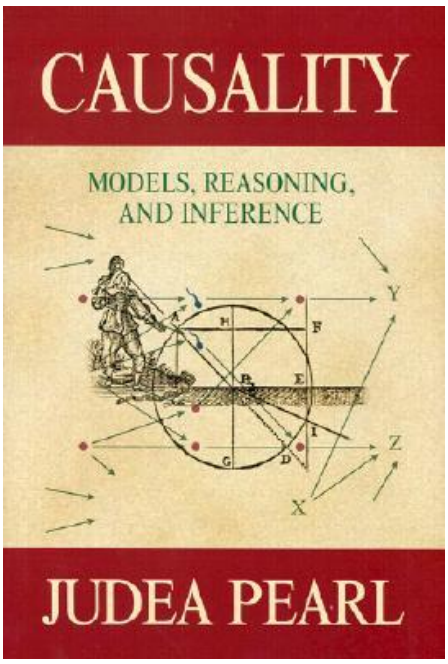

Causality: Models, Reasoning, and Inference

Judea Pearl

Causality: Models, Reasoning, and Inference

Judea Pearl

Causality: Models, Reasoning, and Inference Judea Pearl
Written by one of the pre-eminent 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. Pearl presents a unified account of the probabilistic, manipulative, counterfactual and structural approaches to causation, and devises simple mathematical tools for analyzing the relationships between causal connections, statistical associations, actions and observations. The book will open the way for including causal analysis in the standard curriculum of statistics, artifical intelligence, business, epidemiology, social science and economics. Students in these areas will find natural models, simple identification procedures, and precise mathematical definitions of causal concepts that traditional texts have tended to evade or make unduly complicated. This book will be of interest to professionals and students 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. Professor of Computer Science at the UCLA, Judea Pearl is the winner of the 2008 Benjamin Franklin Award in Computers and Cognitive Science.

Causality: Models, Reasoning, and Inference Details

Date	:	Published March 13th 2000 by Cambridge University Press
ISBN	:	9780521773621
Author	:	Judea Pearl
Format	:	Hardcover 400 pages
Genre	:	Philosophy, Science, Nonfiction, Mathematics

From Reader Review Causality: Models, Reasoning, and Inference for online ebook

Thomas Eapen says

The classic modern reference on the science and philosophy of causality. However, it can be a challenging read for those who are not familiar with probabilistic models.

David Sundahl says

In the future people will regard this as on the same level as Newton’s “Principia” or Frege’s “Begriffsscrift.”

Chris says

Doesn't answer the question in the title.

Leonardo says

fucking "back door" y "front door" =)

Michael Nielsen says

Historically, it's a strange fact that we developed probability and statistics without also developing a theory of causality. Such a theory would dramatically change science. This book summarizes recent attempts by Pearl and others to develop such a theory. I don't think the theory is complete, but this is a great prelude.

John Ledesma says

A Note On “Causality: Models, Reasoning, and Inference” by Judea Pearl
By Dr. Alex Liu

August 2005 ***

This is a note on my reading Judea Pearl’s book “Causality: Models, Reasoning, and Inference” 1999 Cambridge University Press.

Even it sounds like the book is creating a NEW paradigm of conducting causal research,to many empirical

scholars including me; the main purpose of this book is to:

1) Develop graphical tools in assisting causal analysis
2) Develop a non-linear and non-parametric extension of SEM
3) Discuss about causality
4) Develop an algorithm using partial correlations to discover causal structure under certain assumptions

However, all the above has already made this book a must read for people in empirical research methods. The author made a lot of effort to convince the statistics community for the acceptance of his ideas. I think that is a wrong approach. His work is more useful to people using statistics for empirical research, than to statisticians.

Experts of research methods often say that “research methods do not equal to statistics”. Research methods equal statistics plus something else. Pearl’s work is to formalize this “something else” and provide tools to work on them explicitly. In other words, Pearl’s work can help us processing statistical results for causal analysis, but not much to improve statistical analysis.

In traditional empirical analysis, at least in the mainstream methods teaching, this “something else” for causal analysis is that variable A is a cause of variable B, if:

(1) A and B are correlated.
(2) The association arises because A causes B and not vice versa due to temporal or logical or theoretical reasons.
(3) The association between A and B is not spurious.

It seems to me that at least three parts of Pearl work are worth studying and even being applied to some empirical research projects.

(1) His work of explicitly defining the “something else”
(2) His work of formally representing them
(3) His work of developing rules and tools for us to handle them

After gaining a full understanding of the above three items, I think that we can use Pearl’s work to assist our causal analysis in empirical research.

According to Pearl, statistics deals with mean, variance, correlation, regression, dependence, conditional independence, association, likelihood, collapsibility, risk ratio, odd ratio, marginalization, conditionalization, “controlling for”, … While causal analysis deals with randomization, influence, effect, confounding, “holding constant”, disturbance, spurious correlation, instrumental variables, intervention, explanation, attribution, … The second part minus the first part is the “something else”.

Professor Pearl’s language to formally represent causal analysis and its components include both structural equation models (linear, nonlinear and nonparametric) and graphical diagrams. Pearl uses do(x) to represent intervention. As many methodologists will agree, with Pearl’s work, method concepts like spuriousness and confounding, are
much better formalized than ever before.

His proposed rules include criterion to select covariates for adjustment, intervention calculus, and counterfactual analysis. Professor Pearl also proposed IC* algorithm to discover causal structures.

These are good contributions made by Pearl's work. But, this is just a beginning. In general, I think there are more questions than answers in this book. There are also many missing links we need to bridge, in order to conduct a good causal analysis. For example, indirect effects are not covered as much as the direct effects and total effects. How to
estimate the strength of a causal influence is also left out.

D.A. Freedman of UC Berkeley takes a different view than that of Pearl (Freedman 2004). Many scholars including Freedman mentioned that Pearl did not do any modeling or empirical work, but just talked causation mathematically or philosophically, that may not be a fair comment as theoretical discussion along can be very valuable. Due to this, Freedman claims that Pearl's work is based on many assumptions that are unrealistic and difficulty to confirm in applied research. Freedman claims that Pearl acknowledged some of these assumptions like in page 83 of his book, but did not make all them clear.

Published in 1993 (2nd edition in 2000 by MIT Press), the book Causation, Prediction and Search by Spirtes, Glymour, and Scheines (SGS) is worth reading as they actually developed a software for their developed algorithms and applied to a lot of real research. Between SGS and Freedman, there are also many dialogues in discussing whether the work from statistical evidence to causal inference can be automated without any needs for subject knowledge.

Actually, both the algorithms developed by Pearl and SGS do not work well. Professor Freedman of UC Berkeley claims these algorithms do not work as they are based on false assumptions. As I know, quite many scholars including myself tried these algorithms on some empirical data, and found these algorithms often lead us to nowhere or to some
errors. However, many ideas presented in these algorithms can be used, in combination with subject knowledge and other statistical methods like structural equation modeling method, to aid us in generating hypotheses and also in testing fitted models. Professor Bill Shipley has some good work along this line (Shipley 2000).

In general, I believe to successfully infer causality from statistical evidence like correlation does require some subject knowledge, additional statistical methods and hard work. But, the work of Pearl and SGS can help to improve the current practice greatly.

Reference

Freedman D. A. 2004 Statistical Models for Causation, UCB Statistics Technical Report No. 651 www.stat.berkeley.edu/-census/651.pdf

Freedman D. A. 1998 Are There Algorithms That Discover Causal Structure, UCB Statistics Technical Report No. 514

Halpern, Joseph Y. and Judea Pearl 2001 Cause and Explanations: A Structural-Model Approach, in Proceedings of the Seventeenth Conference on Uncertainy inartificial Intelligence, San Francisco, CA: Morgan Kaufmann

Heckerman, David and Ross Shachter 1995 Decision-theoretic foundation for causal reasoning, Journal of Artificial Intelligence Research 3: 405-430

Pearl, Judea 2003 Statistics and Causal Inference: A Review, Test (2003) Vol. 12, No. 2:281-345

Newberg, Leland Gerson 2003 Review of Causality Econometric Theory, 19, 2003: 675-685

Shipley, Bill 2000 Cause and Correlation in Biology: A User's Guide to Path Analysis, Structural Equations and Causal Inference. Cambridge University Press

Spirtes, P., Glymour, C., and Scheines, R. 1993 Causation, Prediction, and Search. Springer Lecture Notes in Statistics, no. 81, Springer-Verlag, New York. 2nd edn. 2000, MIT Press

***

This note was written when Alex worked in IBM Research from Dec 2004 to April 2005, then was modified in August 2005. The author benefited from discussion on this matter with Dr. Sunil Noronha and Joseph Kramer of IBM Research.

For further work of Dr. Alex Liu on this subject, please visit below for his book ~ From Model Building to Model Mapping:

http://www.researchmethods.org/modeli...

Or visit below for the RM software where causality reasoning and techniques have been incorporated.

http://www.researchmethods.org/rmplat...

Makoto says

very hard to get all of the way through. I think I actually only got 3/4 of the way through in the end...

Moshe says

You really can infer causation from correlation (with a few caveats).

Delhi Irc says

Location: GG6 IRC
Accession no: DL026784

Terran M says

The field of causal inference is important and deserves more attention than it usually gets. I had hoped that this book, which promises to be about "causality: models, reasoning, and inference" would be a comprehensive treatment of the topic. I was badly disappointed.

The book suffers both from decisions about what to include and from the writing. It turns out that Pearl has not actually attempted to provide a comprehensive treatment of the field of causal inference at all, but rather

of *his own* contributions to it — which, while substantial, are narrow and mathematical. He devotes all of four pages to inferring the causal graph from data, and then the rest of the book is predicated on having a complete, unambiguous causal graph; this makes the book irrelevant for empirical work.

Within the scope of what it covers, Pearl's writing is mediocre; he is not a master of exposition, and he offers frequent pot-shots at those with whom he has a professional disagreement. He accepts none of the responsibility for presenting his work in a fairly inaccessible way, and seems to have a grudge that the world has not done more to adopt it. I respect Pearl as a researcher, but he is a poor writer.

What this book is really about is Pearl's mathematical "do-calculus", and how, given a complete causal graph, it can be used to rigorously state what it means to intervene or to assess a counterfactual. This is a valuable contribution, but most empirical practitioners will not require a book-length treatment of this narrow aspect of the field.

For a brief introduction to using causal graphs to select your controls, see Chapter 17 of "Statistical Modeling - A Fresh Approach". That chapter is available free from the author at http://www.mosaic-web.org/go/Statisti...

For more about inferring causal graphs from the data, look for a series of papers by Colombo and Maathuis at ETH Zurich.

For an alternative book which is of more practical relevance for most purposes, you might consider Mostly Harmless Econometrics: An Empiricist's Companion

---

Ari says

The first few chapters are full of ideas, and I found the graphical model of causality a powerful conceptual tool. This is the premiere exposition of that view.

The wife, who is a statistics graduate student, is more skeptical and thinks that other models are as good or better.

I read about half of it; the rest was too technical for my state of mind and needs.
