THEORETICAL ANALYSES OF BOUNDED RATIONALITY AND LEARNING

A Review of Ariel Rubinstein's *Modeling Bounded Rationality,* Cambridge, MA: MIT Press, 1998. Pp.viii + 208. ISBN 0 262 68100 5. Also a review of Drew Fudenberg and David K. Levine's *The Theory of Learning in Games*, Cambridge, MA: MIT Press, 1998. Pp. xiv + 276. ISBN 0 262 06194 5.

Standard microeconomic theory is silent on the procedures that economic agents use in making decisions: the rational economic agent is assumed to have a consistent set of preferences and he simply chooses the optimal action,

subject to the constraints that he faces; in interaction with other such agents, each one successfully deduces what the others will do and an equilibrium results. The uncovering of more and more instances in which the traditional rational actor model predicts actual human behavior rather badly has increased dissatisfaction with the standard model and has stimulated the development of alternative models of choice. As game theoretic models push the perfect rationality assumption to its limits, the limitations have been most visible in that context and alternatives have naturally also been developed there. During 1998, two books appeared with MIT Press: Ariel Rubinstein's 'Modeling bounded rationality' and Drew Fudenberg and David Levine's 'Theory of learning in games' that together provide a good overview of recent work in this area.

The two books are complementary with Rubinstein covering a broader area and Fudenberg/Levine going more into depth on the specific issue of learning in repeated games. The advanced text of Fudenberg/Levine investigates whether learning processes will produce the same type of equilibria as traditional models that assume that players in a game can perfectly predict each other's actions. It will mainly be of interest to researchers in game theory and related fields. The book by Rubinstein, being less specialist and providing material for a one-term graduate course, may be attractive for a wider audience. Rubinstein discusses a collection of models, each one incorporating a specific procedural element of choice making. He deliberately does not touch upon the growing literature that deals with evolutionary and learning models, the main arguments being that that topic deserves a book of its own, that it uses different mathematical tools and that, in the standard literature on evolution and learning, there is little room for deliberations about decisions, rather choice arises by applying a mechanical rule. Fudenberg and Levine's focus is precisely on these learning models. We now discuss both books in turn.

In his 1950 Ph.D thesis John Nash already argued that the Nash equilibrium concept can be interpreted in two completely different ways, viz. as an outcome of deliberation and as outcome of learning. The game theoretic literature since then has mainly relied on Nash's first, rationalistic, interpretation. In this interpretation, players are perfectly rational and can predict correctly what their opponents will do; an equilibrium is a situation in which, given these correct predictions, each player maximizes his payoff. Nash's second 'mass-action' interpretation involves the game being repeatedly played over time with players being boundedly rational. These players observe the outcomes through time, learn and gradually adjust their behavior. Nash claimed that if behavior converges it must converge to a Nash equilibrium. In their book, Fudenberg and Levine investigate in what contexts the original intuition of Nash can be confirmed. The book thus addresses two main questions: (i) when and why should we expect play to converge to a Nash equilibrium? and (ii) if the game has multiple Nash equilibria, which ones should we expect to observe? In addition, it addresses such issues as how long will it take before players

have learned to play an equilibrium. It, however, does not extensively address the question of what we might expect when behavior does not settle down.

The Fudenberg/Levine book makes a useful distinction between learning models and evolutionary models. In the former class, learning is modeled at the level of the individual; there is an explicit model about how an individual incorporates new information and how the adjusted state leads to a new action. An example, that plays an important role throughout the book, is fictitious play in which agents always play a best response against the empirical distribution observed in the past. Hence, if an individual has seen that his opponent has played seven times T and three times B, he will believe that in the next period T will occur with 70 per cent probability and B with 30 per cent. This process can be viewed as a Bayesian learning process. The second class consists of models of 'social learning'. Here the adjustment is described directly at the aggregate level of the population without going into details about how such changes result from changes at the micro level. The question, of course, is why such a dynamic would be interesting for economic settings. As the authors show, such processes might result when, at the micro level, agents learn by asking around and by imitating other successful agents, or when there is learning until a certain aspiration payoff has been reached.

The first part of the book is devoted to normal form games. Chapter 2 deals with fictitious play. The basic result is that, if the empirical distributions of actions converge, the product of these distributions must be a Nash equilibrium. While this is a positive result, the authors point out its limitations: actual play need not converge (it may cycle), nor need the payoffs converge to the equilibrium payoffs. Furthermore, because players base themselves on the same history, correlations and correlated equilibria arise naturally. Hence, a first conclusion is that Nash's original intuition cannot be sustained in general. Chapter 3 then reviews the results that are available for the 'replicator equation' and for related models with monotone dynamics, i.e. actions that do better spread more quickly through the population. Again, stable steady states of these systems must be Nash equilibria. Chapter 4 discusses models with a persistent stochastic component, this randomness, for example, being caused by mutations or errors in decision making. The book reviews the models of Kandori, Mailath, Rob, Young and Ellison that establish equilibrium selection results, showing specifically that in the long run, the system will end up in the risk dominant equilibrium.

The second part of the book investigates the same set of issues, but now in the context of extensive form games. The important insight is that, when players only learn the actual outcomes of play, they need not have common expectations about behavior of the equilibrium path, hence, that outcomes need not be Nash equilibria. (The Nash equilibrium concept insists on common beliefs, both on and off the equilibrium path.) Nash equilibria are obtained only if there is sufficient player experimentation so that enough data off the path are generated. In general, however, only the weaker concept of

self-confirming equilibrium that was introduced by the authors is obtained.

The book concludes with a chapter on 'sophisticated learning' that deals with the general question of whether and how a player can exploit the fact that his co-players are learning as well. For example, a player could learn to detect patterns in past play and exploit these. Alternatively, a player who plays the game more than once, against short-term players, need not play myopically, it is better for him to 'teach' his co-players, in this way he can obtain his more preferred Stackelberg outcome.

Overall, the Fudenberg and Levine book provides an excellent overview of the recent theoretical literature on learning and evolution in games. The book focuses on the mathematical aspects. While there is some discussion about observations from experimental games at places one may, perhaps, regret that there is not more discussion of actual learning processes and of the speed of these. Some of these experiments have taught us that, even in relatively simply contexts, learning may be very slow, so that the relevance of the long-run equilibrium may be limited. (See, for example, the discussion on the simplified version of Akerlof's lemons problem in Selten (1998).) A more extensive discussion of this experimental literature might have enabled the reader a better judgement on the domain where the theory that is developed here is relevant.

As stated above, Ariel Rubinstein sees the main distinction between his work and the type described in the book by Fudenberg/Levine in that his deals with deliberate choice making using procedures and not with mechanical adaptation to the environment.

Rubinstein starts by reviewing the main implicit assumptions in the standard rational actor model: a rational agent is assumed to know (to have a full overview of) the problem, to know what he wants, to do what he wants (i.e. he has the ability to optimize) and not to be misled by the way in which the problem is presented to him (i.e. framing effects do not play a role). The problem with this model is not just that the assumptions are unrealistic but also, and more importantly, that it produces results which differ significantly from actual human behavior, hence, that it may not offer much help in understanding human behavior. For example, the psychological literature provides convincing evidence that framing effects are important: the way the situation is presented may determine the choice that the subject takes in that situation. The reason for this 'frame dependence' may lie in the heuristics that subjects use in simplifying the original problem that is presented to them: one frame may lead to a different simplified problem that is solved by cognitive effort than another frame. Alternatively, subjects search for a way to rationalize their choices and one frame may enable a different rationalization than another.

The challenge that Rubinstein sets for himself is 'to model formally procedures of choice that exhibit a certain procedural element and then to investigate whether or not such procedures are compatible with rationality' (p. 25). Various chapters of his book describe different aspects of and approaches

to this problem. The first half devotes itself to choice problems (one-person decision problems), while the second half addresses these issues in a game context.

Chapter 2 deals with a formalization of the idea that individuals simplify problems by using the notion of similarity. A complicated problem may be related to a similar, but less complicated one. For example, in choices between lotteries, probabilities may be judged to be similar or prizes may be considered to be so. As in one case one similarity relation may be considered more prominent and in the other case another, this provides an explanation for choice behavior that conflicts with that of the rational agent but that is more in agreement with actual behavior.

Chapters 3–6 deal with various aspects of knowledge: they review the basic models of Hintikka and Kripke and consider extensions (non-partition models) which could then explain such phenomena as pure speculative trade. Standard models assume perfect memory, i.e. knowledge once acquired remains available intact. Chapter 4 discusses models of imperfect recall and shows again that these lead to qualitatively different conclusions such as, for example, the possibility to improve payoffs by making use of a randomizing device. As the discussion of the 'absent minded driver paradox' illustrates, such models raise challenging conceptual issues. Chapter 5 considers the question of what knowledge to acquire, given the constraint that only a limited amount can be acquired. Furthermore, it gives a specific example of how economic institutions could be explained from such bounds on knowledge. Concretely, if some consumers face knowledge constraints, a seller could increase profits by deliberately reducing market transparency and making price comparison more difficult. Chapter 6 discusses similar issues in the context of the theory of the firm: how to structure communication between units of the firm when gathering and communicating the information is costly?

The chapters 7–10 deal with games. In chapter 7 each player is characterized by a simple rule (try out all of your actions once, see which action did best and continue to play that one) and it is investigated what outcomes will result in that case. The procedures lead to an equilibrium, but this need not be the same as the usual Nash equilibrium. Chapter 8 focuses on the complexity of repeated game strategies. From the so-called 'folk theorem' it is well known that cooperative outcomes may be obtained in repeated games when players threaten to punish deviations by moving to unattractive outcomes. In equilibrium such punishments will, however, not be carried out and when memory is costly, players will be tempted to drop the punishment clauses, making the original equilibrium unstable. The chapter addresses the question of what outcomes can be sustained when players are not just concerned to maximize payoffs, but also want to use simple (low memory) strategies. It is shown that the folk theorem does not survive. Chapter 9 addresses the question of whether the models from the book could offer a convincing explanation for the 'finite horizon paradoxes' from game theory. While models can be constructed that

produce cooperation in the finitely repeated prisoner's dilemma, Rubinstein himself is not convinced of their relevance, arguing that we do not yet have a good understanding about the reasoning procedures of humans in finite horizon games. Chapter 10 briefly discusses some 'computability' issues, arising out of the question 'is there a rational player?' but this part seems only loosely related to the rest of the book.

The book concludes with a very informative discussion between the author and the founding father of the field, Herbert Simon, about what the present book has contributed and about the aims of economic theory more generally. In essence, Simon accuses Rubinstein of 'armchair theorizing' on the basis of casual empiricism; he argues that the latter is no good basis on which to erect a theory and urges for more detailed observation and experimentation first. While admitting that the departures of the rational actor model have to be based on some empirical observations about actual behavior, especially since the space of possible deviations is so large, Rubinstein defends himself in this debate by taking the theorist's position that the logic underlying these procedural models deserves to be explored as it clarifies the concepts that we use and sharpens our intuitions. Hence, while Simon's first priority is to find out the kind of reasoning procedures that people actually use (and why and how these are shaped by experience and the social environment), Rubinstein' s is to analyze properties of tractable processes that are inspired by, and relate in their basic concepts to, actual decision-making processes. Such discussions on priorities are, of course, not new, not even in game theory where the founding fathers discussed them (see Von Neumann an Morgenstern (1944 Sect. 2.1.).

Eric van Damme
Tilburg University
eric.vandamme@kub.nl.

REFERENCES

Selten, R. (1998) 'Features of experimentally observed bounded rationality', *European Economic Review* 42: 413–37.

von Neumann, J. and Morgenstern, O. (1947) *Theory of Games and Economic Behaviour*, Princeton: Princeton University Press.