Discussion Paper No. 676

STRATEGIC BEHAVIOR AND COMPETITION: AN OVERVIEW*

by

Ehud Kalai**

February 1986

---

*Papers in this volume were presented at a Northwestern University 1984 summer workshop having the same title. We wish to thank Dean Donald P. Jacobs of the J. L. Kellogg Graduate School of Management for his encouragement and financial support. Also special thanks are due to Professors Kenneth L. Judd and Steven A. Matthews who were the scientific and administrative organizers of this workshop. Finally we wish to thank many friends and colleagues in the M.E.D.S. Department and elsewhere who helped with the editorial and refereeing work.

**Department of Managerial Economics and Decision Sciences, Kellogg Graduate School of Management, Northwestern University, Evanston, Illinois 60201.

Kalai's work is partly supported by a grant from the National Science Foundation N.S.F. No. SES-8409798.

Table of Contents for the JET Symposium Volume Entitled

STRATEGIC BEHAVIOR AND COMPETITION

1. Introduction

In his presentation to the 1980 Journal of Economic Theory symposium, Noncooperative Approaches to the Theory of Perfect Competition, Andreu Mas-Colell writes:

> [I]n order to explain perfect competition . . . [Cournot and Edgeworth], respectively, invented the concepts that modern game theorists call Non cooperative Equilibrium and the Core. . . . Edgeworth's theory of perfect competition has been subject to extensive and rigorous elaboration. . . . Nothing of this sort has happened with Cournot's equilibrium concept.

Mas-Colell then notes a few cases in which the use of the Cournot-Nash Paradigm serves as the basis for a competitive theory; cf., e.g., works due to Gabszwicz-Vial, O. Hart, Novshek-Sonnenschein, Shubik, Shapley, Dubey-Shapley, and Jaynes-Okuno-Schmeidler). He presents the 1980 JET symposium volume as a contribution to what might be thought of as the Cournot-Nash foundation of competitive theory.

Six years later, we find ourselves in the middle of an explosion of research activity on the role of strategic behavior in economics. An attempt to survey the applications of Cournot-Nash equilibrium and Harsanyi-Bayesian-Nash equilibrium in economics would be a major undertaking. The extent of the strategic behavior revolution is substantial. It has led to some fundamental methodological changes in major areas of economic theory including general equilibrium, decentralization, industrial organization, and social choice. In the present symposium volume, a small sample of the current state-of-the-art contributions to strategic behavior and competition are presented. Given the intensity of research in this area, "state of the art" could quite possibly be "out of date" by the time of publication.

I believe that applying game theory to economics is more than a passing fad. What we are witnessing may be the social sciences analog to the long-standing interaction between mathematics and physics. As will become apparent to a reader of this volume, we see here mutual nourishing between economic theory and mathematical game theory. It is by now obvious that the language and methodology of game theory is useful in the analysis of economic problems. And, in the other direction, we see that economists--by selecting from various game theoretic models the ones that are useful in economics--are carving a path for game theory to follow. In a way, we may be seeing economics helping to create its own branch of mathematics, game theory.

We can use the 1980 and this 1986 symposium volumes as snapshots in order to see the development of the theory of strategic behavior and competition. One observes the following trends. First, there is a growing tendency to use the Harsanyi-Bayesian-Nash equilibrium concept as opposed to the simpler Nash equilibrium concept. (The majority of papers in this symposium use and study this Harsanyi-Bayesian-Nash concept.) This allows researchers to model informational issues in an analytical way. A second development is in the use and investigation of refinements of the Nash equilibrium concept. (The middle group of four papers in the present symposium follow this research strategy.) A third tendency is to use dynamic rather than one-shot or steady-state analysis. (Two of the three groups of papers here consist entirely of dynamic models.) It seems very likely that during the next few years research in this area will be concerned with dynamic aspects of games of incomplete information. In these future works, it is likely that some emphasis on low-complexity strategies and information will be attempted in order to make the models more realistic. Indeed, the last paper in the present symposium takes an important step by finding simple strategies that solve a dynamic oligopoly model with incomplete information.

The contributions to this symposium are organized into three groups. The first group of papers deals with informational incentives issues within the context of general equilibrium theory. The second group introduces and uses more refined notions of Nash equilibrium in order to analyze applications having to do with bargaining and cooperation. The final group of papers deals with dynamic, oligopolistic competition.

2. Equilibrium, Decentralization, and Differential Information

The three papers in this group deal with implementation issues in markets that allow for differential information on the part of their traders. The approach they take is to analyze the relevant issues using the Harsanyi-Nash notion of Bayesian equilibrium.

The first paper, by Postlewaite and Schmeidler, introduces the notion of a differential information economy and addresses the question of which social welfare functions can be implemented in such environments. Essentially their model is of a finite, exchange economy into which a probability space of states of the world is incorporated. The agents' endowments and von Neumann-Morgenstern preferences are random variables which depend on the state of the world. Differential information is modeled by including for every agent a partition of the set of states. The usual interpretation is then that when a state occurs each agent is informed only of the element of his partition containing this state. The fact that the information structure and the random variables are assumed to be common knowledge and the assumption that the agents are Bayesian maximizers yield the notion of Bayesian equilibrium in this context.

The Postlewaite-Schmeidler model includes and generalizes major results of the complete information implementation literature. The authors discuss two varieties of conditions, called monotonicity and self-selection, which are imposed on the social welfare correspondence and on the economy, respectively. These conditions turn out to be important for the implementation question. The monotonicity conditions are generalizations to the incomplete information case of the monotonicity condition studied by Maskin and Maskin-Hurwicz-Postlewaite. The Postlewaite-Schmeidler self-selection condition requires that large subsets of the agents can collectively deduce all the relevant information on their own. The authors show that weak versions of these conditions constitute necessary conditions for the implementability of a social welfare correspondence and that stronger versions of the same conditions constitute sufficient conditions for implementability.

The notion of implementability used by Postlewaite and Schmeidler is an attractive one. It requires that the set of equilibria of the implementing mechanism exactly coincide with the social welfare correspondence. As such it does not suffer from some of the well-known shortcomings of weaker notions of implementation. The authors discuss results which are relevant to these types of issues and in particular to the revelation principle, which is sensitive to a particular choice of the notion of implementability.

The paper by Palfrey and Srivastava studies the effect of differential information on a market with a large number of traders. In the model they develop, each agent is "informationally small." Their significant message is that in such an environment there is no serious social welfare loss due to the differential information. The "informational smallness" assumption is crucial because in situations involving many traders, some of which have relatively "big" information (e.g., traders in the stock market with only one having "inside" information), the effect of information is bound to be significant.

The markets that Palfrey and Srivastava study are similar to the ones studied by Postlewaite and Schmeidler. However, rather than assuming that the partition of each agent is fixed, they assume that each agent draws a random partition from a publicly known probability distribution on partitions. Every agent is informed of the partition he drew and then of the element of his partition containing the state of the world. However, his competitors do not know what he knows or even "how much" (the partition) he knows. The authors chose this particular model because it allows them to replicate the economy so that replicated agents do not necessarily share the same information. Their r-fold replication of the economy is the following. Every agent in the original economy is replaced by r agents. These r agents have the same utility functions, the same endowments, and the same prior probability distribution on states and on private information partitions. A highly replicated economy of this type captures the notion of "informational smallness," because the information that an agent has is likely to be known to other agents (his replicas), yet is not insignificant because ex post (i.e., after the state is revealed) his replicas may be different from him.

Palfrey and Srivastava concentrate on the question of what allocations may be "attainable" in a highly replicated economy under a direct revelation mechanism which they describe. It turns out that under their mechanism the agents' incentives to misrepresent their information go to zero with the increased replication of the economy. Consequently, what society can achieve using the actual reports of the agents is close to what society could achieve if it had the actual aggregate information of the agents.

The paper by Ledyard deals with the scope of the hypothesis that Bayesian equilibrium is an appropriate tool to describe players' behavior. The underlying question may be phrased as follows. Suppose that in some well-defined environment involving uncertainty (Bayesian game of incomplete information), we observe a particular behavior by a group of players. Could such behavior contradict our belief that these players are following some

Bayesian equilibrium strategy of the underlying game if we do not know the beliefs and utility functions of these players? Or, equivalently, we could ask: Is there some observable group behavior that would not be generated as a Bayesian equilibrium of this situation as we vary the utility functions and the beliefs of the players?

These questions are obviously important in testing for the Bayesian-equilibrium hypothesis. If the answers to the questions are positive, then rejecting the hypothesis may be easy. If the answers are negative, however, then one must know the players' utility functions and beliefs before having any hope of rejecting the hypothesis. Using the methodology due to Myerson, Ledyard converts these questions to questions about what behavior may be observed under revelation games. He then obtains negative answers to the questions raised above. Given the negative answers he goes on to obtain even more striking negative statements. Even if ordinal preferences over the pure outcomes were known and the players had a consistent information structure, the only type of behavior that could be ruled out by Bayesian equilibrium considerations is one in which players use dominated strategies. In other words, given the players' ordinal preferences and a behavior which does not use dominated strategies, we could always find cardinalizations of their preferences that would make their behavior consistent with the Bayesian-Nash equilibrium of the given underlying game. As Ledyard argues, the results of this research point out that special care must be taken when testing a hypothesis of Bayesian equilibrium behavior in positive economic theory. In such models, making deductions from stylized facts and casual observations can be dangerous.

3. Rationality, Equilibrium Notions, Cooperation and Bargaining

The paper by Rubinstein takes an important first step in the difficult process of incorporating "bounded rationality" into the formal framework of repeated games. He is able to do so by combining an economic game-theoretic problem using the tools of computer science.

Rubinstein uses the repeated prisoner's dilemma game (with the limit of the average payoff criterion) to illustrate his ideas and carry on the initial analysis of such a model. The implementation of a strategy in this game is carried out by a finite automaton. The automaton consists of a finite number of states with one initial state. For every state of the automaton and an action of the opponent the automaton transits to a new state (according to the fixed rule of the automaton) and outputs an action which is a function of the new state. Following this process repeatedly, with a specification of an initial state and reacting to the opponent's sequence of actions, the automaton generates a sequence of plays in the repeated game.

An equilibrium solution in this model consists of a pair of automata, one for each player. At equilibrium, the automata are required to use all of their states infinitely often and also to satisfy lexicographically the following two optimization criteria: (1) No player can increase his payoff by a multilateral change of his machine; (2) No player can reduce the number of states used by his machine without strictly decreasing his payoff. Thus the notion of bounded rationality is modeled into the solution concept, but the complexity costs are assumed to be secondary to "real" profits.

It turns out that the set of payoffs induced by Rubinstein's solutions includes more than just the noncooperative payoff, but does not include all of the possible cooperative payoffs (the individually rational feasible payoffs) and thus can be viewed as an anti-folk theorem. This set includes, in addition to the noncooperative payoff, payoffs on the individually rational segment of the line connecting the two payoffs that correspond to having one player cooperate and the other defect. As the author suggests, the meaning of his results, and further formal studies of bounded rationality would hopefully follow in order to bring us better understanding of these important issues.

The two papers by Grossman and Perry study a refinement of the Kreps-Wilson notion of sequential equilibrium within the context of signaling and bargaining games, and generalize the results of Rubinstein regarding full information dynamic bargaining games. The equilibria in their refined class are called perfect sequential equilibria. In order to define this notion they first extend the set of strategies to what they call metastrategies. In addition to being a function of histories of past plays, a metastrategy is also a function of the beliefs the player has regarding his opponent's type. Bringing in the beliefs explicitly into the domain of the strategy function allows the authors to include restrictions on the beliefs at equilibrium in a formal and effective way. One specific restriction on beliefs that they require at perfect sequential equilibrium is the following. Suppose that a player observes an out-of-equilibrium move by his opponent. How should he modify his beliefs regarding his opponent's possible types? We can partition the opponent types into two possible sets K and K^c where K includes precisely the types that would benefit from the change of beliefs of the player (given his equilibrium metastrategy). It is then required that the player's beliefs be modified by conditioning them on the event that the opponent's type was drawn from the benefiting set K (when not empty).

Grossman and Perry then proceed to study the performance of this equilibrium notion on a discounted sequential bargaining game with incomplete information on one side. A deterministically known seller bargains with a buyer who can be of many possible types. The seller and the buyer alternate in making price offers to each other. After an offer is made by a player his opponent has the option to accept it and terminate the game or to reject it and make a counter offer. Just as in the Rubinstein deterministic model of this kind, the game stops only if an offer is accepted, and, because of the discounting, both players have an incentive to end it early.

In the deterministic version of this game Rubinstein has shown that there is a unique subgame perfect equilibrium which depends on the impatience rate of the bargainers. An objective of Grossman and Perry is to generalize this result to the one sided incomplete information case by using more sophisticated equilibrium notions. They restrict themselves to a special type of equilibrium which they show to be essentially unique in their incomplete information bargaining game and then show that every perfect sequential equilibrium is of this sort. It turns out that in any sequential equilibrium the buyer types are partitioned into three groups, A, B, and C, depending on their impatience in the bargaining. The buyers in group A accept the seller's offer immediately; the buyers in group B respond with an acceptable counteroffer; and the buyers in group C respond with an unacceptable counteroffer. An acceptable offer reveals to the seller that the buyer belongs to set B. However, in a perfect sequential equilibrium the seller cannot revise his beliefs in an arbitrary manner and the authors show that he cannot credibly threaten to reject an offer above the discounted value of the game in which he moves first and the buyers belong to the set B. If the seller gets an unacceptable offer he revises his belief accordingly to C, and the whole process repeats itself with the seller's new belief.

The perfect sequential equilibrium notion is strong. It refines other known equilibrium restrictions in the current literature and generates intuitive results in the games discussed by Grossman and Perry. However, as the authors point out, it may be too strong and may fail to exist in some games. They take the viewpoint that this is the type of criterion that should be applied when possible.

The paper by Gul, Sonnenschein and Wilson deals with a single producer selling a durable good to infinitely many buyers with different valuations for his good. It formally confirms the Coase conjecture that the monopolist will sell at a price close to his production cost and also relates (with an interpretation that is given in the next paragraph) to Rubinstein's results regarding dynamic bargaining with complete information. Each one of the many buyers in this market needs at most one unit of this monopolist seller's good. The seller repeatedly makes price offers for the good. After each such offer is made, every buyer (who has not purchased one yet) has the option of buying at the stated price or waiting further. Being a strategic equilibrium analysis, the paper captures the ideas that consumers correctly anticipate the price sequences and maximize according to this anticipation and that the seller actually offers the sequence of anticipated prices. While only the seller has the power to make offers in this model, he is unable to commit in advance to the entire sequence of prices. In the case where his offers are made very frequently, this lack of commitment turns out to override completely the fact that only he makes offers. This happens because buyers know that eventually the monopolist will end up selling cheap to the low valuation buyers, and given the equilibrium strategies of buyers, profit maximization by the seller is inconsistent with prices that start high and fall slowly. (The authors cite recent papers on durable goods monopoly by Bulow, Kahn and Stokey that have concentrated on these issues. This is the so-called Coase conjecture which states that in a durable goods monopoly the seller will sell immediately at a price close to the minimum of the buyers' valuations.)

A second contribution of the Gul-Sonnenschein-Wilson paper is to the dynamic bargaining literature with incomplete information. In the original Rubinstein bargaining model it was assumed that the seller and the buyer alternate in making the offers to each other until one is accepted. As we mentioned earlier, it was shown there (under the assumption that the buyer's valuation is higher than the seller's) that there is a unique subgame perfect equilibrium which is determined by the impatience rate of the bargainers. In recent models of the dynamic bargaining literature (works of Sobel-Takahashi, Cramton, and Fudenberg-Levine-Tirole) as in the current model, the offers are made repeatedly by the seller to a buyer whose valuation is a random variable. Also, all players have the same discount parameters. It is shown in the current model that if the lowest value of the buyers' valuations is higher than the valuation of the seller, then generically there is still a unique sequential equilibrium. This equilibrium depends on the buyers' valuations distribution, the seller's valuation, and the discount parameter. Moreover, the market will clear after a finite number of periods. If a valuation gap between the seller and the buyers does not exist, the analysis and results are more complicated and there may be a (generic) continuum of equilibria; however, conditions are given (relating the strategies of different buyers) under which the generic uniqueness of equilibrium is restored.

Finally, an alternative interpretation of the result discussed in the preceding pargraph shows that for one-sided offer bargaining games with incomplete information on one side, as the frequency of offers increases, the first offer approaches the minimum of the buyers' valuations, and therefore, the expected time that the market is open converges to zero.

4. Dynamic Oligopolistic Competition

The first paper in this group studies the structure of extremal subgame perfect Nash equilibria of discounted repeated Cournot games. It applies an approach developed in area work by Abreu for general discounted repeated games, which focuses on optimal (in sense of most severe) punishments and "simple strategy profile." He shows that the latter family suffices to generate all subgame perfect equilibrium payoff, and even subgame perfect equilibrium outcome paths. Such profiles are simple because apart from the main equilibrium path, they involve only finitely many paths, namely, one (punishment) path for each player. Any deviation by player i from prescribed play triggers off the single punishment (path) for player i, any deviation from which is responded to in the same manner, and so on. These strategy profiles are easily checked for subgame perfection, making the calculation of optimal punishments more tractable than it would otherwise be.

The present Abreu paper addresses symmetric Cournot games. The author first studies symmetric paths. So-called stick-and-carrot punishments play a central role here. In the first phase of such punishments, production levels are very high and unprofitable; in the second phase firms return to the most attractive stationary collusive output levels. These punishments are easy to calculate, they are more severe than Cournot-Nash reversion (and thus sustain more collusion) and moreover they are optimal symmetric punishments. (For some values of the parameters of the game they are actually globally optimal.)

In general, however, it is shown that optimal punishments must be asymmetric. In this setting, it is argued however, that one can restrict one's attention to the case where the asymmetries are only between the deviating (punished) firm and the others (so that there is symmetry among the nondeviating firms). The stick-and-carrot structure of the symmetric case is partially preserved. Namely, the payoff from the second period and on is not interior in the set of perfect equilibrium payoffs and it is "locally Pareto optimal." However, for both asymmetric Pareto optimal paths and asymmetric optimal punishments, the structure is complicated and a high degree of nonstationarity is involved.

The two papers by Stanford study reactive strategies subgame perfect equilibrium in a model of repeated Cournot duopoly. Reactive strategies are simple in the sense that a player's action in a given period depends only on the action of his opponent in the previous period. Examples of such strategies are various types of "tit-for-tat" strategies and "conjectural variation strategies" as applied to repeated games. A long standing open question posed by Friedman was whether in a discounted duopoly Cournot game there exists a subgame perfect equilibrium which consists of reactive strategies. Stanford's first paper gives essentially a negative answer to this question. He shows that the only equilibria of this type are trivial in the sense that they must coincide with the repeated one shot Cournot-Nash play.

In the second paper, Stanford studies reactive strategies in repeated Cournot duopoly games but with the limit of the average payoff criterion. He first studies linear reaction strategies and shows that in the context of the limit of the average payoff criteria these strategies constitute a positive answer to Friedman's open question. A linear reaction strategy is composed of three parameters, a quantity assigned to the player, a quantity assigned to his opponent, and a number representing the player reaction (punishing) rate. When playing this linear reaction strategy the player starts by playing his assigned quantity. In subsequent periods he deviates from his assigned quantity by his punishing rate times the deviation of his opponent from his assigned quantity in the previous period. It is shown then, that if the parameters of two reactive strategies are chosen appropriately to supplement each other, then subgame perfect equilibrium results. Moreover, if severe reaction rates are chosen then highly collusive production levels (close to the monopoly one) result, zero reaction rates yield the Cournot quantities, and negative reaction rates can yield quantities close to the competitive ones. In addition to generating a large set of subgame perfect payoffs, the linear reaction function exhibits nice convergence properties. On the equilibrium path, the players go on producing the stationary quantities assigned to them. But they also have the property of continuously punishing and correcting defection from the equilibrium quantities. In particular, any deviation from the equilibrium quantities will result in a sequence of quantities that "monotonically" converge back to the equilibrium ones.

Stanford also shows that this property is shared by a larger family of strategies. Namely, any subgame perfect equilibrium which consists of continuous reactive strategies must have a pair of stationary quantities which are fixed points for the equilibrium strategy and for which the fixed point is asymptotically stable in the same self-correcting sense as for the linear reaction strategies.

The paper by Abreu, Pearce and Stacchetti was put at the end of the symposium issue because it captures three important ingredients that future research is likely to follow. It is a dynamic, competitive model with players having imperfect ability to monitor their rival actions. Moreover, the model's assumptions about players' informational requirements are modest and quite realistic. The plausibility of the analysis is underlined by the simplicity of the results that emerge.

In the Abreu-Pearce-Stacchetti model, a finite number of firms repeatedly play a Cournot quantity game with stochastic noise in the price. Specifically, in each period the price that forms is a random variable whose distribution depends upon the aggregate level of production in this period. Firms do not observe past quantities produced by their rivals but only the market prices. Thus, after every history of the game a firm's strategy dictates a level of production which may depend only on the history of the market prices and its own history of production levels. This game generalizes the Porter and Green-Porter games of the same type. In explaining solutions to this game, the authors dispose of a number of major restrictions placed on equilibria in earlier studies.

Their main objective is to find the optimal symmetric sequential equilibria of the repeated game with discounted payoffs. It turns out that there exists a family of very simple equilibria meeting this goal. When playing these symmetric strategies, the firms switch between only two production levels. The decision of which production level to use at a given period depends only upon the assigned production level in the previous period and the observed price of the previous period. Basically, the decision whether or not to "trigger" the punishing quantity in this period depends upon whether the price observed in the previous period fell in a certain critical region of prices.

The authors obtain this result while allowing for more severe punishments than just the Cournot-Nash one (as discussed in the earlier Abreu paper). The technique of proof used by the authors involves reduction of the infinite game to a family of static games which are easier to solve. The authors anticipate that this technique can be useful for the analysis of other repeated games.

5. Summary and Constructive Criticism

As illustrated by the papers in this volume, over the last several years, there has been tremendous progress in understanding strategic behavior both in economic settings and in general contexts. Most of the papers in this volume were presented at a 1984 summer workshop on strategic behavior and competition held at Northwestern University. During and following this workshop, discussions about shortcomings of existing models took place. Some of the major points made by the workshop participants are summarized here with the hope that this constructive criticism may encourage future research.

a. Complexity of the strategies and information. As we incorporate more realistic details (specific rules of the game, information held by the players, dynamics, etc.) into our models, we increase the domains of the players' strategies. Consequently the strategies become very complex. The assumption of full rationality of the players then becomes less realistic. Similarly, as we allow for the players to have incomplete information, we still have a tremendous information requirement imposed on them. In a Bayesian game of incomplete information as in the Schmeidler-Postlewaite, Palfrey-Srivastava, and Ledyard papers, a player must know the information structure of all the other players. This is often unrealistic. The Rubinstein paper deals with the complexity issue head on by incorporating a complexity minimization objective into the solution concept. Other papers in this volume, for example the Abreu, Stanford, and Abreu-Pearce-Stacchetti papers, deal with complexity by pointing out that the solutions they obtain are relatively simple, and thus the complexity issue does not arise. Formal methods of analyzing complexity are beginning to emerge and hopefully will be brought into the analysis.

b. Robustness of the Model. Solutions of extensive form games tend to be very sensitive to the specifications of the game tree. So when we model a situation by such a game, some of the model specifications become critical. As Grossman-Perry and Gul-Sonnenschein-Wilson illustrate in their bargaining papers, the sequencing of offers, the choices of where and how to model the information, and the ability of players to make binding commitments make tremendous differences to the outcome of the game. When the economic situation clearly determines these choices in a unique way (as for example in auctions), then the ability to model these details is a clear asset. However, when these choices are not that clear, then the extensive form game model is less desirable. In such situations it may be better to use solution concepts of the cooperative and semi-cooperative type. For example, Aumann's correlated equilibria may prove to have better robustness properties for these types of applications (assuming that the equilibrium-multiplicity difficulty introduced by these solutions can be disposed of by more general principles). These solutions depend to a greater extent on the feasible set of outcomes and less on the rules of the game.

c. Refinements of the Nash Equilibrium Concept. The concept of a Nash equilibrium is justified by imposing necessary conditions on the solution to a game. However, as has been pointed out by Selten and many other authors, it is not restrictive enough. Refinements of this concept, such as perfect and sequential equilibria, are often still unsatisfactory as illustrated by papers in this volume. It is not clear whether the current on-going search for the best solution concept would actually yield a unique satisfactory choice. It is quite possible that in order to find the correct concept, we have to make distinctions between the situations that the games describe. For example, in the game where players have to choose whether to drive their cars on the left or on the right side of the road, two different solution concepts may apply. If the game has been played repeatedly (even if by different players), the usual Nash solution seems quite satisfactory and the prediction that tomorrow all drivers in the United States will drive on the right side is credible. On the other hand, if people start driving tomorrow for the first time, and no pre-communication took place, then the prediction that they would all be driving on the same side seems less realistic, and solution concepts such as Bernheim and Pearce's notions of rationalizable strategies seem more appropriate. It would be nice if one solution concept would apply to all situations by being sensitive to new types of considerations (for example, "how long has this game been played?"), but it is also possible that until such a unified theory is found, we will build on—as papers in this volume do--a variety of solution concepts chosen appropriately according to the economic situation.