

Preface

This book grew out of attempts to explain what we do outside our profession – economics – and from the belief that many of our insights do not, or should not, require much mathematical sophistication. Attempts to explain the winner’s curse in simple words and to model it as a simple selection bias was a point of departure: in competitive auctions, bidders who are more optimistic about the value of an object are more likely to win the auction (this is the selection bias), and winners are thus likely to be disappointed when they subsequently learn their value of the object. A formal statement of this logic should not require knowing differential equations.

In the pursuit of simpler models, we introduce *direct strategy restrictions* on agents’ behavior. That is, rather than assuming that an individual tries to learn how much to bid when his value for a painting is \$1,100 and how much to bid should his value be \$1,200, we examine individuals trying to figure out what fraction of their valuation they should bid, whatever that value is. Mathematically the simplification is that we are trying to determine the value of a scalar, a fraction, rather than a huge number of scalars, one for each possible value of the painting that the individual might have.

From a pragmatic perspective, much of our effort boils down to investigating models in which a single dimension of behavior is endogenized at a time.

This book also grew out of discomfort with the predictions of some of our models, a rising gap between economic intuition and the output of models, the feeling that conclusions sometimes reflect the prodigious cognitive capabilities with which we endow the agents in our models, rather than true economic forces, and that the level of sophistication often draws attention to less relevant mathematical details at the expense of economic intuition. Part II illustrates our discomfort as simply as possible in the hope that it can be read without prior knowledge of the models discussed.

Last, this book critiques the near universal way that our discipline models ignorance and information: we mostly abandon the central tenet that an agent’s ignorance be modeled as a probability distribution over states. We focus instead on what agents can plausibly perceive, without making *a priori*

recommendations on the mathematical form that these perceptions should take, and on what agents can plausibly make of these perceptions.

Modeling the unknown seems to be a contradiction in terms. How can one describe the unknown without providing overly detailed contours to it? A central step in understanding how this is achieved in economic models is in realizing that a model contains two perspectives: that of the agent whose behavior we try to understand or characterize, and that of the analyst who acts as an omniscient outsider, making precise the economic environment and what is unknown to the agent.

For some problems, the two perspectives may coincide. When a die is thrown, both the analyst and the agent may be ignorant of the face that will show up, and it seems fine to assume that each face has a one-in-six chance, and that the agent understands this as well.

Most economic situations fall short of such easily quantifiable ignorance. Typically, neither the analyst nor the agent knows the chances of each outcome, nor even all possible outcomes. For the sake of modeling, the analyst generally posits a specific set of outcomes, and specific chances of those outcomes. This mathematical object, a probability distribution, is a modeling convenience for the analyst. Yet, behavior is then analyzed as though agents could exploit the precise specification, as if they were as omniscient as the analyst who sets out the model.

A central contribution of this book is to give direct strategy restrictions a prominent role. These restrictions may not only lead to more parsimonious models, and models that are less subject to the criticism expressed above – the dependence of predictions on less relevant aspects of the model. They may also constitute an effective way to disentangle the analyst's and the agent's perspectives by limiting the ability of agents to exploit parameters that the analyst introduced for her convenience, and which were not meant to be known by the agent, or at least not intended to significantly drive the agents' decisions.

Last but not least, asking that one define the range of available strategies prompts the analyst to think about what constitutes a reasonable range, taking into consideration the agent's thinking or limitations. In this way, sophistication or bounded rationality considerations can be introduced into an otherwise standard model. Defining the range of available strategies also highlights something omitted from standard models: information lies not only in what one observes and perceives, but also in what one makes of one's perceptions.

Comments and Warnings.

1. *Ignorance and Uncertainty.* The book is about modeling, and the title highlights one of the main challenges in constructing a model: how does one reconcile the analyst's need to define the economic environment in detail and yet keep agents ignorant of many aspects of the environment that, for the sake of mathematical modeling, the analyst has made artificially precise? How can one reconcile the analyst's need to quantify the uncertainty that agents face

through the definition of exact probability distributions, and yet avoid making agents too knowledgeable of the details of the model itself? The book proposes that one may keep agents somewhat ignorant of the model's fine details by limiting their ability to exploit these details.

2. *Credits.* Modeling ignorance is central to our discipline and many of the issues we discuss have been mentioned by others. We don't suggest that we are the first to address them. We sometimes cite papers that can be identified as seminal for an issue, but we do not cite all the papers that are related to our discussion. The same comment applies to the many applied topics covered in this book: we build on decades of modeling efforts, and do not provide an extensive discussion of the literature on each of the topics covered.

3. *Audience.* We wrote this book in the hope that it can reach multiple audiences, including people with only a knowledge of basic economics (though some understanding of the objectives of modeling will help). For this reason, the more technical discussions, or the discussions that assume some knowledge of the literature have often been relegated to the comments sections of each chapter, or to footnotes.

The book is intended to be a complement to a textbook, not a substitute, with the hope that it will help readers better understand economic modeling and its foundations, that it will suggest an alternate path, or at least trigger a critical appraisal of the literature we discuss.

4. *Warning.* While this book is a critique of the standard way to model ignorance – the Bayesian methodology – we do not advocate discontinuing that line of research. Rather, we view the book as highlighting difficulties in interpreting the results obtained from standard models and in applying them to real problems. Our view is that the book provides a guide for questioning or reevaluating the plausibility of the predictions obtained in such models, and that it provides one (of possible many) alternatives to the standard methodology to address the difficulties that we identify.

5. *Organization.* Part I is meant to provide an easy access to the main issues in constructing a model, in particular the link between the set of strategies available to the agents, and the agent's knowledge about the structure of the model implicit in his ability to optimize among these strategies. Chapter 1 describes this link in an environment with no uncertainty. Chapters 2 and 3 explain the challenges of modeling an uncertain environment. Chapter 4 motivates direct strategy restrictions while Chapter 5 analyzes some implications of our modeling approach.

Part II discusses important issues addressed by the discipline (information aggregation, mechanism design, surplus extraction, folk theorems, comparative statics) in light of our central concern – the possibility that models provide

agents too many instruments, and that insights are sometimes driven by this richness.

Part III discusses a number of standard applications, suggesting various ways to improve the balance between a sufficiently rich environment and a reasonably rich set of instruments.

Apart from Part I, most chapters are self-contained, and can be read independently of each other, although the reader may benefit from reading Chapters 12 to 14 sequentially.

6. *Acknowledgments.* The authors wish to thank Vincent Crawford, Itzhak Gilboa, Ron Spiegler, Vincent Crawford and especially Larry Samuelson, for detailed and useful comments on a preliminary version of the book. We have benefited from discussions with Andrew Caplin, Guillaume Fréchette, Harold Cole, Dirk Krueger, Dan Levin, Laurent Mathevet, Paul Milgrom and Jean Tirole; from extensive correspondence and/or exchanges with Alessandra Casella, Paul Klemperer and Michael Woodford; and, lastly, from encouragement from Jeff Ely. Olivier Compte thanks Philippe Jehiel for more than two decades of invaluable interactions, and Diego Aranega for agreeing to be distracted from his usual targets. Andy Postlewaite would like to thank David Schmeidler for decades of discussions. We also thank Joonbae Lee for proofreading and assistance on figures, and Alecia McKenzie for numerous editorial improvements.

7. The cartoons which appear at the beginning of Chapters 6–10 have been used with permission from Diego Aranega.