Introduction

1. This book has three goals: methodological, critical, pragmatic. It questions the way economic models deal with ignorance and uncertainty, and proposes alternative modeling strategies. It questions some of the main insights of the literature, and observes that they are sometimes a result of special features of our models. Last, it questions the high degree of mathematical technicality embodied in many standard economic analyses and suggests how one might reduce it.

These goals are linked. One element of the connection stems from how ignorance is typically modeled, via the *Bayesian methodology*, which consists of making precise what is not known: if you don't know the number of balls in an urn, define the range of plausible numbers and a probability distribution over them. This allows the analyst to artificially quantify the unknown but it embodies significant technical sophistication, often foregoing plausibility and parsimony.

A second element lies in the exercise of modeling itself, which attempts to combine two often conflicting perspectives, that of the analyst, who takes an omniscient perspective, describing in detail an artificial economic environment to be analyzed, and that of the agent, not meant to be omniscient, who is nevertheless assumed to know or behave as if he knew the details of the model.

A third element rests on how perceptions are handled in standard models. Typically, perceptions are modeled as a probability distribution over consequences and come with the recipe that defines how to use them, namely through expected utility maximization. In practice, a perception is akin to an ingredient in search of a recipe. A perception is useless if one does not know how to use it. A perception may sometimes be misleading, and better discarded. Information has to do both with what one perceives, and with one's ability to use and benefit from what one perceives.

2. This book is more generally about modeling. As emphasized in many economics and game theory textbooks, models are meant to illuminate some aspects of an economic problem or strategic situation, to shape intuition, or to help us detect when intuition is flawed. In this quest, the analyst inevitably

makes simplifications and approximations, and focuses on the ingredients that seem most relevant while omitting those that seem of second-order importance and would only distract us.

Simple models have a virtue. However, eventually, we are interested in a simple model only to the extent that we believe that it captures an essential aspect of behavior present in real-world problems. A challenge is that in the process of simplifying the environment, some irrelevant details may become salient, or acquire undue importance. We may think that a simple model of auction should consider agents whose values are either high or low, say 100 or 10. If we analyze such a model, behavior will eventually be driven by the particular values 100 and 10, and the hoped-for simplification may become the seed for complex behavior driven by the analyst's particular choice of values (100 and 10).

A seemingly easy fix to such a concern would be to define a richer environment in which values can take any value from 10 to 100. Yet, the fix is often an illusion. Other, less obvious details such as the value of the upper limit 100 may acquire undue importance, and without care, we may make predictions that hinge on agents behaving as though they could determine, among other things, that having a value close to 100 also means likely having the highest value for the object.

This book is about understanding why and when some details of models have undeserved importance (Part I and Part II), and about suggesting means to avoid this (Part III).

3. In short, a modeling exercise consists of finding an appropriate balance between a possibly rich economic environment – yet simple enough to be handled mathematically – and a set of instruments (or strategies) with which we endow the agents, thereby enabling them to adjust to that environment. When the balance is off, for example when an agent is endowed with a set of instruments that gives him implausible powers of discernment, some irrelevant details of the environment description can take on unwarranted importance because agents end up behaving as if they were knowledgeable of special aspects of the environment that were introduced solely as simplifying modeling devices. When this happens, a model's predictions can be driven by the inner structure of the model rather than reflecting true economic forces, or they can fail to account for relevant forces (Part II).

How can an analyst restore the balance between the economic environment considered and the instruments or strategies provided to each agent to exploit it?

The usual path consists of further enriching the environment. As suggested in our auction example, if it is too easy for agents to exploit the structure of the basic model, one might decrease their knowledge about their environment. The Bayesian methodology accomplishes this by introducing further uncertainty, defining a more complex model in which parameters of the basic model are realizations of a random variable.

Often, however, enriching the environment comes at the expense of parsimony, with no guarantee, as our auction example suggests, that the balance is restored. Also, in standard models, a richer environment often implies a richer set of instruments as well: rather than choosing a bid when one's value is 10 and a bid when one's value is 100, in the enriched environment, the potential buyer may in principle choose a bid for each value realization between 10 and 100.

This book suggests an alternative. It advocates direct restrictions on the set of instruments with which agents are endowed. We do not necessarily argue against enriching the environment, but rather against simultaneously enriching the set of instruments available to the agent. In Part III, we provide several standard economic environments and suggest plausible strategy restrictions. The models considered are possibly quite rich, in the sense that the economic environment (what the agent cares about, what the agent perceives) takes values in a continuum. *Despite the richness*, they are strategically simple.

In summary, analysts typically prefer simple models. What makes a model complex is not the profusion of data assumed to be available to the agent, but the degree to which agents are able to process the data: what is the range of possible behaviors allowed? A model may be descriptively rich, yet strategically simple. Part III considers models that are rich enough to avoid triviality, and yet simple strategically, with a single dimension of behavior examined at a time.

4. *Ignorance and uncertainty.* Our work has been driven by pragmatism, and the desire to convey economic intuition as simply as possible. We believe the exercise also contributes to the debate on how to model ignorance and uncertainty.

There has been a huge effort to extend economic models to worlds in which agents face uncertainty or ignore some aspects of the situations that they face. The work of Savage (1954) is the culmination of this effort, leading to a representation of decision making under uncertainty in which agents make choices as if they had a utility function over consequences and a personal probability distribution (or belief) over consequences, and used that belief to maximize expected utility.

This way of thinking about ignorance remains today at the heart of the economist's toolbox. We refer to this as the *Bayesian methodology*: for any aspect of a problem where the analyst thinks the agent lacks complete knowledge, define the possible realizations of that aspect and a probability distribution over them.

Another challenge has been to *incorporate* the methodology into economic models. In itself, the methodology places no restrictions on what agents believe. Beliefs mirror what agents choose. An agent may choose to bet all his fortune on a horse named Daisy, thereby reflecting his belief that Daisy is a sure win. However, if one wants to study horse races in which there is some uncertainty about the winner, a model that assumes that an agent could have arbitrary beliefs such as the above one seems too flexible.

In practice, analysts typically impose discipline on how beliefs are incorporated into models. Discipline can be imposed by assuming that the agent's belief bears some relationship with the actual "chance" that this particular horse wins, however questionable defining this "chance" objectively might be. More generally, it is achieved by assuming a *consistency* condition on beliefs, taking the form of a possibly stochastic relationship between signals (capturing perceptions) and states (capturing elements or aspects of the environment).[1]

For example, suppose we wish to model how an agent reacts to hunger, where hunger is used as a proxy for the level of depletion of the reserves in the agent's body. As an analyst, one may posit a stochastic relationship (or probability distribution) between hunger levels and the level of reserves. Then, to any hunger perception, the analyst (and the agent if he knows the model) may associate a belief – a probability distribution over reserve levels.

5. There are two well-known difficulties with the approach described above. First, it often gives rise to an *overly precise* representation of what the agent is supposed to be ignorant about. Probabilistic beliefs are an instrument invented by analysts to structure, analyze and describe the behavior of agents. Yet, by and large, as a *positive* description of behavior, beliefs remain a somewhat implausible construct. Not because agents would *not* form beliefs; agents undoubtedly have some elements of likelihood in mind when making a decision. But whatever form agents' beliefs actually take, they are surely more casual than what the Bayesian methodology assumes.

Second, when the consistency route mentioned above is adopted, the approach potentially gives rise to an *overly accurate* representation of what the agent knows.[2] To illustrate with an extreme example, one might be ignorant of whether the square root of 73,057 is above or below 281.56, or unable to say whether it is above or below 281.56 within 5 seconds.[3] Since there is an objective answer to that question (281.56 is larger than the square root of 73,057), the consistency route that assumes there is a stochastic relationship

[1] In games, incorporating the Bayesian methodology has been done by assuming a consistency condition among beliefs held by different players (Harsanyi). Myerson (2004, page 1824) justifies this consistency condition as follows: "If we can assume any arbitrary characteristics for the individuals in our model, then why could we not explain the surprising behavior even more simply by assuming that each individual has a payoff function that is maximized by this behavior? Thus, to avoid such trivialization of the economic problem, applied theorists have generally limited themselves to models that satisfy Harsanyi's consistency assumption."

[2] Or, more generally, driven too much by the distribution chosen by the analyst, an assumption made for lack of a better model, for her convenience only.

[3] This example can be seen as a variation of one provided in Lipman (1999), which we shall discuss in Chapter 21. See also Chapter 5. The example is extreme because, to the analyst, there is no underlying uncertainty about the correct answer.

between the correct answer and the agent's perception necessarily restricts beliefs to the correct answer, in spite of the fact that the agent might only have a vague impression about whether one number exceeds the other.

More generally, even in problems where the analyst assumes some uncertainty over the underlying state, the consistency route does not simply provide a mechanism to translate signals or perceptions into a precise belief; it also provides a perfect guide to using these perceptions, as though that information was always immediately available to the agent – the agent just needs to maximize expected utility given the belief associated with the perception.

In practice, I can have erroneous perceptions that a particularly dangerous activity is safe, or that betting on Daisy is a sure win: perceptions are signals that one might use, and sometimes profitably ignore or take with caution. The conclusion that some perceptions should be ignored or taken with caution constitutes information, and this information may or may not be easily available to the agent, or easily quantified.

Said differently, a perception is like an ingredient without a recipe. To accompany your asparagus with a sauce mousseline, you may know that you need oil, butter, whipped cream, eggs and heat, but if you don't know how to combine those ingredients, you won't come close to a sauce mousseline, and if you don't keep track of temperature, you will end up with a greasy omelet.

This book is an attempt to address the idea that perceptions are one thing, learning to deal with them is another. It is an attempt to disentangle perceptions and information, to disentangle ingredients and recipes, and to disentangle the data the agent gets from the various ways he may use it. Accurate perceptions are valuable, as quality products may improve a meal. But information also stems from the ability to determine the profitable uses of perceptions.

6. This book argues that one cannot eschew drawing a precise connection between the specific conditions agents face and how they perceive them. As for the consistency route, this connection is defined by the analyst and represented as a joint distribution over specific conditions and perceptions. However, in light of the discussion above:

(i) we are agnostic about the mathematical form taken by perceptions, favoring, when possible, simpler mathematical objects having plausible interpretations;
(ii) we do not assume that agents know the distribution, but assume that by comparing experience from different strategies, they identify which performs best;
(iii) we avoid assumptions that facilitate agents' overly exploiting that distribution, and we do this by *limiting a priori* the strategies available to each agent. In this constrained world, "information" will embody not only the perceptions or the data that an agent gets, but also what she can make of them – for example, the set of possible recipes available (or made available) to her.

In short, we do not propose a universal way to deal with ignorance and uncertainty. Different economic environments and different degrees of sophistication call for different ingredients and different recipes. Part III is like a cookbook. For each kind of dish, it proposes basic ingredients and a basic set of recipes, from which each agent finds the one most suited to her taste. We do not claim that there is a unique way to define the set of recipes. We aim for a basic cookbook, characterizing what seems to us essential aspects of a number of strategic phenomena. Other sets of recipes, possibly characterizing other ways by which agents comprehend their strategic environment, and possibly conforming better with experimental evidence, would shape results and intuitions differently. In the end, one must judge the various restrictions on the set of strategies allowed to the agent by their usefulness in shaping our thinking and understanding of strategic behavior, and their empirical support.

References

Lipman, B. L. (1999). Decision theory without logical omniscience: toward an axiomatic framework for bounded rationality. *The Review of Economic Studies*, 66 (2), 339–361.

Myerson, R. B. (2004). Comments on "Games with Incomplete Information Played by 'Bayesian' Players, I–III." *Management Science*, 50 (12), 1818–1824.

Savage, L. J. (1954). *The Foundations of Statistics*, New York: John Wiley and Sons, 1954.

PART I

Modeling Challenges

As analysts, we use formal models to describe the situations that agents face. In doing so, we take the perspective of an omniscient outsider who would have superior knowledge of the details of the interaction: we choose the strategies available to agents, as well as the payoff structure, that is, the payoffs that each agent gets as a function of the strategies played. Most often, we also wish to prevent agents from knowing with precision the payoff structure. To that end we introduce an *information structure*, which describes the uncertainty about the payoff structure, and the (possibly stochastic) process that generates what agents perceive or observe conditional on each possible payoff structure.

Although a model is an analyst's tool, not meant to be known with precision by the agents being described, the mere assumption that an agent can compare the performance of the strategies available to him grants him some implicit knowledge of the model, or *inside knowledge*. The objective of Part I is to explain the connection between optimization and inside knowledge, to explain how modeling choices have consequences on the degree to which agents can exploit the structure of the model itself.

Why is this important? By constructing a model that endows agents with overly precise knowledge of the model itself, we run the risk of deriving insights that hinge on that unrealistic assumption. Models need not be realistic, but we hope that the insights that we derive from them are not too strongly driven by unrealistic features of the model.

Cambridge University Press
978-1-108-42202-4 — Ignorance and Uncertainty
Olivier Compte, Andrew Postlewaite
Excerpt
More Information

CHAPTER 1

Action Space

1. Economic theory studies decision problems and strategic interactions, with the objective of understanding and/or predicting the behavior of agents involved in these situations.

Modeling a decision problem or a strategic interaction begins by specifying a set of *actions*, which agents may choose from, and a *payoff structure*, which is a formal description of how actions translate into payoffs for each agent. There exists an extraordinarily vast array of decision problems or strategic situations because, in principle, there is no *a priori* limit on the space of possible actions available, nor limits on the possible mappings from actions to payoffs.

Analysts have considerable freedom in choosing the action space and the payoff structure when constructing models, and a great achievement of game theory has been to identify, within that vast array of situations, simple situations that provide insight into important real-world problems, and for which behavior can be described or characterized. One such example is the Prisoner's Dilemma.

2. The prisoner's dilemma is a two-player game in which each player has only two actions (i.e., a two-by-two game), confess (C) or not confess (N), with the property that confessing is a better option for each individual whatever choice the other makes, and where *both not confessing* is a better outcome than *both confessing*.

The following matrix describes payoffs that are consistent with the properties described above:[1]

	C	N
C	$1,1$	$4,0$
N	$0,4$	$3,3$

Given these payoffs, each player individually finds that confessing is a better strategy, because $1 > 0$ and $4 > 3$. Confessing is said to be a dominant

[1] The matrix indicates that if player 1 chooses N and player 2 chooses C, player 1 gets 0 and player 2 gets 4.

strategy. The prediction is that both confess despite *both not confessing* being a better joint outcome, thereby reflecting the conflict between private objectives (confessing is individually better) and social objectives (both not confessing is jointly better).

3. A virtue of the formal description above is that it can be used across disparate applications: beyond the dilemma that prisoners might face, there are many interactions that fit naturally in the two-by-two game above, with "not confessing" characterizing a *cooperative* strategy, and "confessing" characterizing a *selfish* strategy. These broadly defined labels (cooperative and selfish) may capture different behaviors depending on context. But this is precisely why the model is useful, making it easily applicable across a large variety of situations.

4. Another virtue is that the prediction holds not just for a single specification of the payoff parameters, but for a large range of values, as long as the dominance relations hold. In particular, the agents need not have precise knowledge of the parameters of the model for the prediction to hold.

We emphasize the latter point, as this is central to the critique of the literature that we address in this book. In writing down a payoff structure, we take an outsider's perspective, defining what each player gets as a function of the pair of actions played. In solving the game, we derive "optimal" choices for each player as though they knew the payoff structure. As analysts, we avoid complicating the model further; we avoid being precise about what each player actually knows about the payoff structure.[2]

The reason for doing so is parsimony. A model is meant to be an analyst's tool, a parsimonious way to represent reality, helping to explain economic insights which seem relevant. For the sake of parsimony, we generally solve the model as if the agents had precise knowledge of its parameters, hoping (without formally verifying) that the insights drawn from the model do not hinge on this questionable assumption. The prisoner's dilemma safely passes this test.

5. The restriction to two actions ("not confess," the cooperative strategy, and "confess," the selfish strategy) provides a parsimonious model of the conflict between private and social objectives. There are many contexts, however, in which one could imagine varying degrees of cooperative behavior, and where the restriction to two actions could be viewed as unrealistic. In an attempt to assess the strength of the forces away from "full" or efficient cooperation, one may want to enrich the model with multiple levels of cooperation.

But there is a tension: while the restriction to two actions enables the analyst to capture the basic strategic effect, further quantification of this effect

[2] Analysts sometimes take a different view, assuming that payoffs are precisely known, and known to be known, etc. We discuss this alternative view at the end of this chapter.

Cambridge University Press
978-1-108-42202-4 — Ignorance and Uncertainty
Olivier Compte, Andrew Postlewaite
Excerpt
More Information

through a more "realistic" action space is subject to the criticism that the solution implicitly assumes that the agents of the model have substantial knowledge of the structure of the model. The next example illustrates that tension.

6. *A partnership game.* Consider a standard partnership problem in which each of two agents $i = 1,2$ picks an effort level e_i where the effort level can be any non-negative real number. Agent i's gain from the pair (e_1, e_2) is defined as $g_i^\gamma(e_1, e_2)$, with:

$$g_i^\gamma(e_1, e_2) = \gamma\sqrt{e_1 e_2} - e_i^2.$$

An *equilibrium* outcome in this model is a pair of actions (e_1^*, e_2^*) from which neither player wants to deviate unilaterally.[3] The equilibrium effort levels satisfy $e_i^* = \frac{1}{4}\gamma$, while the socially efficient levels would satisfy $e_i^{**} = \frac{1}{2}\gamma$.[4] The model allows one to quantify the effect of both agents following private objectives, each ignoring the positive externality on the other and the higher welfare that would result from a marginal increase in effort.

Despite being possibly more realistic in terms of the strategy space, the model implicitly makes implausible cognitive assumptions: the equilibrium outcome relies on agents behaving as though they knew the details of the model (e.g., the functional forms associated with gain and cost functions for each player) or as if they had learned which effort level was optimal among all possible levels.

7. *Coming to play an equilibrium.* Models are typically silent about how players come to play according to the equilibrium strategies the analyst identifies. One natural hypothesis is that equilibrium is the outcome of a learning process, the stable point from which individual experiments with alternative strategies are unprofitable. In this view, the cognitive assumption is not that players know the parameters of the model, but that they have learned which of their available strategies is best.

There are at least two dimensions that make learning difficult: the number of alternatives to be compared, and changes in the environment. The plausibility of the implicit cognitive assumption may therefore differ a great deal across models. In the prisoner's dilemma, the agent need only compare two actions, and best responses are unaffected by changes in the underlying payoff structure as long as the dominance relations continue to hold. In the partnership game, the cognitive assumption is stronger: the agent must compare many effort

[3] See the note at the end of this chapter for some history and motivation.

[4] Formally, an equilibrium is a pair (e_i^*, e_j^*) such that, for each i, the gain $g_i^\gamma(e_i, e_j^*)$ is maximum at e_i^*. The social optimum is a pair (e_i^{**}, e_j^{**}) such that the total gain $g_i^\gamma(e_i, e_j) + g_j^\gamma(e_i, e_j)$ is maximum at (e_i^{**}, e_j^{**}).