

## INTRODUCTION

JOHN O. LEDYARD AND THOMAS R. PALFREY

*California Institute of Technology, Pasadena, California 91125*

This special issue of *Games and Economic Behavior* is devoted to what we call experimental game theory, and the intention is threefold. First, we want to provide a window into one of the exciting new areas in game theoretic research. Second, we try to collect in one place some of the best current work that spans many of the interesting topics currently being studied. Third, we hope to encourage more game theorists to think about the behavioral and empirical content of their models, particularly in relation to data from carefully controlled laboratory experiments. The belief is that theory and experiment do not develop independently and the expectation is that better interaction will ultimately lead to both better theory and better experiment.

Game theory and experimental research are two rapidly growing areas of the social sciences, so it should come as no surprise to anyone that there is a significant overlap in what is being studied. While game theory is of considerable interest and value as a normative theory, its greatest application to economics and social science may well be as a predictive theory. For one example, we already see widespread attempts to use game theoretic reasoning to create policy prescriptions and to justify legal decisions. We see this being done in a wide range of applied problems of industrial organization, regulation, litigation, labor negotiation, auction design, and so forth. Implicitly underlying these attempts is the supposition that the theory is good at predicting behavior across the wide variety of settings to which it is being applied.

For a second example, an entire subfield of economics has grown around the use of game theory as a predictive tool. Mechanism design uses game theoretic equilibrium concepts (such as Nash equilibrium, or iterated removal of dominated strategies, or refinements of these) as constraints on a planner who is trying to design institutions to implement particular social choice functions. To the extent that the

goal of this approach is to provide a scientific foundation for the rational choice of institutions and procedures, the entire exercise will be on shaky grounds if the equilibrium concept is not a good predictor of behavior in the particular domain under consideration. Thus, if policy and legal analysis or mechanism design theory are to be useful, it is important to have a good predictive theory of strategic behavior which can be consistently applied across all games in a domain. At the very least, we need to know where the theory's strengths and weaknesses lie. Experimental game theory is studying precisely how good (or how bad) the theory is and under what conditions.

While many experiments can be, and sometimes are, carried out without any real theoretical analysis, it is increasingly the case that research involving carefully designed and controlled laboratory experiments with financially motivated human subjects is directly inspired by game theoretic models of strategic interaction between economic decision makers. This experimental research not only provides evidence about the value and accuracy of game theoretic models of equilibrium behavior as predictive tools, but also provides directions for interesting avenues of theoretical research by identifying the domains where the current theory performs well and where it does not perform well. Theoretical game theory should be studying precisely how well the experimental research is designed and how to fit the theory closer to those facts that have been sufficiently replicated.

If we were to classify the main thrusts of current research in experimental game theory, we would identify the following three categories:

1. *Development and Testing of Alternative Models of Behavior in Games.* This involves a combination of theory, experiment, and econometrics. There are many obvious deficiencies of Nash equilibrium (and its refinements) as a predictive tool. Thus it makes sense to develop alternative models of behavior in games and contrast their predictions to the predictions of Nash equilibrium. It is typical in such papers to attempt, using standard econometric models, to sort between systematic deviations from what is usually considered "rational" behavior or "equilibrium" behavior and those deviations that can be attributed to what is essentially statistical noise. For example, in early auction experiments attempts were made to estimate bidding functions in sealed bid auctions and to identify not only aggregate deviations from the predicted bidding functions, but also the heterogeneity of bidding functions across bidders. The selections by McKelvey and Palfrey and Stahl and Wilson represent an attempt to modify the theory to better fit the data. They take different approaches to studying such statistical variation in the context of simple matrix games, and they fit experimental data to some specific structural models.

2. *Learning in Games.* Experimental games are nearly always conducted by having subjects play the same game many times. The idea behind this is borrowed from earlier economics experiments, where it was found that convergence to competitive equilibrium, while pretty fast, did not happen instantaneously. Presumably, convergence to equilibrium in games, if it occurs at all, will also take time as subjects learn how to solve the problems they face. This is also a very active area of theoretical research. Models are being developed and experiments are being conducted to develop a better understanding of the dynamics of learning in games. Andreoni and Miller study whether certain patterns of bidding behavior that have been observed in experiments can be explained by an evolutionary model of learning. Ochs explores some aspects of the dynamics of convergence in several experimental variations of the matching pennies game and finds that the behavior of the subjects in those games does not appear to converge to the predicted Nash equilibrium. McKelvey and Palfrey identify some systematic trends in their estimated parameters that are suggestive of some kind of learning process. Most of the other papers also identify systematic features of the time-series data, which provide some casual evidence about the nature of learning in experimental games.

Related to the issue of learning is behavior in games (and making predictions) when there are multiple noncooperative equilibria. This arises in coordination games, signaling games, bargaining games, repeated games, and many other settings. Theoretically, there is no consensus about what one might observe in games like this. The combination of experimental research and careful examination of the data from well-designed experiments seems to be the most promising route we have right now for better understanding play in games with multiple equilibria.

3. *Bargaining and Related Applications.* Not surprisingly, many "applied" game theory experiments study bargaining behavior in one form or another.<sup>1</sup> In most bargaining experiments, there is a conflict between the equilibrium predictions of noncooperative equilibrium, when choice behavior is sequentially rational, and the concepts of efficiency and fairness. Experimental research has typically found the greatest divergence from game theoretic predictions when such a conflict arises. There have been many experiments, especially in the areas of bargaining and the voluntary provision of public goods, demonstrating

<sup>1</sup> This is not to say that there are not other areas of applied game theory that are being examined in the laboratory. For example, other areas of specific application to economics which are being studied in game theory experiments include auctions, oligopoly, political economy, and public goods provision, but these are not well represented in this volume.

this divergence. However, although there is a recent surge of experimental work on these problems the systematic exploration of these issues remains largely undone. In this issue there are four bargaining experiments (Knez and Camerer; Bolton and Zwick; Berg, Dickhaut, and McCabe; and Battalio and Van Huyck) and one public good experiment, which can equally well be reinterpreted as a multilateral bargaining game over a pie of unknown size (Budescu, Rapoport, and Suleiman). The first four of these look at some simple variations in The Ultimatum Game, in which one side makes a take it or leave it offer to the other side. This has been of interest in large part because the observed behavior frequently looks much different from the very lopsided prediction of the perfect Nash equilibrium, where one side is supposed to get nearly all the surplus. Some of the experiments are designed to identify systematic psychological factors other than expected utility maximization that could be driving the results.

Our trichotomy of current trends in experimental game theory does not do full justice to the great variety of research being conducted in this area, nor is it the only way to conceptualize the organization of research in the area. There are other ways to categorize experiments. For example, some experiments focus on methodological issues such as optimal experimental design, subject pool effects, and payoff methods. Some (including a couple in this issue) focus almost exclusively on psychological factors. There is a wide range in the complexity and level of abstraction of the games under study. Some experiments are designed to directly test certain game theoretic hypotheses, others are really intended to replicate and extend past findings, and still others are more or less just exploring the frontier, i.e., documenting behavior in a new class of games for which we have no data. A number of recent experiments are motivated more by specific mechanism design issues than by explicitly game theoretic questions. All can provide fascinating observations that game theorists should find provocative.

The obvious links between game theory and laboratory experimentation give game theorists good reason to pay attention to the interesting and growing collection of facts about how people actually behave in strategic settings. We hope this special issue of *Games and Economic Behavior* inspires some interest of this sort among theorists. If history is any guide, good theorists often come up with really good ideas about informative experiments to conduct. Reciprocally, social scientists who concentrate on experimental research can certainly benefit from keeping abreast of the latest theoretical developments.

Finally, a special word of thanks to the many referees who participated in the editorial process. Their timely and thorough help was much

appreciated by the authors as well as by us. We are also especially grateful to Victoria Mason and Marta Goodman for their valuable assistance with the process at Caltech. We thank the editor of *Games and Economic Behavior*, Ehud Kalai, for the good idea of having a special issue on experimental game theory and also thank Marilyn Lott for her excellent editorial and administrative help along the way.