CONTENTS

Introduction

There are three broad branches of decision analysis: normative, descriptive, and prescriptive. Normative analysis seeks to establish ideal properties of decision making, often aiming to give meaning to the terms "optimal" and "rational." Descriptive analysis seeks to understand and predict how actual decision makers behave. Prescriptive analysis seeks to improve the performance of actual decision making.

One might view normative and descriptive analysis as entirely distinct subjects. It is not possible, however, to cleanly separate prescriptive analysis from the other branches of study. Prescriptive analysis aims to improve actual decisions, so it must draw on normative thinking to define "improve" and on descriptive research to characterize actual decisions.

This book offers prescriptive analysis that seeks to improve patient care. My focus is decision making under uncertainty regarding patient health status and response to treatment. By "uncertainty," I do not just mean that clinicians and health planners may make probabilistic rather than definite predictions of patient outcomes. My main concern is decision making when the available evidence and medical knowledge do not suffice to yield precise probabilistic predictions.

For example, an educated patient who is comfortable with probabilistic thinking may ask her clinician a seemingly straightforward question such as "What is the chance that I will develop disease X in the next five years?" or "What is the chance that treatment Y will cure me?" Yet the clinician may not be able to provide precise answers to these questions. A credible response may be a range, say "20 to 40 percent" or "at least 50 percent."

Decision theorists use the terms "deep uncertainty" and "ambiguity" to describe the decision settings I address, but I shall encompass them within the broader term "uncertainty" for now. Uncertainty in patient care is common and has sometimes been acknowledged verbally. For example, the Evidence-Based Medicine Working Group asserts that (Institute of Medicine, 2011, p. 33): "clinicians must accept uncertainty and the notion that clinical decisions are often made with scant knowledge of their true impact." However, uncertainty has generally not been addressed in research on evidence-based medicine, which has been grounded in classical statistical theory. I think this a huge omission, which this book strives to correct.

Surveillance or Aggressive Treatment

I pay considerable attention to the large class of decisions that choose between surveillance and aggressive treatment of patients at risk of potential disease. Consider, for example, women at risk of breast cancer. In this instance, surveillance typically means undergoing periodic mammograms and clinical exams, while aggressive treatment may mean preventive drug treatment or mastectomy.

Other familiar examples are choice between surveillance and drug treatment for patients at risk of heart disease or diabetes. Yet others are choice between surveillance and aggressive treatment of patients who have been treated for localized cancer and are at risk of metastasis. A semantically distinct but logically equivalent decision is choice between diagnosis of patients as healthy or ill. With diagnosis, the concern is not to judge whether a patient will develop a disease in the future but whether the patient is currently ill and requires treatment.

These decisions are common, important to health, and familiar to clinicians and patients alike. Indeed, patients make their own choices related to surveillance and aggressive treatment. They perform self-surveillance by monitoring their own health status. They choose how faithfully to adhere to surveillance schedules and treatment regimens prescribed by clinicians.

Uncertainty often looms large when a clinician contemplates choice between surveillance and aggressive treatment. The effectiveness of surveillance in mitigating the risk of disease may depend on the degree to which a patient will adhere to the schedule of clinic visits prescribed in a surveillance plan. Aggressive treatment may be more beneficial than surveillance to the extent that it reduces the risk of disease development or

the severity of disease that does develop. It may be more harmful to the extent that it generates health side effects and financial costs beyond those associated with surveillance. There often is substantial uncertainty about all these matters.

Evolution of the Book

I am an economist with specialization in econometrics. I have no formal training in medicine. One may naturally ask how I developed an interest in patient care under uncertainty and feel able to contribute to the subject. It would be arrogant and foolhardy for me to dispense medical advice regarding specific aspects of patient care. I will not do so. The contributions that I feel able to make concern the methodology of evidence-based medicine. This matter lies within the expertise of econometricians, statisticians, and decision analysts.

Research on treatment response and risk assessment shares a common objective: probabilistic prediction of patient outcomes given knowledge of observed patient attributes. Development of methodology for prediction of outcomes conditional on observed attributes has long been a core concern of many academic disciplines.

Econometricians and statisticians commonly refer to conditional prediction as *regression*, a term in use since the nineteenth century. Some psychologists have used the terms *actuarial prediction* and *statistical prediction*. Computer scientists may refer to *machine learning* and *artificial intelligence*. Researchers in business schools may speak of *predictive analytics*. All these terms are used to describe methods that have been developed to enable conditional prediction.

As an econometrician, I have studied how statistical imprecision and identification problems affect empirical (or evidence-based) research that uses sample data to predict population outcomes. Statistical theory characterizes the imprecise inferences that can be drawn about the outcome distribution in a study population by observing the outcomes of a finite sample of its members. Identification problems are inferential difficulties that persist even when sample size grows without bound.

A classic example of statistical imprecision occurs when one draws a random sample of a population and uses the sample average of an outcome to estimate the population mean outcome. Statisticians typically measure imprecision of the estimate by its variance, which decreases to zero as sample size increases. Whether imprecision is measured by variance or another

way, the famous "Laws of Large Numbers" imply that imprecision vanishes as sample size increases.

Identification problems encompass the spectrum of issues that are sometimes called *non-sampling errors* or *data-quality problems*. These issues cannot be resolved by amassing so-called big data. They may be mitigated by collecting better data, but not by merely collecting more data.

A classic example of an identification problem is generated by missing data. Suppose that one draws a random sample of a population, but one observes only some sample outcomes. Increasing sample size adds new observations, but it also yields further missing data. Unless one learns the values of the missing data or knows the process that generates missing data, one cannot precisely learn the population mean outcome as sample size increases.

My research has focused mainly on identification problems, which often are the dominant difficulty in empirical research. I have studied probabilistic prediction of outcomes when available data are combined with relatively weak assumptions that have some claim to credibility. While much of this work has necessarily been technical, I have persistently stressed the simple truth that research cannot yield decision-relevant findings based on evidence alone.

In Manski (2013a) I observed that the logic of empirical inference is summarized by the relationship:

$$\text{assumptions} + \text{data} \Rightarrow \text{conclusions}.$$

Data (or evidence) alone do not suffice to draw useful conclusions. Inference also requires assumptions (or theories, hypotheses, premises, suppositions) that relate the data to the population of interest. Holding fixed the available data, and presuming avoidance of errors in logic, stronger assumptions yield stronger conclusions. At the extreme, one may achieve certitude by posing sufficiently strong assumptions. A fundamental difficulty of empirical research is to decide what assumptions to maintain.

Strong conclusions are desirable, so one may be tempted to maintain strong assumptions. I have emphasized that there is a tension between the strength of assumptions and their credibility, calling this (Manski, 2003, p. 1):

> *The Law of Decreasing Credibility*: The credibility of inference decreases with the strength of the assumptions maintained.

This "Law" implies that analysts face a dilemma as they decide what assumptions to maintain: Stronger assumptions yield conclusions that are more powerful but less credible.

I have argued against making precise probabilistic predictions with *incredible certitude*. It has been common for experts to assert that some event will occur with a precisely stated probability. However, such predictions often are fragile, resting on unsupported assumptions and limited data. Thus, the expressed certitude is not credible.

Motivated by these broad ideas, I have studied many prediction problems and have repeatedly found that empirical research may be able to credibly bound the probability that an event will occur but not make credible precise probabilistic predictions, even with large data samples. In econometrics jargon, probabilities of future events may be *partially identified* rather than *point identified*. This work, which began in the late 1980s, has been published in numerous journal articles and synthesized in multiple books, written at successive stages of my research program and at technical levels suitable for different audiences (Manski, 1995, 2003, 2005, 2007a, 2013a).

Whereas my early research focused on probabilistic prediction per se, I have over time extended its scope to study decision making under uncertainty; that is, decisions when credible precise probabilistic predictions are not available. Thus, my research has expanded from econometrics to prescriptive decision analysis.

Elementary decision theory suggests a two-step process for choice under uncertainty. Considering the feasible alternatives, the first step is to eliminate dominated actions—an action is dominated if one knows for sure that some other action is superior. The second step is to choose an undominated action. This is subtle because there is no consensus regarding the optimal way to choose among undominated alternatives. There are only various reasonable ways. I will later give content to the word "reasonable."

Decision theory is mathematically rigorous, but it can appear sterile when presented in abstraction. The subject comes alive when applied to important actual decision problems. I have studied various public and private decisions under uncertainty. This work has yielded technical research articles and a book on public policy under uncertainty written for a broad audience (Manski, 2013a).

I have increasingly felt that patient care is ripe for study as a problem of decision making under uncertainty. I therefore have sought to learn enough about research on evidence-based medicine to make original contributions that build on my methodological background in econometrics

and decision analysis. The results include studies of diagnostic testing and treatment under uncertainty (Manski, 2009, 2013b), personalized care with partial assessment of health risks (Manski, 2018a), analysis and design of randomized clinical trials (Manski, 2004a; Manski and Tetenov, 2016, 2019), drug approval (Manski, 2009), and vaccination policy with partial knowledge of disease transmission (Manski, 2010, 2017). I have also written a review article (Manski, 2018b).

The idea of writing a book evolved as I have accumulated background in evidence-based medicine and have developed an enlarging set of original research findings. A book provides the space to present major themes and to show how they become manifest in various contexts. A book enables an author to speak to a broader audience than is possible when writing research articles on particular topics.

I hope that this book will prove useful to a spectrum of readers. I would like it to help clinicians and public health planners recognize and cope with uncertainty as they make decisions about patient care. It may help patients to become informed about and participate in their own care. I anticipate that the book will help medical researchers design randomized trials and interpret the evidence they obtain from trials and observational studies. I will be pleased if the book encourages the biostatisticians who assist medical researchers to make constructive use of modern methodological advances in econometrics and statistical decision theory.

Some readers with certain types of expertise will correctly view the book as critical of the methodologies they have advocated. These include biostatisticians who have used the statistical theory of hypothesis testing to advise medical researchers on the design and analysis of randomized trials. They include personnel at the US Food and Drug Administration and other governmental agencies who regulate approval of new drugs, biologics, and medical devices. They include developers of clinical practice guidelines who have argued that evidence-based medicine should rest either solely or predominately on evidence from randomized trials, disregarding or downplaying evidence from observational studies. I hope that these readers will make the effort to understand the bases for my criticisms and that they will view the prescriptive decision analysis presented here as constructive suggestions.

Summary

Aiming to make the book accessible to a wide readership, the exposition in the main text is almost entirely verbal rather than mathematical. For readers

who want to dig deeper, I include a set of complements that formalize or elaborate on key parts of the discussion in the main text. I also provide references to the technical articles that present the full analysis.

The eight chapters of the book move from review and critique in chapters 1 and 2 to prescription in chapters 3 through 7 and conclusion. Chapter 1 reviews the continuing discourse in medicine regarding the circumstances in which clinicians should adhere to evidence-based practice guidelines or exercise their own judgment, sometimes called "expert opinion." Chapter 2 critiques how evidence from randomized trials has been used to inform medical decision making.

Chapter 3 describes research on identification, whose aim is credible use of evidence to inform patient care. Chapter 4 develops decision-theoretic principles for reasonable care under uncertainty. Chapter 5 considers reasonable decision making with sample data from randomized trials. Moving away from consideration of a clinician treating an individual patient, chapter 6 views patient care from a population health perspective. Chapter 7 considers management of uncertainty in drug approval. The final chapter provides concluding suggestions that encourage putting the themes of the book into practice.

INDEX