Journal of Economic Perspectives—Volume 11, Number 3—Summer 1997—Pages 159–174

A Nobel Prize for Game Theorists: The Contributions of Harsanyi, Nash and Selten

Faruk Gul

In 1994, the Royal Swedish Academy of Sciences awarded the Nobel Prize in economics to John C. Harsanyi, John F. Nash and Reinhard Selten for their contributions to noncooperative game theory. This was the first time the academy had awarded a Nobel Prize to game theorists. It is difficult to imagine more appropriate recipients. The purpose of this essay is to discuss and to celebrate the contributions of these three researchers to game theory and economic analysis.

I will not attempt to provide a full catalogue of the achievements of the three Nobel prize winners, nor will I attempt to discuss their individual papers in detail. A detailed analysis of a large portion of the work of Harsanyi, Nash and Selten can be found in two excellent surveys, Guth (1994) and Van Damme and Weibull (1995). A more concise treatment of the same material is provided in Van Damme (1995). Although I have benefited greatly from these three papers in writing this essay, my focus here is somewhat different: to identify the impact of the work of Harsanyi, Nash and Selten on microeconomic theory in the last half-century and to point to the parts of their papers that contained the most influential ideas. Thus, I begin by reviewing the development of the role of game theory within modern economic analysis by focusing on the contributions of researchers other than Harsanyi, Nash and Selten. This will provide the proper context for the subsequent discussion the work of the three Nobel Prize winners. The final section contains a few closing comments on evaluating the role of game theory and the style of research championed by Harsanyi, Nash and Selten.

Game Theory and Economic Analysis

The first use of game theory in the analysis of an economic problem is Cournot's classic analysis of duopoly. In this work, Cournot considers quantity

■ *Faruk Gul is Professor of Economics, Princeton University, Princeton, New Jersey.*

competition among two firms. The price corresponding to a given level of aggregate output is determined by the market demand. Cournot computed what we would now call the Nash equilibrium of this game.

John von Neumann and Oscar Morgenstern are justly viewed as the founders of modern game theory. The work of von Neumann and Morgenstern differs from the previous examples of game-theoretic analysis both in terms of the generality of their approach and the ambitiousness of their project. In von Neumann and Morgenstern (1944), they define two-person zero-sum games in normal form and showed the existence of a solution, or equilibrium, in such games.[1] They also define extensive form games, discuss cooperation and coalition formation, and much more. They state that their objective is to "find the mathematically complete principles which define 'rational behavior' for the participants in a social economy and to derive from them the general characteristics of that behavior."

Signs that game theory might fulfill the promise that its founders had foreseen can be found as early as the 1960s. In a path-breaking article on auction theory, Vickrey (1961) describes his purpose as investigating Abba Lerner's claim that a socialist government armed with a sufficiently powerful computer could reproduce market outcomes through a centralized process.[2] What is most striking is that Vickrey's investigation of this issue focuses not on the computational difficulty of finding the right prices or on transaction costs, issues that had been the focus of earlier discussions on the workability of socialism, but on incentive issues. Thus, in the absence of markets, incentive constraints and strategic considerations are viewed as *defining* the set of feasible outcomes from which socially desirable outcomes are to be chosen.

In another paper written about the same time, Vickrey (1960) addressed the problem of misrepresentation of preferences in a social choice setting. Arrow's (1963) celebrated theorem establishes the impossibility of having a social choice function that satisfies Pareto efficiency, an axiom concerning the independence of irrelevant alternatives and is not dictatorial. The desirability of Pareto efficiency and not having a dictator are obvious. But in spite of its intuitiveness, the normative appeal of the axiom of the independence of irrelevant alternatives, which states that the social ranking of two alternatives should depend only on individuals' ranking of those two particular alternatives, is less clear. However, Vickrey pointed out that a social welfare function not satisfying this property might be susceptible to strategic misrepresentation of preferences. Hence, Vickrey helped bring strategic

[1] A normal form or strategic form game refers to a static description of the strategic interaction. In a normal form game, the players are to choose their actions simultaneously. A single static strategy may in principle correspond to a complicated contingent plan. Therefore, a dynamic strategic interaction can also be modeled as a normal form game. However, in a normal form game, the dynamic nature of the interaction is not made explicit. By contrast, an extensive form game is a representation of the strategic interaction that renders the timing of moves and the information of the players when they move explicit.

[2] Vickrey was to win a Nobel Prize of his own in 1996, awarded jointly with James Mirlees.

considerations to the forefront in the analysis of allocative efficiency and in the abstract analysis of social choice problems.

In 1972, Hurwicz provided a detailed example showing that the competitive mechanism might itself be susceptible to strategic manipulation. By the mid-1970s, researchers felt the need to augment their models of competitive equilibrium with game-theoretic models of competition among producers. It was thought that price-taking behavior was unwarranted when the number of producers was small. By the end of the decade, it was felt that competitive equilibria in general could benefit from noncooperative foundations. This development was significant: the idea of seeking justifications for the most basic model of economic analysis, perfect competition, would have been considered absurd a decade earlier. Recall that even Vickrey had used game theory to study what might happen when markets could not operate, not to justify market equilibrium. Thus, game theory became far more central to economic analysis in the 1970s.

Another significant development of the early 1970s was the emergence of information economics. After Akerlof's (1970) path-breaking paper on the used car market, a number of prominent papers modeled asymmetric information between the buyers and sellers of labor (Spence, 1976) or insurance (Rothschild and Stiglitz, 1976). The key common characteristics of the latter papers were that they made novel and surprising observations regarding markets and market failure and that they employed a new formalism. This formalism was neither a model of perfect competition nor a game. The timing of actions or choices was often not clear. It was not obvious who knew what at the time the choices were made and who acted first and who reacted. The difficult trick in these models was to come up with a notion of equilibrium that captured some of the unmodeled dynamics. Market clearing, zero profit and best response conditions were used simultaneously in defining equilibrium. It was no coincidence that there often seemed to be some difficulty with defining and establishing the existence of equilibrium in these models.

Prior to the 1970s, the theory of industrial organization was dominated by mostly descriptive analysis loosely built around the ideal of perfect competition. In this setting, it was difficult to address many of the main issues of industrial organization. Economists had developed a rich theory of perfect competition and the case of monopoly was straightforward enough. However, the intermediate cases appeared troublesome. In the absence of a theory of strategic behavior, it was difficult to construct a useful model of competition among a small number of agents. There were attempts to reshuffle elements of perfect competition and monopoly to create a non-game-theoretic theory for the intermediate case, like the theory of monopolistic competition or the idea that oligopolies faced kinked demand curves. In retrospect, these models appear mildly amusing and perhaps a little embarrassing, rather than genuinely enlightening.

There were other more subtle difficulties in attempting to study industrial organization without the benefit of game theory. For example, without a dynamic theory of strategic interaction, collusion seemed impossible or irrational or at least not explainable within the realm of standard economics. Moreover, in the absence of a model of information, certain issues like the analysis of research and development or the role of patent laws and advertising were impossible to understand. The well-documented cases of perversely low prices in secondary markets appeared as anomalies. This is not to say that none of these issues could possibly be discussed fruitfully without the aid of game theory. However, game theory provided a unified and effective tool for addressing all of these problems.

Perhaps the most important development of the 1970s was the work on "supergames"—that is, infinitely repeated games—and noncooperative collusion. Repeated interaction provides agents with the opportunity to make credible threats or promises regarding future behavior. The availability of such threats and promises enables higher (and lower) payoffs to be sustained in a noncooperative equilibrium than would be possible in a one-shot interaction. The first papers to establish this point were Friedman (1971) and Aumann and Shapley (1976). (Both of these papers used Selten's notion of a subgame perfect Nash equilibrium, to be discussed in the next section.) The opportunity for rationalizing behavior that is inconsistent with short-term incentives as a part of a dynamic equilibrium, and the resulting theory of noncooperative collusion is perhaps the most important single achievement of game theory.

It was the work of Harsanyi and Selten that provided the tools needed for the modern theory of industrial organization. However, a critical step in the incorporation of Selten's ideas on dynamic strategic interaction into the theory of industrial organization was the idea of sequential equilibrium due to Kreps and Wilson (1982). Sequential equilibrium was closely related to Selten's notion of trembling-hand perfect equilibrium (again, to be discussed in the next section). Unlike Selten's notion of equilibrium, however, sequential equilibrium proved to be often easy to compute and easily adapted to games in which either the choice of actions, such as price or quantity, or some other variable was a continuous variable.[3] These turned out to be the type of games that were most useful for the theory of industrial organization.

In the 1980s, armed with models of asymmetric information, dynamic strategic interaction and insights from the information economics of the 1970s, economists developed truly game-theoretic models of industrial organization.

Meanwhile, Rubinstein (1982) investigated a problem that had been, for the most part, thought to be outside the domain of economic analysis: suppose two agents need to agree on a division of a unit of surplus. On what division will they agree, if any? The traditional response to this question was that bilateral monopoly was the quintessential example of the absence of market forces, and hence economics had nothing to say on the matter, except perhaps that an efficient outcome would be reached. However, Rubinstein modeled this interaction as a dynamic

[3] Sequential equilibrium has the additional advantage of emphasizing beliefs rather than trembles. The meaning and interpretation of such small probability mistakes were less clear to most economists, especially prior to the work of Kreps and Wilson.

game with impatient players and concluded that there would exist, in many cases, a unique equilibrium. In this equilibrium, agreement is reached immediately, and the pie is shared roughly in inverse proportion to the level of the agents' impatience. This was a remarkable development. Game theory had yielded a determinate, reasonable answer to one of the most intractable problems of economic analysis. Once an answer to the problem of bilateral monopoly was found, then it seemed plausible that answers to many other problems could be built up from that.

Spurred by the success of the new theory of industrial organization, game theory became a popular tool in many other fields of economics as well; finance, public finance, international trade, labor economics, political science and macroeconomics experienced an infusion of game-theoretic ideas. Of course, descriptive, competitive and other forms of analysis that have nothing to do with game theory continue to play an important role in these disciplines, but more and more, the language of game theory has become the language of economics. In his article on Paul Krugman, Dixit (1993) quotes Krugman's recommendation: "The point . . . is to wear one's theory lightly." Since the 1980s, chances are that the lightly worn theory has in it a good deal of game theory.

However, perhaps the biggest impact of game theory has been on what one might call pure theory; that branch of economics that focuses on developing tools and investigating the foundations of economic analysis. The extensive use of game-theoretic models in applied theory led to a great deal of research interest in refinement and critique of these tools, as well as a search for alternatives. More and more, "pure" theorist came to mean game theorist, even though other kinds of theory continued to prosper as well. Relatedly, game theory has become the center of a particular approach to economics; one that values precision and generality of the underlying ideas and emphasizes the usefulness of abstraction and formalism. To put it another way, game theory has become the meeting place for those who do not wish to wear their theory so lightly. This is a role that game theory has, at least partially, taken over from general equilibrium analysis since the 1970s.

The Contributions of Nash, Harsanyi and Selten

The three most important ideas in noncooperative game theory are equilibrium, asymmetric information and credibility. These also happen to be the contributions of John F. Nash, John C. Harsanyi and Reinhard Selten, respectively.

In his paper "Non-Cooperative Games," John Nash (1950b) defined the concept of an n-person (finite) game and an equilibrium point of such a game. Nash summarized his contribution modestly in the introduction as follows: "The notion of an equilibrium point is the key ingredient in our theory. This notion yields a generalization of the concept of the solution of a two-person zero-sum game. It turns out that the set of equilibrium points of a two-person zero-sum game is simply the set of all opposing 'good strategies.' " In other words, an equilibrium point is a profile of strategies; that is, a strategy for each player, in which each agents' strategy is a best response to the strategies of the others. These equilibrium points are now called Nash equilibria.

Nash proved that all finite n-person games have at least one equilibrium. He generalized the theory of two person zero-sum games to allow for multiple players and arbitrary payoffs. Before his work, much emphasis was placed on the fact that equilibrium, when attention is restricted to two-person zero-sum games, provides safe payoffs to agents. The more general setting enabled Nash to isolate the idea of "opposing good strategies," or a best response. Hence, the modern theory of rational strategic interaction, one that is capable of dealing with economic problems, not just "games," was born. To this day, the notion of a Nash equilibrium continues to be the main vehicle for determining the social consequence of conflicting individual interests. Finite n-person games are still the main tools for summarizing many of the essential insights of game theory such as the conflict between individual incentives and efficiency, the problem of coordination, the importance of reasoning about opponents' reasoning and so on. The theory of finite normal form games, exactly as formulated by Nash, is also the initial testing ground for most new game-theoretic ideas.

In two related papers, "The Bargaining Problem" and "Two-Person Cooperative Games," Nash (1950c, 1953) defined the two-person bargaining problem, proposed his axiomatic solution and provided a cooperative theory for two-person non-zero-sum games that leads to his bargaining solution. The first of these papers has had an enormous impact. The axiomatic approach initiated by this paper continues to be an active area of research. The particular solution proposed by Nash is still the main benchmark for this subfield and is a frequently used tool in more applied research. In terms of noncooperative game theory, however, the more significant contribution of these two papers is the initiation of what is now called the "Nash program."

The meaning of this term can be understood by studying the two papers together: The bargaining problem is a description of the utilities that can be achieved. Nash in the first paper considers an abstract rule that picks a solution from any possible bargaining problem. His simple and reasonable axioms yield a unique solution now called the Nash bargaining solution. The second paper takes as its starting point a normal form game, views the set of possible outcomes of this game as the attainable payoffs of a bargaining problem, describes a new noncooperative game to model the negotiation procedure, computes a Nash equilibrium of the negotiation game and observes that this Nash equilibrium results in the payoffs associated with the bargaining solution. Thus, the axiomatic solution is justified by showing that it is the Nash equilibrium outcome of some noncooperative game. This is the Nash program.

Stated in its most extreme form, the view underlying the Nash program is that all nonstrategic (positive or normative) theories, such as the theory of perfect competition, are in fact goals or guesses as to what might happen.[4] The feasibility of

[4] This is a much more comprehensive definition of the Nash program than is customary. Usually, the Nash program refers only to research aiming to establish noncooperative foundations of cooperative

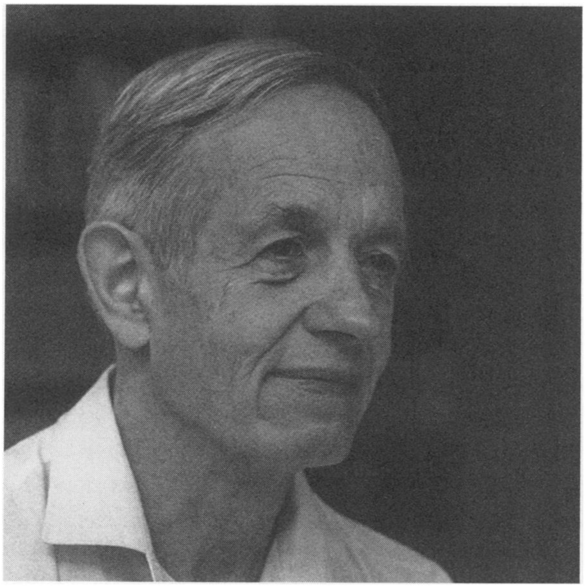

John Nash

these goals or the validity of the guesses is to be established by showing them to be the Nash equilibrium outcomes of a sensible noncooperative model of strategic interaction—that is, of a noncooperative game. Vickrey's work discussed in the previous section can be considered to be within the Nash program. The work on auctions aims to determine if a socialist government can achieve market outcomes, while the discussion of Arrow's impossibility theorem suggests that social choice rules that fail the independence of irrelevant alternatives condition are not feasible. The work of the 1970s and 1980s on decentralization and noncooperative foundations of competitive equilibria are readily seen to be examples of the Nash program with this extended definition. Here we are verifying that competitive equilibria are reasonable outcomes of economic situations by noting that they can be supported as the Nash equilibrium of certain games. Rubinstein's (1982) classic paper is easily seen to fall within this program as well.

By the 1960s, it seemed clear to many if not most economists and game theorists that much of economic theory was based on an implausible and perhaps misleading assumption of complete or symmetric information. Three seminal papers by Akerlof (1970), Spence (1976) and Rothschild and Stiglitz (1976) were among the earliest attempts to deal with the problem of asymmetric information. By the time these papers were written, most of the relevant game-theoretic tools had already been developed by Harsanyi and Selten. However, these tools were not yet sufficiently understood or refined and consequently, they had not made their way into the tool chest of most economists.

---

solution concepts. For the purpose of identifying the impact of Nash's work on game theory and economics, this more comprehensive definition may be more suitable.

In 1967, John Harsanyi published the first in a sequence of three papers entitled "Games of Incomplete Information Played by 'Bayesian' Players." His work would soon provide economists with the much needed tools and vocabulary for studying the problems of asymmetric information. Prior to the work of Harsanyi, economists were forced to work either with models in which there was absolutely no uncertainty or there were objective probability distributions describing the commonly shared view of the underlying uncertainty. Hence, until Harsanyi's work made its way into the mainstream, it was difficult and rare for researchers to study situations in which one person knew something that others did not. It was not that economists found situations in which one agent was more informed than another to be unimaginable; the difficulty was coming up with a coherent modeling device for such situations. It was clear that postulating that 1 knows X and 2 does not was not enough in deciding what 1 or 2 would do in a strategic setting. It was necessary to specify what 2 knew about what 1 knows, and what 2 knew about what 1 knows about what 2 knows, and so on. Harsanyi provided a way of specifying all these layers of information in a consistent manner.

To illustrate the importance of Harsanyi's contribution, imagine two gamblers contemplating betting on the outcome of a role of a die. They can bet either "High," which wins if the outcome is 4 or higher, or they can bet "Low," which wins if the outcome is 3 or less. Now suppose that the gamblers are in one room while the die is to be rolled in another. We agree that after the die is tossed but before the gamblers see the outcome or have to place their bets, we will let the first gambler know if the outcome is odd or even and the second gambler know if the outcome is a prime number (2, 3 or 5) or not.

Given this structure of the interaction, suppose that the outcome is 2. Then, gambler one is told that the outcome is even while gambler two learns that the outcome is prime. What else can gambler one deduce at the time she is informed that the outcome is even? If she initially assigned equal probabilities to each of the six outcomes, presumably she now believes that there is a 1/3 chance of 2, 1/3 chance of 4 and 1/3 chance of 6. Hence, there is a 2/3 chance that the outcome is "High" and 1/3 chance that it is "Low." More importantly, if gambler one has understood the rules of this interaction she can conclude also that there is a 1/3 chance that gambler two is told that the outcome is prime and 2/3 chance that the second gambler is told the outcome is not prime. This means that for gambler one, there is a 1/3 chance that gambler two assigns a 1/3 chance to the outcome "High" and 2/3 chance that gambler two assigns a 2/3 chance to the outcome "High."

The exact values of these probabilities are not important. What is important is that at the interim stage of this interaction—that is, at the stage at which the gamblers are given their private information but have not yet learned the outcome of the roll—we are in exactly the kind of situation that is of interest to economists and game theorists; the agents have different information regarding the outcome

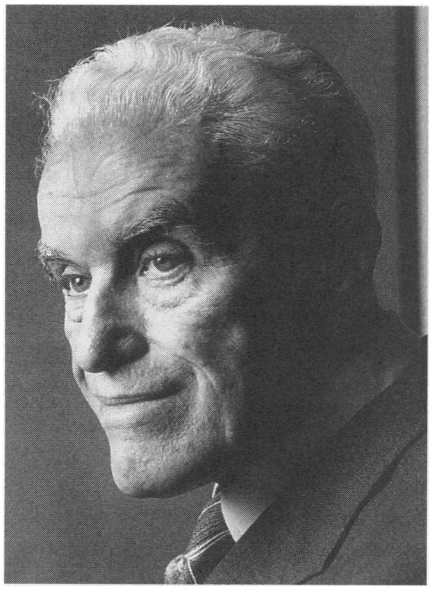

John C. Harsanyi

and are in a position to reason about the opponent's information and what the opponent is thinking about the other's information. For example, if gambler one were to learn that gambler two is betting on "Low," which presumably is what gambler two would do given his information, then she, gambler one, would bet "Low" as well, even though her own information favors betting "High." A decision to bet "Low" on the part of gambler two would signal to gambler one that the outcome is prime. Together with her own information that the outcome is even, this would mean that the outcome is 2. Hence gambler one would bet "Low."

One may be tempted to argue that while possibly interesting, the story of differential information here is a rather exceptional one. In most economic situations, we do not have such well-specified outcomes and clear and common understanding of the "rules of the game." Economic agents are typically not sure about what information their opponents have or even what kind of information they might have. Moreover, they are not simply interested in things like the probabilities of various outcomes but also the way they and others value these outcomes, the options available to their opponents and the choices their opponents and themselves will have in the future. What is so striking about Harsanyi's contribution and what makes it so significant is that he successfully argued that all situations of asymmetric information, whether the relevant uncertainty is about the set of possible outcomes, or the agents' preferences over these outcomes, or the set of choices available to the agents, or agents' uncertainty about the information of their opponents, can be described as the interim stage of some suitably constructed model of asymmetric information, essentially no different than the problem confronting our gamblers above.[5] In other words, Harsanyi developed a general model of strategic interaction that incorporated asymmetric information. He called this model a game of incomplete information. He showed that Nash's concept of an equilibrium point could be extended in a straightforward manner to games of incomplete information.

It is worth noting that the story of differential information that our gamblers are confronted with above was not unknown to statisticians before Harsanyi's work; Vickrey had even used such a model to study auctions. What Harsanyi showed was that the model was perfectly general; all problems of asymmetric information could be modeled as the interim stage of such a model. He also showed a way to combine the model with the notion of a game. Thus economists not only gained a formalism to discuss asymmetric information, but also a way to incorporate that formalism with their theories of strategic interaction.

Harsanyi provided a detailed discussion and interpretation of each of the constructs in his games of incomplete information and described the relation between these constructs and a particular view of probability. One of the many important results of these discussions was the concept of a "type" that summarizes all of the relevant characteristics of a particular agent. This notion of a type is now a standard and almost universal concept in economics.

The information economists of the 1970s who did their work after the work of Harsanyi, but without knowledge of it, had difficulties with the very issues that Harsanyi had resolved. Spence (1976), in his path-breaking analysis of job market signaling, struggles with the problem of finding an interpretation for the probability distributions that appear in his analysis. He suggests that these probabilities might be rest points of some adjustment process of population frequencies and that a fixed point argument might be used to establish the existence of such a rest point. Had he been familiar with the work of Harsanyi, Spence would have observed that the interaction he had been investigating could have been modeled as a game, and that a frequentist interpretation for the probabilities was possible but not necessary. Moreover, he would have realized that the fixed point argument that he was searching for was none other than the celebrated existence theorem of Nash, as utilized by Harsanyi. Similarly, Rothschild and Stiglitz (1976) in their paper are troubled by the possibility of numerous different notions of equilibria that in their view represented various different levels of rationality on the part of the agents. Rothschild and Stiglitz also suggest the choice of equilibrium might depend on how quickly agents can respond to deviations of others. Had the work of Harsanyi been internalized in 1976, Rothschild and Stiglitz would have known that the "right" equilibrium notion for their context was Bayesian-Nash equilibrium—that is, Harsanyi's notion.

The potential multiplicity of equilibrium outcomes that Rothschild and Stiglitz (1976) attributed to different possible adjustment lags was in fact due to different

[5] Harsanyi's argument regarding the generality of his model of asymmetric information was later formalized by Mertens and Zamir (1985) and Brandenburger and Dekel (1993).

possible extensive form game representations of roughly the same economic interaction. What Rothschild and Stiglitz were observing was that rationality in the dynamic interaction entailed anticipating the future rational behavior of others. What was not yet clear was that given rationality—that is, the ability to anticipate opponents' behavior—the appropriate question was who could commit and who had to respond rather than how quickly could agents readjust their strategies. For understanding this last issue, Reinhard Selten's contribution to game theory was needed.

The notion of an extensive form game existed before Selten's first major contribution to this field. However, it was taken for granted that all the relevant aspects of dynamic interaction are captured by the corresponding normal form game. The detailed, dynamic extensive form description was viewed as a strategically irrelevant notational device. Thus, when confronted with an extensive form game, one needed to find the suitable normal form representation, which was the type of representation studied in Nash's work, and compute the corresponding Nash equilibria. Selten's first major contribution, the notion of subgame perfection, disproved this claim.

Consider the strategic problem confronting two brothers. The older brother has grabbed a toy belonging to the younger. The parents have this policy: if either child complains, the children are punished by having all their toys taken away and are sent to their room. This is the least desirable outcome for either boy. However, conditional on not being punished, each child would like possession of the toy that is currently in the hand of the older one. The younger one says to the older that unless he gets back the toy, he will complain to the parents. The older brother reasons as follows: "If I do give back the toy, I will be fine but will not get to play with it today. If I do not give the toy back then my brother gets to decide whether to complain or not. If he does, then I am in trouble. If he does not, I end up with the toy, which is my favorite outcome. But my brother cannot gain by complaining to our parents, so I do not think he will, and therefore I am not going to give back the toy." The conclusion is that the older brother considers the threat of the younger sibling not credible and hence dismisses it.

The simplicity of the above story masks the generality of the underlying idea. Note that the correct solution of the problem entails the older brother realizing that if his sibling were faced with the choice of complaining to the parents, he would not do so since it is not optimal given his situation. Thus, Selten formulated the idea that strategies must not only be best responses given the anticipated behavior of the opponents, but in all contingencies. The candidate equilibrium in which the older brother turns over the toy to the younger, fearing that he will complain, fails Selten's criterion since the predicted behavior of the younger sibling is not optimal in the event that the older sibling fails to give him the toy. Note that given the candidate equilibrium, the older brother is not expected to deny his sibling the toy. Nevertheless, Selten's notion of equilibrium insists that behavior be optimal, even in such unanticipated contingencies or *subgames.*

Selten applied his notion of subgame perfection to the analysis of the theory of entry deterrence. Prior to 1965, the analysis of entry deterrence was based on the notion that incumbent firms artificially kept prices low to keep out competitors, without questioning why the potential entrants would expect the prices to stay low if and when they entered. It was not at all clear why the potential entrant should make any inference regarding what would happen after entry based on pre-entry behavior. Selten's analysis and his use of the concept of subgame perfection exposed the weakness of the existing views on entry deterrence.

Another example demonstrating the importance of dynamics in strategic interaction is the repeated prisoner's dilemma. The standard analysis of the one-shot prisoner's dilemma yields the obvious conclusion that neither agent will cooperate. However, if this game is repeated infinitely many times and if both agents place a sufficiently high value on the future compared to the current period, then it is possible to sustain cooperation in every period with the following strategies: if anyone ever defects then everyone will defect forever; both players will cooperate until somebody defects. Hence, the outcome of these strategies is that cooperation will continue forever. Notice that both the threat to defect forever after one player has defected and the promise to cooperate forever if no one deviates are credible (that is, a part of a subgame perfect equilibrium): if one player were defecting in every period, it is optimal for the second player to defect in every period as well. Similarly, given that one player is threatening to defect forever after a defection but promising to cooperate as long as no one defects, it is optimal for a player who values the future sufficiently to cooperate in every period. This way of sustaining apparently collusive outcomes—that is, outcomes inconsistent with short-term incentives—with credible threats of retaliation is called noncooperative collusion. Some of the most significant accomplishments of game theory in the 1980s entailed establishing the limits of noncooperative collusion for a much wider class of dynamic games.

The idea of an infinitely repeated game, or "supergame," is not due to Selten (Aumann, 1960). Selten's contribution is defining subgame perfection and viewing subgame perfect Nash equilibria as the only relevant equilibria of all dynamic games. Requiring the strategies to be a Nash equilibrium after every conceivable history (that is, subgame perfect) is the key to understanding the idea of a credible threat (or promise) as well as to understanding the significant effect of dynamic interaction on incentives.

Selten (1978) made a second very significant contribution along these lines. He noticed that his earlier notion of subgame perfection failed to capture the underlying idea in some extensive form games. In his 1975 paper, "Reexamination of the Perfectness Concept for Equilibrium Points," he defined a new notion of perfection, called trembling-hand perfection, that captured the basic idea of credibility in all extensive form games (Selten, 1975a). The main idea was roughly the following: to eliminate those equilibria that entail threats or promises that are not credible, Selten focused on equilibria that were the limits of equilibria when players were prone to making mistakes with very small probabilities.

In the problem facing the two brothers, trembling-hand perfection would require the younger brother to make a best response not just to the older brother's equilibrium strategy, but also to a strategy that has the older make a mistake with

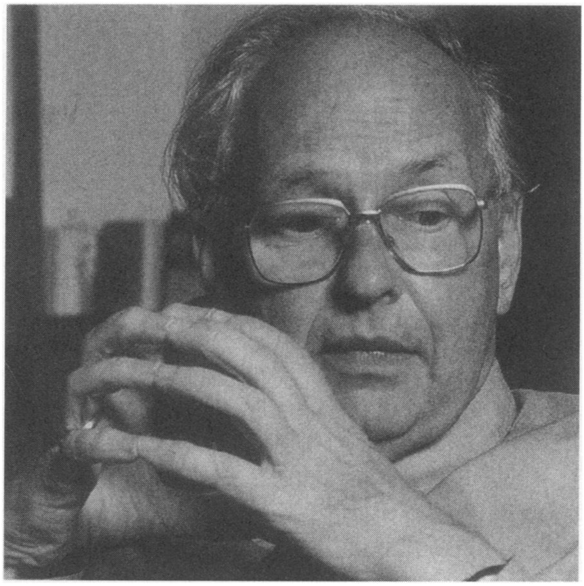

Reinhardt Selten

some small probability. Again, consider the candidate equilibrium in which the older brother is to turn the toy over for fear that the younger brother will complain, and the younger brother is to complain if the older brother does not give him the toy. The younger brother's strategy is optimal only because he does not expect the older brother to hold out. However, if the older brother were to make a mistake and with some small probability, refuse to yield, then the younger brother's strategy would no longer be optimal, even if the probability of this mistake is very small. Thus, in this example, trembling-hand perfection also rules out the equilibrium with the noncredible threat. It turns out that any equilibrium that is not subgame perfect will also fail the test of trembling-hand perfection. However, trembling-hand perfection will rule out certain equilibria that are not excluded by subgame perfection.

In introducing the new concept, Selten also introduced three ideas that would dominate game-theoretic discussion for the next two decades: that small probability mistakes could capture forward-looking rationality (that is, credibility and subgame perfection) in dynamic interaction; that such strategy perturbations could be used to eliminate equilibria; and finally that a sensible definition of rationality must pay some attention to possible deviations from rationality by other players.

The work of Harsanyi and Selten had given to economists all the tools needed for studying industrial organization: that is, a way of dealing with asymmetric information and dynamics in strategic situations. Positive microeconomic theory of the 1980s can more or less be summarized as the applications of these tools to various concrete economic questions. To address any question in a number of fields, one had simply to write a suitable extensive form game of asymmetric information describing the interaction and then apply a refinement criterion such as subgame perfection or trembling-hand perfection. Much of economic theory consisted of investigating the limits on efficiency imposed by the need for credibility in dynamic strategic interaction (that is, the limits of noncooperative collusion) and the constraints imposed by asymmetric information (that is, incentive constraints).

Conclusion

In recent years, discussions on the achievements of game theory have often emphasized the practical significance of game theory. Thus, it is customary to start off such discussions by paying tribute to the success of game theory in the design of the auction of the radio spectrum by the Federal Communications Commission (FCC). Similarly, one recent book on game theory states, "By the end of this book we hope that you will emerge a more effective manager, negotiator, athlete, politician, or parent." I suppose that all this is a good sign, showing that game theory and its practitioners have arrived and that the theory is now less likely to be viewed as esoteric or irrelevant.

However, the exercise of identifying the contribution of a field of research with its perceived practical consequences can be problematic. This juxtaposition makes it impossible to distinguish between influence and benefit, or between the success of game theorists and the achievements of game theory, and can cause short-term success to be overvalued at the expense of long-term contributions. Unlike academic researchers, those who have the responsibility of making a particular policy decision or carrying out a given task are often under severe time constraints and hence their choice of theoretical tools or advisers is at best a limited signal of the worth of the given theory. Furthermore, evaluating the success of a particular theory is in itself a difficult and fundamental scientific question for which the skills (and interests) of practitioners may be ill suited.

Consider the following example. In describing one FCC auction, one newspaper observed: "By combining cutting-edge economic theory with auction procedures, the FCC wrung top prices out of bidders." It is not clear how the author of these lines concluded that prices were higher in the FCC auction than they would have been in the absence of the "cutting-edge" theory. Nor is it easy to figure out what the cutting-edge theory was that went into the design of the auction. My purpose here is not to deny that game theory has had such "practical" impact, nor to dismiss criticism regarding what some might consider the limited extent of this impact. I wish only to justify my decision to focus exclusively on the influence of the work of the three Nobel Prize winners on research in microeconomic analysis. I feel strongly that the proper way to celebrate the work of these three researchers is to concentrate on resulting changes in the way economists reason.

The work of Harsanyi, Nash and Selten has had an major impact on nearly all fields of economic analysis. However, their greatest effect has been through their role as leading practitioners of a formal, rigorous and abstract approach. This approach has yielded benefits not only within game theory but also in other areas. For example, the techniques developed by Nash, Harsanyi and Selten to study strategic equilibria have been used to study competitive equilibria. The modeling of asymmetric information pioneered by Harsanyi was later used and developed to study problems of knowledge and rationality. Directly or indirectly, the work of these three outstanding researchers has had an effect on fields ranging from biology to philosophy.

That a formal and abstract approach emphasizing the generality of the underlying ideas rather than immediate answers should yield long-lasting benefits is no surprise to many economists. To convince others of the virtues of this approach, one can do no better than to point to the remarkable achievements of Harsanyi, Nash and Selten.

References

Akerlof, George A., "The Market for 'Lemons': Qualitative Uncertainty and the Market Mechanism," *Quarterly Journal of Economics,* August 1970, *84,* 488–500.

Arrow, Kenneth J., *Social Choice and Individual Values.* 2nd ed., New York: John Wiley and Sons, 1963.

Aumann, Robert, "Acceptable Points in Games of Perfect Information," *Pacific Journal of Mathematics,* 1960, *10,* 381–417.

Aumann, Robert, and Lloyd S. Shapley, "Long Term Competition: A Game Theoretic Analysis," mimeo, 1976.

Brandenburger, Adam, and Eddie Dekel, "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory,* February 1993, *59,* 189–98.

Cournot, Antoine Augustin, *Researches into the Mathematical Principles of Wealth.* N. T. Bacon, translator, Macmillan, New York: Macmillan, 1897.

Dixit, Avinash, "In Honor of Paul Krugman: Winner of the John Bates Clark Medal," *Journal of Economic Perspectives,* Spring 1993, *7*:2, 173–88.

Dixit, Avinash, and Barry Nalebuff, *Thinking Strategically: The Competitive Edge in Business, Politics, and Everyday Life.* New York and London: Norton, 1991.

Friedman, James, "A Non-Cooperative Equilibrium for Supergames," *Review of Economic Studies,* January 1971, *28,* 1–12.

Guth, Werner, "On the Scientific Work of John C. Harsanyi, John F. Nash, and Reinhard Selten." Discussion Paper No. 35, Humboldt-University of Berlin, 1994.

Harsanyi, John, "Games with Incomplete Information Played by 'Bayesian Players': I. The Basic Model," *Management Science,* November 1967, *14,* 159–82.

Harsanyi, John, "Games with Incomplete Information Played by 'Bayesian Players': II. Bayesian Equilibrium Points," *Management Science,* January 1968a, *14,* 320–34.

Harsanyi, John, "Games with Incomplete Information Played by 'Bayesian Players': III. The Basic Probability Distribution of the Game," *Management Science,* March 1968b, *14,* 486–502.

Kreps, David M., and Robert Wilson, "Sequential Equilibria," *Econometrica,* July 1982, *50,* 863–94.

Mertens, Jean-Francois, and Shmuel Zamir, "Formulations of Bayesian Analysis for Games of Incomplete Information," *International Journal of Game Theory,* 1985, *14*:1, 1–29.

Nash, John F., "Equilibrium Points in N-Person Games," *Proceedings from the National Academy of Sciences,* 1950a, *36,* 48–49.

Nash, John F., "Non-Cooperative Games," Ph.D. dissertation, Princeton University, 1950b.

Nash, John F., "The Bargaining Problem," *Econometrica,* April 1950c, *18,* 155–62.

Nash, John F., "Non-Cooperative Games," *Annals of Mathematics,* 1951, *54*:2, 286–95.

Nash, John F., "Two-Person Cooperative Games," *Econometrica,* January 1953, *21,* 128–40.

Neumann, John von, and Oscar Morgenstern, *Theory of Games and Economic Behavior.* Princeton, N.J.: Princeton University Press, 1944.

Rothschild, Michael, and Joseph E. Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics,* November 1976, *90,* 629–49.

Rubinstein, Ariel, "Perfect Equilibrium in a Bargaining Model," *Econometrica,* January 1982, *50,* 97–109.

Selten, Reinhard, "Spieltheoretische Behandlung eines Oligopolmodels mit Nachfragetragheit," *Zeitschrift fuer die Gesamte Staatswirtschaft,* 1965, *121,* 301–24, 667–89.

Selten, Reinhard, "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory,* 1975a, *4*:1-2, 25–55.

Selten, Reinhard, "The Chain Store Paradox," *Theory and Decision,* 1975b, *9*:2, 127–59.

Spence, Michael, "Informational Aspects of Market Structure: An Introduction," *Quarterly Journal of Economics,* November 1976, *90,* 591–98.

Van Damme, Eric, "On the Contributions of John C. Harsanyi, John F. Nash and Reinhard Selten," *International Journal of Game Theory,* 1995, *24*:1, 3–12.

Van Damme, Eric, and Jorgen Weibull, "Equilibrium in Strategic Interaction: The Contributions of John C. Harsanyi, John F. Nash and Reinhard Selten," *Scandinavian Journal of Economics,* March 1995, *97,* 15–40.

Vickrey, William, "Utility, Strategy and Social Decision Rules," *Quarterly Journal of Economics,* November 1960, *74,* 507–35.

Vickrey, William, "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance,* March 1961, *16,* 8–37.

LINKED CITATIONS

A Nobel Prize for Game Theorists: The Contributions of Harsanyi, Nash and Selten
Faruk Gul
The Journal of Economic Perspectives, Vol. 11, No. 3. (Summer, 1997), pp. 159-174.
Stable URL:
http://links.jstor.org/sici?sici=0895-3309%28199722%2911%3A3%3C159%3AANPFGT%3E2.0.CO%3B2-H

References

The Market for "Lemons": Quality Uncertainty and the Market Mechanism
George A. Akerlof
The Quarterly Journal of Economics, Vol. 84, No. 3. (Aug., 1970), pp. 488-500.
Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28197008%2984%3A3%3C488%3ATMF%22QU%3E2.0.CO%3B2-6

In Honor of Paul Krugman: Winner of the John Bates Clark Medal
Avinash Dixit
The Journal of Economic Perspectives, Vol. 7, No. 2. (Spring, 1993), pp. 173-188.
Stable URL:
http://links.jstor.org/sici?sici=0895-3309%28199321%297%3A2%3C173%3AIHOPKW%3E2.0.CO%3B2-T

A Non-cooperative Equilibrium for Supergames
James W. Friedman
The Review of Economic Studies, Vol. 38, No. 1. (Jan., 1971), pp. 1-12.
Stable URL:
http://links.jstor.org/sici?sici=0034-6527%28197101%2938%3A1%3C1%3AANEFS%3E2.0.CO%3B2-2

Games with Incomplete Information Played by "Bayesian" Players, I-III. Part I. The Basic Model
John C. Harsanyi
Management Science, Vol. 14, No. 3, Theory Series. (Nov., 1967), pp. 159-182.
Stable URL:
http://links.jstor.org/sici?sici=0025-1909%28196711%2914%3A3%3C159%3AGWIIPB%3E2.0.CO%3B2-P

LINKED CITATIONS

Games with Incomplete Information Played by "Bayesian" Players, I-III. Part II. Bayesian Equilibrium Points
John C. Harsanyi
Management Science, Vol. 14, No. 5, Theory Series. (Jan., 1968), pp. 320-334.
Stable URL:
http://links.jstor.org/sici?sici=0025-1909%28196801%2914%3A5%3C320%3AGWIIPB%3E2.0.CO%3B2-0

Games with Incomplete Information Played by "Bayesian" Players, I-III. Part III. The Basic Probability Distribution of the Game
John C. Harsanyi
Management Science, Vol. 14, No. 7, Theory Series. (Mar., 1968), pp. 486-502.
Stable URL:
http://links.jstor.org/sici?sici=0025-1909%28196803%2914%3A7%3C486%3AGWIIPB%3E2.0.CO%3B2-4

Sequential Equilibria
David M. Kreps; Robert Wilson
Econometrica, Vol. 50, No. 4. (Jul., 1982), pp. 863-894.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28198207%2950%3A4%3C863%3ASE%3E2.0.CO%3B2-4

The Bargaining Problem
John F. Nash, Jr.
Econometrica, Vol. 18, No. 2. (Apr., 1950), pp. 155-162.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28195004%2918%3A2%3C155%3ATBP%3E2.0.CO%3B2-H

Non-Cooperative Games
John Nash
The Annals of Mathematics, 2nd Ser., Vol. 54, No. 2. (Sep., 1951), pp. 286-295.
Stable URL:
http://links.jstor.org/sici?sici=0003-486X%28195109%292%3A54%3A2%3C286%3ANG%3E2.0.CO%3B2-G

Two-Person Cooperative Games
John Nash
Econometrica, Vol. 21, No. 1. (Jan., 1953), pp. 128-140.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28195301%2921%3A1%3C128%3ATCG%3E2.0.CO%3B2-N

LINKED CITATIONS

Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information
Michael Rothschild; Joseph Stiglitz
The Quarterly Journal of Economics, Vol. 90, No. 4. (Nov., 1976), pp. 629-649.
Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28197611%2990%3A4%3C629%3AEICIMA%3E2.0.CO%3B2-N

Perfect Equilibrium in a Bargaining Model
Ariel Rubinstein
Econometrica, Vol. 50, No. 1. (Jan., 1982), pp. 97-109.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28198201%2950%3A1%3C97%3APEIABM%3E2.0.CO%3B2-4

Informational Aspects of Market Structure: An Introduction
Michael Spence
The Quarterly Journal of Economics, Vol. 90, No. 4. (Nov., 1976), pp. 591-597.
Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28197611%2990%3A4%3C591%3AIAOMSA%3E2.0.CO%3B2-1

Utility, Strategy, and Social Decision Rules
William Vickrey
The Quarterly Journal of Economics, Vol. 74, No. 4. (Nov., 1960), pp. 507-535.
Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28196011%2974%3A4%3C507%3AUSASDR%3E2.0.CO%3B2-B

Counterspeculation, Auctions, and Competitive Sealed Tenders
William Vickrey
The Journal of Finance, Vol. 16, No. 1. (Mar., 1961), pp. 8-37.
Stable URL:
http://links.jstor.org/sici?sici=0022-1082%28196103%2916%3A1%3C8%3ACAACST%3E2.0.CO%3B2-L