

The editors promise at least one more volume, covering such additional topics as equilibrium concepts, incomplete information, stochastic games, games with many players, bargaining, oligopoly, implementation, inspection games, the Shapley value and its applications, history of game theory, macroeconomics, experimentation, psychology, and law. I hope that this additional volume(s) is at least as successful as the current two.

Preparing these handbooks must have been a tremendous undertaking for the editors and contributing authors. For working game theorists, these books are indispensable (even at their high purchasing price of \$110 per volume). I wish to thank all the contributors. I hope that in ten years we will have another comprehensive updated snapshot of game theory, prepared by equally qualified and committed contributors. While this appears unlikely, given the required coordinated effort, the successful record of *Handbooks in Economics* in publishing follow-up volumes makes one more optimistic.

Ehud Kalai

*Northwestern University
Evanston, Illinois 60208*

A Course in Game Theory. By Martin J. Osborne and Ariel Rubinstein, MIT Press, Cambridge, MA, 1994.

Two hypotheses readily come to mind for the recent spate of game theory texts. On the supply side, perhaps a large cohort of game theorists has matured to the stage at which pedagogy pays; prolific researchers pause to bear their contributions to the classroom. On the demand side, perhaps game theory itself has matured; a new consensus creates a desire for new textbooks. Both hypotheses are supported by this new text by Osborne and Rubinstein (OR).

If game theory has matured, it has done so only recently (recent Nobel prizes notwithstanding). Both its substance and perceived purpose have changed dramatically over its short history. Initially, it was often viewed as an attempt to extend, for prescriptive purposes, the decision-theoretic ideal of individual rationality to groups. Then ascended the view that game theory should be a tool for description and prediction. Game-theoretic models of practically everything arose: consumers, firms, politicians, voters, insects, plants, and even supernatural beings. Entities in these models embody the decision-theoretic idealizations: they maximize expected utility, possess unbounded computational ability, and have common knowledge of payoffs, type distributions, strategies, and beliefs. Accumulating evidence,

much of it experimental, brings into question the descriptive adequacy of these idealizations. Theorists have recently responded with developments of non-expected utility, machine players, non-common knowledge, learning, and evolutionary games.

Perhaps the most notable feature of OR is its cautious introduction of some of these recent developments. It presents primarily standard topics (with some exceptions), but in a way that will sensitize students to the strength of rationality and knowledge assumptions. This aspect of the book sets it apart from the other textbooks at this level, as do its clarity and stress on interpretation.

* * *

Before further discussing these aspects, I should describe the book's level and coverage. I take the viewpoint of an instructor contemplating adoption.

The book rigorously develops game theory from scratch. It presents proofs with unusual care, though routine arguments are left to the reader. (Relative to some texts, such as Binmore (1992), OR does not coddle the mathematical tyro.) Nonetheless, behavioral matters are given priority over mathematical generality. For example, action and state spaces are generally finite, and the only fixed point theorem is Kakutani's. The book is appropriate for a game theory course in the second year of an economics or political science Ph.D. program.

The introductory chapter sets the tone. It defines game theory, says a few words about decision theory, says a few more words about bounded rationality, presents the steady-state (evolutive) and deductive (eductive) interpretations of solutions, and lays out notation. This is all done in just seven pages. As in every chapter, ideas are presented clearly and precisely, without verbiage.

Section I treats strategic games. The first three chapters deal with Nash, correlated and ESS equilibria, and rationalizable strategies. Bayesian games are presented early; Bayesian equilibria are defined directly and simply as Nash equilibria of the strategic game obtained by identifying each type of a player as a different player. Most notably, the section ends with a chapter on knowledge that presents an axiomatic treatment of information (P) and knowledge (K) functions, the derivation of solution concepts from knowledge assumptions (as surveyed in Brandenberger, 1992), and the electronic mail game (Rubinstein, 1989). This is probably the first introductory text to discuss non-partitional knowledge.

Section II treats "extensive games with perfect information," which OR defines to be games in which players move simultaneously within each of a sequence of stages and observe at each time of acting all the moves taken

in previous stages.¹ The first three chapters cover the basic theory of these games, the alternating offer bargaining model of Rubinstein (1982), and a nice treatment of perfect folk theorems for repeated games. The “one deviation property” (principle of optimality), which is so useful in dynamic models for checking subgame perfection, is emphasized. Also notable is the formulation of repeated game strategies as automata, a neat way of conveying the essence of a strategy. The automata formulation is crucial to the fourth chapter, which presents a simplified version of the material in Rubinstein (1986) and Abreu and Rubinstein (1988) on games in which the cost of implementing a strategy increases with its complexity. The section concludes with a concise chapter on non-Bayesian implementation, including the Gibbard–Satterthwaite theorem, Groves mechanisms, revelation principles for both dominant strategy and (full information) Nash mechanisms, the Maskin theorem on Nash implementation, and subgame perfect implementation.

Section III presents the basics of extensive games with imperfect information. The first of its two chapters contains foundations, including the notions of equivalent extensive forms and the relationship between mixed and behavioral strategies. The second chapter is called “Sequential Equilibrium,” but it also discusses perfect Bayesian and trembling-hand perfect equilibria. A brief three pages uses well-known examples to introduce belief-based refinements (Cho and Kreps, 1987).

Section IV, the final section of the book, treats cooperative game theory, a neglected topic in other recent texts.² The first chapter is a straightforward presentation of the core, including market games and core convergence in exchange economies. The second chapter considers stable sets, the bargaining set, the nucleolus, and Shapley values. The final chapter presents the Nash bargaining solution, following the reformulation of Rubinstein *et al.* (1992); the usual treatment of Nash bargaining is only briefly treated, with Nash’s characterization left as an exercise.

Relative to texts such as Myerson (1991) or Fudenberg and Tirole (1991), OR chooses to omit many topics and applications. For example, little is said about games of timing, differential games, repeated games with imperfect observations, Bayesian mechanism design, and strategic stability. According to the preface, if it were written today OR would include three other

¹ The pedagogical conservative will be disconcerted by this deviation from the long-standing definition that declares a game to have perfect information if all its information sets are singletons. OR attribute their definition to Dubey and Kaneko (1984); it corresponds to the definitions of “multi-stage games” in Myerson (1986, 1991) and Fudenberg and Tirole (1991).

² Fudenberg and Tirole (1991) have no words on the subject, and Binmore (1992) has only slightly more.

omitted topics: learning, evolution, and experiments.³ Missing applications include bargaining and trade under asymmetric information, oligopoly, and voting.⁴ The back-cover blurb is accurate when it states that OR

presents the main ideas of game theory, . . . emphasizing the theory's foundations and interpretations of its basic concepts. The authors provide precise definitions and full proofs of results, sacrificing generalities and limiting the scope of the material in order to do so.

The result is that OR is a thin 352 pages, versus the 568 pages of Myerson (1991), 579 pages of Fudenberg and Tirole (1991), or 642 pages of Binmore (1992).

To my mind, the lack of applications is not a drawback. A book without applications frees the instructor to discuss his or her favorite ones separately or to use as a second text, one that stresses applications (such as Fudenberg and Tirole (1991) or, at a lower level, Gibbons (1992)). Besides, there are good arguments, in addition to "lack of time," for not including many applications in a game theory course. In my experience, interjecting what should be serious applications between chunks of theory tends to diminish the applications—they can appear to be no more than finger exercises, undercutting (as well as cream-skimming) their development in later courses.

To be quite clear, let me distinguish between "examples" and "applications." Examples, as opposed to applications, are short, illuminate theory, and need have no substantive interest. Examples are pedagogically necessary, and OR has plenty of good ones.

* * *

Perhaps the most striking feature of OR, and for an introductory textbook its most surprising feature, is its skepticism about the rationality and knowledge of players. The clearest evidence of this is the presence of chapters on knowledge functions and strategy complexity.

In addition, a cautious attitude toward the expected utility hypothesis is evident throughout the book. Consider three examples. First, strategic games are presented initially in terms of preferences over action profiles and later, when mixed strategies and random moves by nature are introduced, in terms of preferences over lotteries. Payoff functions are used initially only to describe the classic game matrices, as examples.

³ The other texts also contain little on these topics. OR actually does include some material on large population models and evolutionary stable strategies.

⁴ The book is not devoid of applications. It treats some important ones, such as education signaling (Spence, 1973) and chain store reputations (Kreps and Wilson, 1982; Milgrom and Roberts, 1982).

Second, the alternating-offer bargaining game is presented in terms of preference relations on an abstract set of agreements; the discounted utility representation is invoked only later for the uniqueness proof. Third, the Nash bargaining solution is characterized in terms of preferences on lotteries, following Rubinstein *et al.* (1992).

Intertemporal preferences are also treated gingerly. Repeated games are presented initially in terms of general preference relations over sequences of action profiles. Even after utility representations are admitted, there is reluctance to restrict attention too quickly to the discounting criterion. The limiting average and the overtaking criteria are considered first and are given comparable space to the discounting criterion.⁵

OR's use of general preferences whenever possible reflects a growing skepticism in the profession about our usual preference assumptions. The case is strong for students to learn of this skepticism at the onset. Students who learn their game theory from OR will become comfortable working with preference relations, and they may be among the first to formulate applied models that do not rely on such strong assumptions as expected utility.

However, the pedagogic cost of using preference relations must be admitted. Their abstraction can be a stumbling block. Graduate students who intend to do theory will have no problems, of course, but it will be a pity if OR does not reach those who intend to do applied work, given the book's otherwise exceptionally clear exposition and stress on interpretation (see below).

* * *

Great lengths have been taken to make OR a pleasure to use for both student and instructor. Little things are noticed first, such as the exceptional diagram (page xii) of the connections between chapters—at least six dimensions of information are conveyed in this clever figure. Then there is the delightful numbering of lemmas, propositions, definitions, and exercises. Proposition 144.3, for example, is the third item on page 144 that is numbered; the reader always knows exactly on which page a cited result can be found. A journal editor once asked me to number results and equations so that “the paper could be nearly read backwards.” OR almost achieves this goal.

I have not seen a textbook better designed for reference. At the back appears a list of all results and the references for all cited articles, both of which are fully cross-referenced to the relevant pages in the text. The index is unusually detailed.

⁵ The strategy of presenting the perfect folk theorem first for the limiting average criterion, then for the overtaking criterion, and only then for the discounting criterion is a wonderful pedagogical device. Understanding why the proof for one criterion fails for the next one in this sequence conveys a deep appreciation of the theorem.

The exposition is lean and careful; its leanness serves well to reveal the bones of arguments. Every sentence contributes; students will want to read each passage more than once. Rarely are important points not discussed. Well-chosen examples are liberally sprinkled throughout the text.

* * *

Finally, let me turn to one of OR's greatest strengths, its stress on interpretation. From beginning to end, matters of interpretation are seriously addressed. The first chapter starts the discussion by alluding to the steady-state and the deductive interpretations of solution concepts. The final chapter focuses on interpretation, as it reinterprets the Nash bargaining solution first in terms of the risk of bargaining breakdown (Rubinstein *et al.*, 1992) and then in terms of alternating offer games with outside options (Binmore *et al.*, 1986).

More fundamental are the discussions of mixed strategies and extensive game strategies. In the first, mixed strategies are discussed alternatively as objects of choice, as stochastic steady states of a large matching population, as pure strategies of a game with payoff-irrelevant types, as pure strategies of a perturbed game with payoff-relevant types (Harsanyi's purification argument), and as the other players' beliefs about how the mixing player will behave. The presentation of each is cogent and illuminating. Particularly useful is a device that I at first thought to be too cute, namely, that of breaking the text into paragraphs labeled "AR" or "MJO" wherever the authors have agreed to disagree. I eventually came to appreciate these mock debates—they just might waken the lazy reader who generally assumes text is truth.

Regarding extensive form strategies, OR first observe that the prevalent view of them as plans of action cannot be entirely correct, since a player's strategy must specify what he does in parts of the game that cannot occur if he plays this strategy. The alternative interpretation is that a player's strategy captures the beliefs of others about what he would do in each contingency. Important implications of this view are brought to light: beliefs must be unanimously held, a player cannot literally choose a strategy, and the rationale for subgame perfection is questionable.⁶

OR's discourses on interpretation are valuable. One might object, however, on the grounds that a textbook should leave philosophical points until after the student masters the mechanics of definition and proof. This argument has merit, and I was happy to see OR treat the philosophical issues in their usual parsimonious way. But matters of interpretation are too important to gloss over. It is time for applied models to reflect a deeper understanding

⁶ Some of OR's observations on interpretation are taken from a 1988 Walras-Bowley lecture (Rubinstein, 1991).

of “strategy” and “equilibrium.” In any case, the discussions of interpretation are a pleasure to read.

* * *

OR is an innovative and wonderfully crafted textbook. It should have a significant effect on the development of both theory and application. I am envious of the students who will learn from it and thankful to the authors for writing it.

REFERENCES

- ABREU, D., AND RUBINSTEIN, A. (1988). “The Structure of Nash Equilibrium in Repeated Games with Finite Automata,” *Econometrica* **56**, 1259–1281.
- BINMORE, K. G., RUBINSTEIN, A., AND WOLINSKY, A. (1986). “The Nash Bargaining Solution in Economic Modeling,” *Rand Journal of Economics* **17**, 176–188.
- BINMORE, K. G. (1992). *Fun and Games*, Lexington, MA: D. C. Heath.
- BRANDENBURGER, A. (1992). “Knowledge and Equilibrium in Games,” *J. Econ. Perspectives* **6**, 83–101.
- CHO, I.-K., AND KREPS, D. M. (1987). “Signaling Games and Stable Equilibria,” *Quart. J. Econ.* **102**, 179–221.
- DUBEY, P., AND KANEKO, M. (1984). “Information Patterns and Nash Equilibria in Extensive Games, I,” *Math. Soc. Sci.* **8**, 111–139.
- FUDENBERG, D., AND TIROLE, J. (1991). *Game Theory*, Cambridge, MA: Harvard Univ. Press.
- GIBBONS, R. (1992). *Game Theory for Applied Economists*, Princeton, NJ: Princeton Univ. Press.
- KREPS, D. M., AND WILSON, R. (1982). “Reputation and Imperfect Information,” *J. Econ. Theory* **27**, 253–279.
- MILGROM, P. R., AND ROBERTS, D. J. (1982). “Predation, Reputation, and Entry Deterrence,” *J. Econ. Theory* **27**, 280–312.
- MYERSON, R. B. (1986). “Multistage Games with Communication,” *Econometrica* **54**, 323–358.
- MYERSON, R. B. (1991). *Game Theory*, Cambridge, MA: MIT Press.
- RUBINSTEIN, A. (1982). “Perfect Equilibrium in a Bargaining Model,” *Econometrica* **50**, 97–109.
- RUBINSTEIN, A. (1986). “Finite Automata Play the Repeated Prisoners’ Dilemma,” *J. Econ. Theory* **39**, 83–96.
- RUBINSTEIN, A. (1989). “The Electronic Mail Game: Strategic Behavior under ‘Almost Common Knowledge,’” *Amer. Econ. Rev.* **79**, 385–391.

- RUBINSTEIN, A. (1991). "Comments on the Interpretation of Game Theory." *Econometrica* **59**, 909-924.
- RUBINSTEIN, A., SAFRA, Z., AND THOMSON W. (1992). "On the Interpretation of the Nash Bargaining Solution and Its Extension to Non-expected Utility Preferences," *Econometrica* **60**, 1171-1186.
- SPENCE, A. M. (1973). "Job Market Signaling," *Quart. J. Econ.* **87**, 355-374.

*Steven A. Matthews**

*Department of Economics
Northwestern University
Evanston, Illinois 60208*

*I thank Eddie Dekel, Ehud Kalai, Kiminori Matsuyama, Wolfgang Pesendorfer, and Rebecca Stein for comments. The views expressed are entirely mine. E-mail: s-matthews@nwu.edu. Fax: (708) 491-2530.