CAUSAL INFERENCE

for Statistics, Social, and Biomedical Sciences An Introduction

Most questions in social and biomedical sciences are causal in nature: what would happen to individuals, or to groups, if part of their environment were changed? In this groundbreaking text, two world-renowned experts present statistical methods for studying such questions.

This book starts with the notion of potential outcomes, each corresponding to the outcome that would be realized if a subject were exposed to a particular treatment or regime. In this approach, causal effects are comparisons of such potential outcomes. The fundamental problem of causal inference is that we can observe only one of the potential outcomes for a particular subject. The authors discuss how randomized experiments allow us to assess causal effects and then turn to observational studies. They lay out the assumptions needed for causal inference and describe the leading analysis methods, including matching, propensity-score methods, and instrumental variables. Many detailed applications are included, with special focus on practical aspects for the empirical researcher.

Guido W. Imbens is Applied Econometrics Professor and Professor of Economics at the Graduate School of Business at Stanford University. He has previously held tenured positions at the University of California at Los Angeles, the University of California at Berkeley, and Harvard University. He is a Fellow of the Econometric Society and the American Academy of Arts and Sciences. He holds an honorary doctorate from the University of St. Gallen, Switzerland. Imbens has done extensive research in econometrics and statistics, specializing in causal inference. He has published widely in leading economics and statistics journals, including the *American Economic Review*, *Econometrica*, the *Review of Economic Studies*, the *Journal of the American Statistical Association*, the *Annals of Statistics*, *Biometrika*, and the *Journal of the Royal Statistical Society, Series A*.

Donald B. Rubin is John L. Loeb Professor of Statistics at Harvard University, where he has been professor since 1983 and Department chair for 13 of those years. He has been elected to be a Fellow/Member/Honorary Member of the Woodrow Wilson Society, the Guggenheim Memorial Foundation, the Alexander von Humboldt Foundation, the American Statistical Association, the Institute of Mathematical Statistics, the International Statistical Institute, the American Association for the Advancement of Science, the American Academy of Arts and Sciences, the European Association of Methodology, the British Academy, and the U.S. National Academy of Sciences. As of 2014, he has authored/coauthored nearly 400 publications (including ten books), has five joint patents, and for many years has been one of the most highly cited authors in the world, with currently over 150,000 citations (Google Scholar). He is also the recipient of honorary doctorates from Otto Friedrich University, Bamberg, Germany; the University of Ljubljana, Slovenia; and Universidad Santa Tómas, Bogotá, Colombia, and has been named an honorary professor at four universities.

Advance Praise for *Causal Inference for Statistics, Social, and Biomedical Sciences*

"This thorough and comprehensive book uses the 'potential outcomes' approach to connect the breadth of theory of causal inference to the real-world analyses that are the foundation of evidence-based decision making in medicine, public policy, and many other fields. Imbens and Rubin provide unprecedented guidance for designing research on causal relationships, and for interpreting the results of that research appropriately."

– Dr. Mark McClellan, Director of the Health Care Innovation and Value Initiative, the Brookings Institution

"Clarity of thinking about causality is of central importance in financial decision making. Imbens and Rubin provide a rigorous foundation allowing practitioners to learn from the pioneers in the field."

– Dr. Stephen Blyth, Managing Director, Head of Public Markets, Harvard Management Company

"A masterful account of the potential outcomes approach to causal inference from observational studies that Rubin has been developing since he pioneered it 40 years ago."

– Adrian Raftery, Blumstein-Jordan Professor of Statistics and Sociology, University of Washington

"Correctly drawing causal inferences is critical in many important applications. Congratulations to Professors Imbens and Rubin, who have drawn on their decades of research in this area, along with the work of several others, to produce this impressive book covering concepts, theory, methods, and applications. I especially appreciate their clear exposition of conceptual issues, which are important to understand in the context of either a designed experiment or an observational study, and their use of real applications to motivate the methods described."

– Nathaniel Schenker, Former President of the American Statistical Association

CAUSAL INFERENCE

for Statistics, Social, and Biomedical Sciences An Introduction

Guido W. Imbens

Stanford University

Donald B. Rubin

Harvard University

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning, and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9780521885881

First published 2015
Reprinted 2016

Printed in the United States of America by Sheridan Books, Inc.

A catalog record for this publication is available from the British Library.

Library of Congress Cataloging in Publication Data
Imbens, Guido.
Causal inference for statistics, social, and biomedical sciences : an introduction / Guido W. Imbens & Donald B. Rubin.
pages cm
1. Social sciences–Research. 2. Causation. 3. Inference. I. Rubin, Donald B. II. Title.
H62.I537 2014
519.5′4–dc23 2014020988

ISBN 978-0-521-88588-1 Hardback

Cambridge University Press
978-0-521-88588-1 — Causal Inference for Statistics, Social, and Biomedical Sciences
Guido W. Imbens , Donald B. Rubin
Frontmatter
More Information

To my parents, Annie Imbens-Fransen and Gerard Imbens, for all their support and encouragement over the years; to my children, Carleton, Annalise and Sylvia, who have provided so much joy in recent years; and to my wife and best friend, Susan Athey.

Guido W. Imbens

To my family, colleagues, and students, and Elizabeth Zell, who has provided ten years of crucial support.

Donald B. Rubin

Contents

Preface

In many applications of statistics, a large proportion of the questions of interest are fundamentally questions of causality rather than simply questions of description or association. For example, a medical researcher may wish to find out whether a new drug is effective against a disease. An economist may be interested in uncovering the effects of a job-training program on an individual's employment prospects, or the effects of a new tax or regulation on economic activity. A sociologist may be concerned about the effects of divorce on children's subsequent education. In this text we discuss statistical methods for studying such questions.

The book arose out of a conversation we had in 1992 while we were both on the faculty at Harvard University. We found that although we were both interested in questions of causality, we had difficulty communicating our ideas because, coming from different disciplines, we were used to different terminology and conventions. However, the excitement about the ideas in these different areas motivated us to capitalize on these difficulties, which led to a long collaboration, including research projects, graduate and undergraduate teaching, and thesis advising. The book is a reflection of this collaboration.

The book is based directly on many semester and quarter-length courses we, initially jointly, and later separately, taught for a number of years, starting in 1995 at Harvard University, followed by the University of California at Los Angeles, the University of California at Berkeley, and Stanford University, to audiences of graduate and undergraduate students from statistics, economics, business, and other disciplines using applied statistics. In addition, we have taught shorter versions of such courses in Barcelona, Beijing, Berlin, Bern, Geneva, Maastricht, Mexico City, Miami, Montevideo, Santiago, Stockholm, Uppsala, Wuppertal, Zurich, and at the World Bank as well as other associations and agencies.

There are a number of key features of the approach taken in this book. First of all, the perspective we take is that all causal questions are tied to specific interventions or treatments. Second, causal questions are viewed as comparisons of potential outcomes, with each potential outcome corresponding to a level of the treatment. Each of these potential outcomes could have been observed had the treatment taken on the corresponding level. After the treatment has taken on a specific level, only the potential outcome corresponding to that level is realized and can actually be observed. Causal effects involve the comparison of the outcome actually observed with other potential outcomes that could have been observed had the treatment taken on a different level, but that are not, in fact, observed. Causal inference is therefore fundamentally a missing data problem and, as in all missing data problems, a key role is played by the mechanism that determines which data values are observed and which are missing. In causal inference, this mechanism is referred to as the assignment mechanism, the mechanism that determines levels of the treatment taken by the units studied.

The book is organized in seven parts. In the first part we set out the basic philosophy underlying our approach to causal inference and describe the potential outcomes framework. The next three parts of the book are distinguished by the assumptions maintained about the assignment mechanism. In Part II we assume that the assignment mechanism corresponds to a classical randomized experiment.

In Part III we assume that the assignment mechanism is "regular" in a well-defined sense, which generalizes randomized experiments. In this part of the book we discuss what we call the "design" phase of an observational study, which we view as extremely important for credible conclusions. In the next part, Part IV, we discuss data analysis for studies with regular assignment mechanisms. Here we consider matching and subclassification procedures, as well as model-based and weighting methods.

In Part V we relax this regularity assumption and discuss more general assignment mechanisms. First we assess the key assumption required for regularity, unconfoundedness. We also explore in this part of the text sensitivity analyses where we relax some of the key features of a regular assignment mechanism.

Next, in Part VI of the text, we consider settings where the assignment mechanism is regular, but compliance with the assignment is imperfect. As a result, the probability of receipt of treatment may depend on both observed and unobserved characteristics and outcomes of the units. To address these complications, we turn to instrumental variables methods. Part VII of the book concludes.

As with all books, ours has limitations. Foremost is our focus on binary treatments. Although many of the results can easily be extended to multi-valued treatments, we focus on the binary treatment case because many critical conceptual issues arise already in that setting. Second, throughout most of the book we make the "stability" assumption that treatments applied to one unit do not affect outcomes for other units and that there are no unrepresented versions of the treatments. There is a growing literature on interactions through networks and peer effects that builds on the notions of causality discussed in this book. Finally, although we designed the book to be theoretically tight and principled, we focus on practical rather than mathematical results, including detailed applications with real data sets, consistent with our target audience of researchers in applied fields.

ACKNOWLEDGMENTS

We are deeply indebted to many collaborators, colleagues, and students from whom we have learned much about the topics discussed in this book over the years. These include Alberto Abadie, Nikola Andric, Joshua Angrist, Susan Athey, Thomas Barrios, David Card, Matias Cattaneo, Gary Chamberlain, Raj Chetty, Rajeev Dehejia, Mike Dickstein,

Peng Ding, David Drukker, Dean Eckles, Valeria Espinosa, Avi Feller, Sergio Firpo, Andrew Gelman, Paul Gift, Paul Goldsmith-Pinkham, Bryan Graham, Roee Gutman, Gary Harvey, Jane Herr, Kei Hirano, Nate Holt, Joe Hotz, Wilbert van der Klaauw, Jacob Klerman, Alan Krueger, Tony Lancaster, Joseph Lee, Victoria Liublinska, Chuck Manski, Eric Maskin, Fabrizia Mealli, Kari Lock Morgan, Eduardo Morales, Julie Mortimer, Eleanor Murray, Whitney Newey, Cassandra Wolos Pattanayak, Jack Porter, Geert Ridder, Paul Rosenbaum, Bruce Sacerdote, Hal Stern, Elizabeth Stuart, Neal Thomas, Sadek Wahba, David Watson, Jeff Wooldridge, Alan Zaslavsky, and Elizabeth Zell.

We are grateful for stimulating discussions over the years with David Freedman, Sander Greenland, James Heckman, Nick Jewell, Mark van der Laan, Judea Pearl, Jamie Robins, Jasjeet Sekhon, Jeffrey Smith, and Edward Vytlacil.

We are also grateful to Joshua Angrist, David Card, Raj Chetty, Rajeev Dehejia, Esther Duflo, Brad Efron, Alan Krueger, Robert Lalonde, June Reinisch, Bruce Sacerdote, Sadek Wahba, and Scott Zeger for making their data available.

Finally, we would like to extend our gratitude to Dean Eckles and Karthik Rajkumar for comments on the first edition, and to Michael Pollman for a remarkably thorough check for errors and typos.