Informational Asymmetries, Strategic Behavior, and Industrial Organization

By PAUL MILGROM AND JOHN ROBERTS*

One of the most active and exciting areas of economic research over the last several years has been the use of noncooperative games of incomplete information to model industrial competition. This work has yielded not only a remarkable number of papers but also several new insights on and explanations of fundamentally important issues. The purpose of this paper is to attempt an appreciation and evaluation of this work. Because most of our individual and joint work since about 1979 has been in this mode, it will come as no surprise that we are proponents of this line of research. However, there are several questions and potential problems that we see as arising in connection with this methodology, and we will attempt to address these.

First, a disclaimer. We are not attempting a survey of the applications of asymmetric information games (AIG) to industrial organization, although we will refer in a highly selective fashion to a number of prominent strands in this literature. (In particular, where any references are provided at all, they are typically only to the earliest contributions to a subject.) Even more, we do not deal with work in which informational asymmetries are important but the analysis is not game theoretic (for example, search and price dispersion, or the early work on the lemons problem and on moral hazard and adverse selection in insurance markets) or with game-theoretic treatments that assume complete information.

*School of Organization and Management and Department of Economics, Yale University, New Haven, CT 06520, visiting the University of California-Berkeley; and Graduate School of Business, Stanford University, Stanford, CA 94305–5015, respectively. We gratefully acknowledge the comments of David Kreps and Robert Wilson, and the financial support of the John Simon Guggenheim Memorial Foundation and the National Science Foundation.

I. AIG Methods and Applications

To get an idea of the role of informational asymmetries in strategic behavior, consider three simple card games. In the first, each player is dealt five cards face up, the players make any bets they want, and then the best hand wins. In the second, each player receives five cards, some of which are dealt face up and the rest face down. Without looking at their hole cards, the players make their bets, then the cards are turned face up and the best hand wins. Finally, the third game is like the second except that players can look at their hole cards. Again there is betting, the hidden cards are revealed, and the best hand wins.

The first game is one of complete (and perfect) information. Everyone knows everything, and as long as we assume that people prefer more money to less, it is fairly trivial to figure out what will happen: there will certainly be no betting, and probably no one will bother to play! Clearly, not all games of complete information are either so uninteresting (witness chess) or so lacking in explanatory power—especially if we consider nonzero sum games and, even more, games with an explicit dynamic structure (Drew Fudenberg and Jean Tirole, 1986a). However, in its informational structure, this game typifies both the sort of game theory that is discussed in intermediate micro texts and, indeed, most of standard microeconomic theory itself.

The second game has uncertainty/informational incompleteness, but no informational asymmetries. Its informational structure puts it in the domain of decision theory and the economics of uncertainty. Games of this sort are useful models for studying such issues as insurance, risky investments, and learning (especially if we revise the game to have the hole cards revealed one at a time, with betting after each is shown). However, its play would not generate any interesting forms of strategic behavior.

The third game involves informational asymmetries: while there is some publicly available information, each player is privately informed about his or her hole cards. (In fact, the informational structure of this game, in which the probability distribution over what the particular private information of the various players could be is common knowledge, corresponds very closely to that in the asymmetric information game models used in most applications to industrial organization.) The existence of this private information can obviously lead to interesting strategic play: bluffing, signaling, reputation building, etc. It is also the reason why poker is of enduring popularity.

As this example is meant to suggest, recognition of informational asymmetries and the strategic possibilities they engender can yield models that begin to capture the richness of behavior that marks the real world. This is the great advantage of these methods: they permit us to model, and thereby start to understand, phenomena that made no sense in terms of complete information analyses or ones based on incomplete but symmetric information (uncertainty).

Perhaps the clearest example of this is predatory pricing. In 1980, the only fully consistent analyses of predatory pricing in the literature (for example, John McGee, 1958, 1980) indicated that predatory pricing could not be expected to succeed, that it was thus not part of a rational competitive strategy, that apparent instances of predatory pricing were consequently likely to be either mistakes or misinterpretations, and that legal prohibitions on predation serve chiefly to protect inefficient firms from the desirable effects of competition. Although these conclusions ran counter to much of the conventional wisdom in the field of industrial organization and left a disconcerting number of mistakes and misperceptions to explain, the logic leading to them seemed compelling. And because no mere fact ever was a match in economics for a consistent theory, these ideas began to represent the basis of a new consensus.

However, these analyses rested on an implicit assumption of symmetric information, and this assumption is crucial. Recent studies by a number of authors (see Roberts, 1987, for a survey and complete references) have relaxed this assumption in various ways and reversed the earlier findings. Pricing below the short-run optimal level aimed at deterring entry, inducing exit, or disciplining rivals so that they compete less aggressively can be part of a rational strategy in the presence of realistic informational asymmetries. Thus we should expect to see predatory behavior being adopted if it is not effectively deterred by legal prohibitions. Moreover, the mechanisms by which these effects come about in these theories correspond well to suggestions found in earlier, less formal discussions of predation in the industrial economics and legal literatures. At the same time, these new analyses indicate that the legal tests that have been proposed for establishing whether predation has occurred may be completely inappropriate in that they may fail in either direction.

Most of the recent, asymmetric information game models of predatory pricing involve generalized signaling: there is an underlying parameter that is of interest to one player (the "receiver") but is not directly observed by this player, and the (costly) actions of the other player (the "signaler" or "sender") can affect the observations made by the receiver. By the choice of actions the sender can thus influence the inferences the receiver makes about the parameter's value and, correspondingly, the receiver's choice of actions. This formulation subsumes the original Spence-type signaling model, in which the signaler (there, the worker) knows the value of the parameter (his or her productivity) and so can condition his or her choice (of education level) on this information. In these circumstances, observation of the sender's choice may allow the receiver to infer the sender's information. It also encompasses "signal jamming" models, in which neither player is informed about the value of the parameter. In these models, the actions of the signaler are not observed by the receiver directly; instead they affect the distribution of a variable that is observed by the receiver and from which he or she must try to infer the value of the parameter. As well, dynamic formulations in which the signaler acts repeatedly and information is revealed over time (perhaps to a sequence of receivers) are also included.

In signaling models of predatory pricing, the sender is the predator firm, the receiver is either the current rival on which the predation is practised or potential rivals that can observe the predator's current behavior, the parameter is a variable that influences the receiver's profit from continued operations, and the signal is the predator's price. The models of Roberts (1986) and Garth Saloner (1986) of predation against a single opponent are Spence-type signaling ones in which the predator is privately informed about demand or cost. The signaling firm is led to lower its price in an attempt to suggest that the value of the parameter is such that either continued competition by the receiver will be unprofitable—thereby inducing exit (Roberts) or encouraging acceptance of a merger offer (Saloner)—or that a reduced level of output would be optimal for the receiver. (See also David Scharfstein, 1984.) Fudenberg and Tirole (1986b) consider a situation where the two firms are symmetrically informed about the random demand, but the predator has an incentive to increase output or lower price unobservably and thereby attempt to bias the receiver's estimates of profitability. Finally, David Kreps and Robert Wilson (1982), our paper (1982b), and Fudenberg and Kreps (1986) consider dynamic models in which the privately informed sender adopts predatory behavior against early entrants, even though it is directly unprofitable, in order to build a reputation for aggressive responses to entry that deters future challenges. See also David Easley, Robert Masson, and Robert Reynolds (1985), where the reputation being built is for having markets into which entry is unprofitable because demand is weak.

Equilibrium in signaling models involves the receiver's having correct conjectures regarding the sender's actions (as a function of the sender's information) and accounting for these in making inferences about the parameter. Thus, both in models in which the sender is uninformed about the parameter and in ones where the sender knows its value, the receiver's estimates are not systematically biased by the signaling behavior. In fact, to the extent that the equilibria are separating (involve a one-to-one map between the parameter's value and the receiver's observations), the receiver will correctly infer the value of the parameter and will thus, in effect, be acting as if he had access to the sender's information. Yet the sender's actions are typically distorted from their full-information levels because the receiver interprets his observations in light of the sender's incentives to attempt to influence his inferences. Thus, for example, in the Roberts model, if the signaler produced the full-information output rather than the equilibrium one, the receiver would interpret the resulting observation as indicating that demand is stronger than it actually is. As a result, the receiver would be less likely to exit, and if it stayed in, it would produce at the higher level corresponding to strong demand.

This property of separating equilibrium has welfare implications: not only is the price lowered during the predatory episode, but there need be no more exit or restriction of output by the prey than there would have been if predation had been effectively forbidden. However, this does not imply that the predatory behavior is socially desirable, because the threat of predation (even if it will fail to induce exit) will deter entry. This is very clear in the Kreps-Wilson and Milgrom-Roberts dynamic reputation models, but is also true in the static models. Thus, legal concern with predation may still be warranted. At the same time, these models indicate that predation may occur without the predator ever pricing below marginal cost or in other ways meeting the tests for predatory behavior often advocated in the literature.

The idea of signaling has also proven fruitful in studying a number of other situations involving pricing under imperfect competition. An early example was our (1982a) rationalization of limit pricing, in which low-cost incumbents are led to price below the short-run monopoly level to signal credibly that the entrant will find their markets unprofitable targets for entry (see our 1982a paper; again see Roberts, 1987, for further references). George Mailath (1985) has used signaling to explain the phenomenon of price wars occurring early in an industry's history: each firm expands output in an attempt to make its competitors believe that its cost are low and that it thus should have a large market share. In a closely related model, Michael Riordan (1985) has used signal jamming to rationalize conjectural variations. A further example is Kyle Bagwell's (1986) signaling explanation of introductory sales: the firm seeks to signal that its costs are low and thus that it will be worthwhile for customers to shop at this firm in the future because its low costs will lead it to set low prices then as well. Our recent treatment (1986a) of both pricing and image advertising as signals for unobservable product quality carries these methods into a multidimensional context.

Models that explicitly incorporate private information but are not of the signaling variety have also been widely used. For example, they have proven useful in examining the incentives to reveal cost, demand and other information to suppliers (Milgrom and Robert Weber, 1982a, b) and to competitors through trade associations and the like (William Novshek and Hugo Sonnenschein, 1982). They have been applied to the problem of maintaining cartel agreements when there are problems in detecting cheating (Edward Green and Robert Porter, 1984, and Dilip Abreu et al., 1986). The burgeoning literatures on nonlinear pricing, priority pricing, durable goods monopoly, auctions, bargaining, etc., are all based in this methodology. Finally, theories of contracting in the presence of informational asymmetries, which are increasingly being based in the methodology of incomplete information game theory, are contributing mightily to our understanding of issues of procurement, regulation, vertical integration, and the functioning and design of economic organizations and institutions more generally. Moreover, all this has been accomplished essentially in the last five years.

The methods of incomplete information game theory have thus allowed us to model formally, often for the first time, issues that are central to industrial organization. Moreover, the behavior that emerges in equilibrium from these models begins to capture, again often for the first time, something of the richness of observed behavior. Indeed, this work holds some promise of yielding a partial and much-belated realization of some of the hopes that were expressed in the early years of game theory that its use would lead to a resolution of the problems of oligopoly and imperfect competition. This certainly (to us) justifies them and their use in industrial organization research.

II. Questions about the AIG Methodology

The questions that can and do arise about these methods seem to us to fall under five headings: the assumption that equilibrium behavior will prevail and the related but more fundamental rationality assumptions; the assumed common knowledge base; robustness of the results; multiplicity of equilibria; and empirical implementation and testing. We discuss each of these in turn.

Almost all of economic theory is equilibrium analysis, and this work is no exception: predictions arise only once equilibrium behavior is assumed. Yet the assumption of equilibrium seems more demanding in many of these incomplete information games than it does in, say, competitive partial-equilibrium models. There seem to be two possible reasons for this.

The first of these relates to the relative specificity of the two types of models. Most standard models of price determination are remarkably incomplete: the timing of actions, the information available to agents when they act, and the consequences of not adopting equilibrium behavior are rarely modeled, and many treatments even leave out such seeming fundamentals as who selects prices, and how the supplies and demands expressed to the market get transformed into actual transactions. This incompleteness facilitates accepting the equilibrium assumption because the model gives nothing else on which to focus. At the same time, it encourages us to comfort ourselves with vague and often unarticulated stories about processes of adjustment that somehow lead to equilibrium. In contrast, the methodology of extensive games forces AIG models to be much more specific and complete. First, this calls to our attention the possibilities of not adopting equilibrium behavior. Secondly, it invalidates appeals to learning through repetition or to adjustment of play over time toward equilibrium, because if there is repeated play, this should be modeled and the resultant game analyzed on its own. As is well known, the equilibria of the repeated-play game may differ substantially from those of the one-shot version.

Of course, to the extent that the assumption of equilibrium seems easier in more standard models because they assume away complicating factors that actually may be important, a finding that the equilibrium assumption is relatively more problematic in asymmetric information-strategic models is hardly a criticism of these models.

The second reason relates to the complexity of the coordination problem in these models. There are often multiple equilibria in AIG models, and one player's adopting his or her strategy from one equilibrium and the others' playing their strategies from a second equilibrium typically does not constitute equilibrium behavior. This problem, of course, is not unique to AIG models: consider the Battle of the Sexes. Still, even if everyone can figure out what strategy n tuples are equilibria, unless there is a unique equilibrium, there are real problems in ensuring that everyone focuses on the same one. There are various stories that game theorists tell in such situations (see Kreps, 1986, for an exposition of these), but none are fully satisfactory. Of course, all these problems disappear if there is a single equilibrium. (We return to the multiplicity issue below.)

Even if the coordination problem can be overcome and if one grants that rational actors would adopt equilibrium behavior, one might still more fundamentally question the rationality assumption as it appears in AIG models.

There is no denying that the sort of inferences, calculations, and forecasts that agents are making in the equilibria of AIG models involve much more sophistication than, say, the agents in an Arrow-Debreu world of complete, competitive markets must show when the equilibrium prices are given. Of course, this latter standard is the extreme, and the demands on rationality increase as soon as we move from this world to ones without complete, perfect markets and symmetric information. Still, equilibrium in AIG models does seem to involve another quantum leap beyond, say, rational expectations models with informative prices. As game models, AIG treatments require players to act as if they anticipate fully the often complex responses of the other players. Further, in AIG models these responses depend on subtle inferences that competitors draw, often by very intricate reasoning, from their conjectures about others' behavior and their observations. The strategic players' decision problems therefore are much more difficult than in more standard models. Correspondingly, the assumption of agents' finding and adopting equilibrium behavior becomes that much more implausible.

In this regard, it seems to us that an appropriate test of the assumptions of rationality and equilibrium is the standard one: if their use aids our understanding, leads to accurate descriptions, facilitates prediction, and generates useful recommendations, then use the assumptions until something better comes along. Nevertheless, the descriptive accuracy of the super-rationality assumption does seem minimal. Indeed, if it did describe actual individuals, the outcome of the game of chess would be totally determinate and obvious, and its play would be as exciting as Tic-Tac-Toe. Given that AIG models have not been around long enough to have yet stood the above test on many occasions, this lack of descriptive accuracy is troubling. Moreover, there do seem to be important phenomena, especially in the economics of organization, that are very hard to explain without retreating at least from the assumption that transferring information, assimilating it, calculating, and deciding can be done instantaneously and without cost.

There have, of course, been some recent models in the AIG framework that include boundedly rational agents. For example, our 1982b model of predation and the Kreps et al. (1982) treatment of the finitely repeated Prisoner's Dilemma involve the possibility of agents' using simple rules of thumb to guide their behavior, and more recent work by a number of game theorists has investigated situations where strategies must be implementable by finite automata or involve limited memory. However, hyperrational agents still play a crucial role in all this work. In the work with which we have been associated, rational agents in equilibrium are led by very complex reasoning to mimic the rule-of-thumb players, while in the later work the machines used to implement strategies are selected by rational agents making the usual sorts of forecasts, inferences, and calculations.

The problem, of course, is that we as yet have little agreement on how to model more descriptively accurate forms of rational behavior, little faith that we can find hypotheses on behavior that will be as tractable and powerful as maximization and equilibrium, and a general fear that by renouncing our standard methods we will forfeit elegance and, in return, get only *ad hockery*.

A related complaint about AIG models is the common knowledge assumption. In applications to industrial organization, it is typically assumed that the private information is of small dimension (for example, cost function parameters) and that the distribution of possible values for this information, as well as everything else in the model, is common knowledge. (Intuitively, an event is common knowledge if each player knows it has occurred, each knows that each knows this, each knows that each knows that each knows this, *ad infinitum*.) The objection to this that one hears is that this is no more realistic than assuming that actual values are common knowledge, and, by implication, analyses based on such an assumption are at least as suspect as ones assuming complete information.

It seems that there are two possible interpretations of this complaint. The first is that assuming that beliefs over the values of the underlying parameter are common knowledge is too simplistic, because there might reasonably by uncertainty about beliefs. Formally, of course, there is no need to assume that the type space is so simple: your type could involve not only your costs but also, for example, your beliefs about what others believe your costs are. Such relatively complicated spaces have in fact been successfully used in applications to industrial organization (see our 1982b paper, appendix B); however, generally they will be incompatible with obtaining the sort of monotonicity properties that have proven so fruitful in signaling models.

The second interpretation relates to the concerns with rationality. The type spaces in AIG models rapidly become extremely complicated mathematical structures as the level at which the uncertainty is assumed to lie is pushed back. For example, if there are two matrices that might give the payoffs, then the possible beliefs over which of these prevails correspond to points in the unit interval, beliefs over beliefs are increasing functions from $[0,1]$ into itself, the next level of beliefs are the measures on this space of functions, and so on. Moreover, whatever the assumed type space, to use AIG methods the distribution of types must be taken to be common knowledge, as must, in fact, the full game tree and the payoffs at each terminal node, the number of which is a function of the number of types. To assume that real people make calculations over such complex spaces seems again to strain the limits of credibility.

This latter point might be thought to be a problem that is particular to the AIG methodology. However, this is not quite the case. One can argue that *any* game-theoretic method of analyzing rational behavior in multiperson decision problems must start from an assumption of what is common knowledge among the agents (see, for example, Wilson, 1986). Thus, AIG models are in no way special in this regard; instead, this methodology, in which the common knowledge assumptions are completely explicit, has made the more general necessity of such assumptions apparent.

While we agree with this position, some of our recent work suggests that it may be important to study the extent to which this point is valid once one moves to an assumption of bounded rationality. In particular, in our 1986b paper, we consider a decision maker whose rationality is decidedly bounded; in particular, he or she does not even know the type space, let alone have a common knowledge distribution over it reflecting beliefs. Yet this decision make can do very well in certain situations by being skeptical and by inducing competition between informed interest parties.

Another potentially disturbing aspect of these models that is less obvious is the apparent sensitivity of the results to alterations in what one might think is fine structure of the models. This is especially clear in finitely repeated games. Kreps-Wilson and our paper (1982b) have shown that the introduction of a tiny bit of private information into such a model can radically change its equilibrium outcomes as the finite horizon becomes "long." However, this technique may be too powerful: Fudenberg and Eric Maskin (1986) have shown that by introducing the right kind of informational asymmetries, one can obtain a Folk Theorem result that almost anything can be made an equilibrium (see also John Ledyard, 1986). What is needed is some way to determine which informational asymmetries have survival value in the sense that if people ascribe positive probability to a variety of possible types, then those with survival value are the ones whose behavior is mimicked (as are the tit-for-tat players in Kreps et al.). See our paper (1982b) and Robert Aumann and Sylvain Sorin (1986) for beginnings in this direction.

Even in one-shot games there are important discontinuities. For example, Mailath has shown that there is a unique separating sequential equilibrium in our limit-pricing model with a continuum of possible values for the incumbent's cost. It involves all but the highest-cost type of incumbent producing more than the monopoly output. This equilibrium is unaffected by changes in the relative likelihood of the incumbent's possible types, so long as the support of the beliefs is unchanged. But suppose that the weight becomes concentrated in the limit on the lowest cost type. At this limit, sequential equilibrium requires a discrete jump in the firm's output choice to the simple monopoly price and output. This same example illustrates the very important point that the addition and deletion of initial nodes (types) having zero prior probability can radically affect the solution of these models. Similarly, the inclusion of the option to take actions that might, on first blush, appear to be dominated—such as publicly burning money—can affect the solution (see our 1986a paper).

The question here is whether these features are artifacts of the models or whether they correspond to something real. If it is the latter (as we suspect it is—the range of possibilities that people consider will affect their decisions and actions), then this sensitivity of behavior is unfortunate for those who hope to draw easy general conclusions, but must be faced. In particular, since one doubts that everyone is always certain that everyone else is super-rational and, more generally, that the model of the world they are using is absolutely accurate, it is crucial to investigate models including informational asymmetries. In this regard, the study of how reasonable forms of bounded rationality affect the sensitivity of models assuming hyper-rationality seems especially important.

A further complaint against the AIG models that have been used in industrial organization is that they admit such a great multiplicity of equilibria that, while they might possibly explain certain phenomena, they are of limited value for prediction because so many patterns of behavior are consistent with equilibria. This is a criticism of which both formal game theorists and those interested in applying AIG models have been very aware, and members of both groups have devoted significant effort to meeting it.

The feature that makes the multiplicity problem especially acute in AIG models compared to complete information games is the multitude of beliefs about the relative likelihood of various events that can be held *off the equilibrium path* and can be consistent with equilibrium. Sequential equilibrium requires that, at each decision point (and not just those reached under equilibrium play), each player find that continuing to use his or her equilibrium strategy is optimal given his or her beliefs about what has happened so far and what the others know. In situations without informational asymmetries, this subgame perfection requirement often has the power to delimit actions quite narrowly. However, in AIG models, it is typically the case that the set of optimal actions varies widely as we alter the players' beliefs about how play has proceeded and what information the others may have. Further, Bayesian updating gives no restriction on beliefs off the equilibrium path, since in such situations we would be conditioning on probability zero events. Thus, a wide variety of beliefs —and a correspondingly wide variety of optimal actions—are consistent with equilibrium.

Of course, if the indeterminacy were confined to situations that are not observed in equilibrium (and if one decides to focus on equilibria), there would be no great problem. However, this cannot be the case. One of the great virtues of game-theoretic methodology is the requirement that behavior be specified in *all* eventualities, not just under the particular circumstances corresponding to some putative equilibrium. Then equilibrium is determined endogenously by considering the implications of deviating from the specified behavior. These implications are fully determinate exactly because players' strategies specify what actions to take at *every* decision point. Thus, in effect, off-the-equilibrium-path behavior determines equilibrium behavior. Correspondingly, in AIG models there can be many equilibria, each supported by different beliefs off the equilibrium path and by the behavior these generate.

Given this diagnosis, most of the effort aimed at overcoming the multiplicity problem has been directed toward narrowing the set of out-of-equilibrium beliefs that are to be considered to represent *reasonable* inferences to have made after observing others' choices (see, for example, Kreps). In some circumstances, relatively simple and intuitively appealing conditions suffice to reduce the set of equilibria significantly, and even to generate uniqueness. For example, in signaling games it is often enough to assume that if a particular message could never be part of a best response for some type of sender but could for others, then if this action is observed, receivers do not attribute it to the sender type for whom it is never a best response. However, in other situations, either the required conditions are very hard to understand or simply are not yet known (for example, bargaining with private information on both sides). Moreover, to the extent that obtaining a small set of equilibria involves agents' making particular, highly sophisticated, extremely subtle inferences from observations, the questions raised about the equilibrium and rationality assumptions arise again.

Finally, there is the issue of testing these models. There has, to date, been relatively little empirical work based on AIG models. In part, this may simply be a matter of time. These models are new and have for the most part been developed by theorists; they may simply not yet have reached the empiricists' agenda. However, it seems that there are inherent difficulties in testing. Some of these arise from the multiplicity of equilibria, but the more central one is that the central object in the theory is, by its very nature, unobservable. How, for example, does one obtain information on what a firm's beliefs about rivals' costs were when it took a particular pricing or entry decision? Note that the sensitivity problem discussed above exacerbates this difficulty, because the predicted results can depend so finely on both the distributions over private information and the fine details of the modeling.

This suggests two approaches: careful case studies and experimental work. The former is clearly very costly, which may be why, to our knowledge, relatively few such studies have been attempted. However, those that have been done (for example, Mark Wolfson, 1985, on contracting; Michael Staten and John Umbeck, 1982, on shirking in labor markets; Robert Porter, 1983, on cartel maintenance) are generally supportive of the theory, as are various other studies not specifically based on theories involving informational asymmetries (for example, Malcolm Burns, 1986). Meanwhile, a significant part of the experimental work on bargaining, auctions, and various market institutions incorporates private information, and while little work directly aimed at examining the influence of information asymmetries on strategic behavior in industrial organization settings has yet been done, those studies that do exist (for example, Ross Miller and Charles Plott, 1985; Colin Camerer and Keith Weigelt 1986) again tend to give some support to the theory. These results, and the exciting insights that the theory has offered, justify the attention that AIG models in industrial organization theory have received and, we trust, will continue to receive.

REFERENCES

Abreu, Dilip, Pearce David and Stacchetti, Ennio, "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory*, June 1986, *39*, 251–69.

Aumann, Robert and Sorin, Sylvain, "Bounded Rationality and Cooperation," paper presented at annual meeting of the American Economic Association, New Orleans, December 29, 1986.

Bagwell, Kyle, "Introductory Price as a Signal of Cost in a Model of Repeat Business," Studies in Industry Economics, No. 130, Department of Economics, Stanford University, 1985.

Burns, Malcolm, "Predatory Pricing and the Acquisition Cost of Competitors," *Journal of Political Economy*, April 1986, *94*, 266–96.

Camerer, Colin and Weigelt, Keith, "Experimental Tests of a Sequential Equilibrium Reputation Model," mimeo., Graduate School of Business, New York University, September 1986.

Easley, David, Masson, Robert and Reynolds, Robert, "Preying for Time," *Journal of Industrial Economics*, June 1985, *33*, 445–60.

Fudenberg, Drew and Kreps, David, "Reputation and Multiple Opponents I: Identical Entrants," mimeo., Graduate School of Business, Stanford University, May 1986.

______ **and Maskin, Eric,** "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, May 1986, *54*, 533–54.

______ **and Tirole, Jean,** (1986a) *Dynamic Models of Oligopoly*, New York: Harwood Academic, 1986.

______ **and** ______, (1986b) "A 'Signal-Jamming' Theory of Predation," *Rand Journal of Economics*, Autumn 1986, *17*, 366–78.

Green, Edward and Porter, Robert, "Noncooperative Collusion under Imperfect Information," *Econometrica*, January 1984, *52*, 87–100.

Kreps, David, "Out of Equilibrium Beliefs and Out of Equilibrium Behavior," Working Paper, Graduate School of Business, Stanford University, September 1986.

______ **and Wilson, Robert,** "Reputation and Imperfect Information," *Journal of Economic Theory*, August 1982, *27*, 253–79.

______ **et al.,** "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma," *Journal of Economic Theory*, August 1982, *27*, 245–52.

Ledyard, John, "The Scope of the Hypothesis of Bayesian Equilibrium," *Journal of Economic Theory*, June 1986, *39*, 59–82.

McGee, John, "Predatory Price Cutting: The Standard Oil (N.J.) Case," *Journal of Law and Economics*, October, 1958, *1*, 137–69.

______, "Predatory Pricing Revisited," *Journal of Law and Economics*, October 1980, *23*, 289–330.

Mailath, George, "Separating Equilibria in Signaling Games: Incentive Compatibility, Existence with Simultaneous Signaling, Welfare and Convergence," unpublished doctoral dissertation, Princeton University, 1985.

Milgrom, Paul and Roberts, John, (1982a) "Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis," *Econometrica*, March 1982, *50*, 443–59.

______ **and** ______, (1982b) "Predation, Reputation and Entry Deterrence," *Journal of Economic Theory*, August 1982, *27*, 280–312.

______ **and** ______, (1986a) "Price and Advertising Signals of Product Quality," *Journal of Political Economy*, August 1986, *94*, 796–821.

______ **and** ______, (1986b) "Relying on the Information of Interested Parties," *Rand Journal of Economics*, Spring 1986, *17*, 18–32.

______ **and Weber, Robert,** (1982a) "A Theory of Auctions and Competitive Bidding," *Econometrica*, September 1982, *50*, 1089–122.

______ **and** ______, (1982b) "The Value of

Information in a Sealed-Bid Auction," *Journal of Mathematical Economics*, June 1982, *10*, 105–14.

Miller, Ross and Plott, Charles, "Product Quality Signaling in Experimental Markets," *Econometrica*, *53*, July 1985, 837–72.

Novshek, William and Sonnenschein, Hugo, "Fulfilled Expectations Cournot Dupoly with Information Acquisition and Release," *Bell Journal of Economics*, Spring 1982, *13*, 214–18.

Porter, Robert, "A Study of Cartel Stability: The Joint Executive Committee, 1880–1886," *Bell Journal of Economics*, Autumn 1983, *14*, 301–14.

Riordan, Michael, "Imperfect Information and Dynamic Conjectural Variations," *Rand Journal of Economics*, Spring 1985, *16*, 41–50.

Roberts, John, "Battles for Market Share: Incomplete Information, Aggressive Strategic Pricing, and Competitive Dynamics," in T. Bewley, ed., *Advances in Economic Theory*, Cambridge: Cambridge University Press, 1987.

______, "A Signaling Model of Predatory Pricing," *Oxford Economic Papers* November 1986, *38*, Suppl., 75–93.

Saloner, Garth, "Predation, Merger and Incomplete Information," Working Paper No. 383, Department of Economics, Massachusetts Institute of Technology, September 1986.

Scharfstein, David, "A Policy to Prevent Rational Test-Market Predation," *Rand Journal of Economics*, Summer 1984, *15*, 229–43.

Staten, Michael and Umbeck, John, "Information Costs and Incentives to Shirk: Disability Compensation of Air Traffic Controllers," *American Economic Review*, December 1982, *72*, 1023–37.

Wilson, Robert, "Competitive Strategies in Business," Working Paper, Graduate School of Business, Stanford University, September 1986.

Wolfson, Mark, "Empirical Evidence of Incentive Problems and Their Mitigation in Oil and Gas Tax Shelter Programs," in John Pratt and Richard Zeckhauser, eds., *Principals and Agents: The Structure of Business*, Boston: Harvard Business School Press, 1985.