Discussion Paper #1426
September 15, 2006

"Some Personal Views of Game Theory"

Ehud Kalai

www.kellogg.northwestern.edu/research/math

CMS-EMS
The Center for
Mathematical Studies
in Economics &
Management Sciences

Northwestern University
2001 Sheridan Road 580 Leverone Hall Evanston, IL 60208-2014 USA

Some Personal Views of Game Theory

by

Ehud Kalai

Abstract: This is a draft of a chapter for a book called "Game Theory, 5 Questions," to be published by Automatic Pressed / VIP. Below are the five questions and my answers. Any feedback is appreciated.

Question 1. Why were you initially drawn to game theory?

The 1960s were turbulent years on American campuses, as students were rebelling against the government, social institutions and "establishment" values. There was a push to increase awareness by the use of new approaches that go beyond the standard "linear thinking." As a PhD student in the mathematics department at Cornell University, I was affected by the spirit of the time, and by the feeling that work in pure mathematics could be isolating. So I decided to look for a career that would be more "relevant" and "interactive." My advisor, Jack Kiefer, suggested that I take a few courses that were more applied than the standard real analysis, algebra, and topology courses taught in the pure-math program.

Looking through Cornell's catalogue of courses, I marked some statistics and optimization courses but then noticed a course called "Game Theory," taught by William Lucas. The name of the course was puzzling and I decided to sit in on the course and see what it was about. It became "love at first sight."

I always loved the axiomatic approach and the logic of geometry, and here was a whole new application of this approach. Bill Lucas started with cooperative game theory, where simple appealing axioms turn out to be powerful enough to offer solutions to questions about individual behavior, group choices, fairness, and the like. It was surprising and fascinating. Although I cannot remember my feelings when I was introduced to geometry, I was amazed by the fact that you can resolve questions about fairness and behavior by simple mathematical axioms.

In addition to the fascinating process of "mathematizing" behavior, it quickly became clear that game theory was a young subject, and that many new ideas had not been tried yet. So creativity was welcome, with mathematics as its canvas. In addition to the wonderful teaching of Bill, young Louis Billera was around and was teaching many of us how to go about doing mathematical research in this subject. The enthusiasm of Bill and Lou, together with that of some of their visitors (including Guillermo Owen and Lloyd Shapley), was captivating, and indeed, during this period they produced some outstanding game theorists including Robert Bixby, Pradeep Dubey, and Robert Weber.

Game theory was still a small and unpopular area of research at American universities when I completed my PhD in 1972. But it was an active area of research in my home country, Israel. I took a faculty position in the Statistics Department of Tel Aviv University, where I could work with David Schmeidler (whose papers were central to my thesis), and arranged a visit with the leader of the Israeli school of game theory, Robert Aumann. The school of Israeli game theorists proved to be one of the most creative research groups. Joining it was a highly rewarding experience, both professionally and personally.

Another important step in my evolution as a game theorist was the move to the business school of Northwestern University in 1975. Game theory was still far from mainstream economics, but the young economists here, with strong support from the administration, were eager to incorporate it into economics. We became a major player in the "game theory revolution" in economics and management of the 1970s and '80s, and we are currently involved in similar revolutions in political science and operations management.

Question 2. What example(s) from your work (or the work of others) illustrates the use of game theory for foundational studies and/or applications?

At this stage of its development, it is useful to distinguish between theoretical applications and real-life applications of game theory. By "theoretical applications" I mean the contributions that game theory makes to other theories. Real-life applications are ones that address specific concrete problems faced by real decision makers.

Most of the applications of game theory until now have been theoretical. I refer to them as applications, despite the fact that they deal with theories, because what is a theory in one field may be an application in another. For example, economic theorists interested in the potential effect of a proposed government regulation study the Nash equilibrium of the game in which the regulation is imposed, and compare it with that of the game without the imposed regulation. Similarly, political scientists interested in comparing parliamentary systems study the games that result under the various parliamentary systems, and compare the efficiency of their game-theoretic solutions.

Game theory has been quite successful in dealing with theoretical applications like the ones above. There is almost no area of economics that was not drastically affected by game-theoretical tools. Similar revolutions are taking place in political and computer sciences. But even in areas not focused on human behavior, such as evolutionary biology, game theory has had a profound effect. At a more foundational level are the contributions of game theory to mathematics and philosophy. Interactive epistemology and the foundation of rational behavior provide a wonderful example. How do we mathematically model an individual's knowledge, including knowledge about the knowledge of others (which includes their knowledge of others, etc.)? What is rational behavior in situations that involve such complex knowledge structure? How does rational behavior respond to demonstrated violations of rational behavior?

I have made my share of contributions to game theory and to its applications to economics, social choice, operations research and computer science. But given that, unlike most game theorists, I have also had the experience of dealing with real-life applications, it may be useful to focus my discussion here on some of these.

As a business school professor, I teach game theory to business executives. Some of them come back to me later for consulting work, when they encounter strategic problems that require sophisticated analytical thinking. I will describe three examples that illustrate some of the benefits and the difficulties of two different game-theoretic methods, and their mixed record of success.

Two cases where game theory was applied successfully involved Baxter Healthcare, one of the world's largest health-care corporations. The first question involved the pricing strategy for Baxter's blood-plasma separating machines in China. The practice of Baxter and its competitor is to lend the machines, for free, to HMOs that use it, but to charge for the new plastic tubing needed for each individual blood donor. Chinese HMOs were gradually shifting to a strategy of taking the free machines from Baxter's competitor, but producing their own plastic tubing in China for substantially less than the price charged by the competitor. While Baxter itself was not cheated in this manner, market share was rapidly shifting to Baxter's competitor. Moreover, the competitor was not fighting this process of being cheated. It was not clear why the competitor lets the process continue, and how Baxter should respond.

Together with a team of experts (in finance, marketing, Chinese HMOs, and more) we developed a large strategy tree[1] that considered the major options available to Baxter and the other players (the competitor, Chinese HMOs, the Chinese government) over the next two years. The tree led to the conclusion that the best course of action was for Baxter to partner with local Chinese producers, and to produce their own tubing in China at a highly competitive price, even when compared with the local production by the Chinese HMOs. While the approach required making

[1] I use the term "strategy tree" to denote an item between a decision tree and an extensive game. It is less sophisticated than an extensive game, because it does not include information sets. It is more sophisticated than a decision tree, since the major choices of opponents are explicitly modeled with some attention paid to the optimality of their choices.

ad-hoc assumptions and simplifications of the problem, it had definite advantages. In addition to leading to a final solution to the concrete problem, it turned out to be a tool that facilitates communication within a large group of experts with diversified expertise.

Different experts contributed their input to their respective parts of the tree. For example, what Chinese HMOs might do in response to various actions of Baxter was assessed by the Chinese HMO expert. How the competitor might respond to various offers from Baxter (given that collusion is legal in China) was assessed by people who have negotiated with this competitor. The strategy tree was a vehicle for putting together all the information obtained from the various experts, and for discussing the various assumptions with management.

About one year later, Baxter's involvement in the production of distilled water gave rise to another game-theoretic study. Baxter market share was being threatened by new entrants with new technologies. Even though the new technologies were at a very early stage, Baxter's strategy experts were concerned that not acting early could lead to a substantial loss in this highly profitable business. Some of the questions raised were the following: Should we buy off the competition? Should we partner with them? Should we compete with them and try to over take them in the development and patenting of the new methods? If we compete, should we develop the technology and try to pass FDA regulation first in the U.S., first in Europe, first in Asia, or simultaneously in several locations?

This was more complicated than the China problem, since it involved a strategy of learning. (Does the new technology work? Can we assess, or even affect, the likelihood of FDA approval in the U.S. by having it tested first in less demanding regions? Can we learn about the demand in the U.S. by observing demands in other parts of the world?) Again, the method of constructing a large strategy tree proved effective. This time, our experts included engineers, FDA experts, marketing experts and more. At the end of several months of work, the tree and its assumptions were presented to upper management. The presentation was effective, and upper management decided not to ignore the new technologies and to devote several million dollars to study the issues in detail.

In addition to being effective, working with strategy trees is natural. In both of the above projects we interacted with two highly capable MBAs from Baxter. They caught on to the approach well, and our services were needed less and less as the work progressed. Unlike traditional game theory, the method involved much subjective optimization with a tree that was trimmed to the most essential issues at hand. We developed effective methods of sensitivity analysis that help in the construction and trimming of strategy trees.

The consulting work I describe next taught me lessons about the need for different methodology, but also pointed to additional difficulties in game-theoretic consulting. Arthur Andersen was a large accounting and consulting partnership headquartered in Chicago. It had two divisions, with hundreds of offices and thousands of partners around the world. Being a partnership, it used a formula to allocate its substantial profits to its partners. The profit-sharing formula was supposed to satisfy several objectives. It should be fair, so the partners would agree to it. It should create incentives to increase Arthur Andersen's profits, and it should have insurance aspects that prevent the income of partners from fluctuating too much due to the random annual fluctuations of the business in various geographical areas. I was hired by one of the senior partners to advise about the profit-sharing formula.

It became clear that the cooperative, rather than the strategic, game-theoretic approach would be more appropriate for the construction of a profit-sharing formula.[2] The monotonicity property of the Shapley value, which I studied earlier by myself and with Dov Samet, creates incentives for partners to exert efforts on behalf of the company. But it also became clear that computing the precise Shapley value of the problem at hand is (still) hopeless, and even if it were possible, the Shapley value theory is (still) not sufficiently developed. First, we needed a Shapley value applicable to a cooperative game with incomplete information. Second, the Arthur Andersen game was one that kept changing and evolving with time, so we needed a dynamic Shapley value with incomplete information. But these were not the only difficulties.

---

[2] A strategic model would require a precise description of the moves available to all the players, the order of moves, the information available at the time of choosing moves, etc. This is clearly impossible.

With a substantial use of ad hoc methods, I developed a Shapley-value-like formula for their dynamic game with incomplete information. However, the senior partner I worked for repeatedly put off considering my solution. It became clear that he really did not want a new solution, but only wanted to use me in order to justify the solution that they were using at the time. He wanted to shut down arguments from the opposing division, by stating that a game-theory expert specializing in bargaining and fairness approved his method. Despite the nice consulting fee I was collecting, I quit this job after learning an important lesson: Game theorists who think they have been hired to solve the game may turn out to be mere pawns in the game.

This phenomenon is not unique to Arthur Andersen. In many situations where companies hire consultants, it is often done for political and strategic considerations, rather than for the advice the consultant may offer. In game theory, where the consultant's recommendation typically affects players with competing objectives, the problem is more severe. We need to find ways to make the decision to hire the game theorist be independent of his recommendations.

Question 3. What is the proper role of game theory in relation to other disciplines?

In my roles of past president of the Game Theory Society and an editor of a game theory journal, I recently constructed a list of areas of activities in game theory and its applications. The list should give the reader an appreciation of the interactive nature of game theory.

1. **Non cooperative game theory** studies the behavior of payoff-maximizing players who take into consideration all strategic and informational parameters.
2. **Cooperative game theory** studies how considerations of efficiency, fairness and stability guide the allocations of profits and costs to coalitions of rational players.
3. **Behavioral game theory** studies how real players play games: **experimental games** played in the lab, and **empirical games** played in the real world.
4. **Evolutionary game theory** studies play guided by imitation, survival of the fittest, etc.
5. **Algorithmic and artificial game theory** study issues of computational, informational, and behavioral complexity in games played by live players or by computing machines.
6. **Interactive epistemology** studies the subject of knowledge, including knowledge about knowledge.
7. **Combinatorial games** deal with mathematical issues unique to games.
8. **Non-Bayesian decision theory** concentrates on decision making under uncertainty, when relaxing or replacing the Bayesian assumptions made in the classical theory.
9. **Neurological studies of games** deal with physiological activities observed during the play of a game.
10. **Economic games** use the above tools to gain insights into strategic economic interaction and the performance of economic systems.
11. **Political games** use the above tools to gain insights into strategic political behavior and the performance of political and social systems.

Game theory has been described by some as the physics of the social sciences. I think that another useful analogy involves probability and statistics. Probability theory offers a language and rules for dealing with uncertainty, and statistics offers tools for real-world applications. These theories are designed to deal with uncertainty, no matter where it arises. Similarly, game theory offers a language and rules to deal with strategic interaction, wherever it arises. The Arthur Andersen and Baxter applications described earlier suggest an interesting parallel with statistics.

When dealing with a real-life application, a statistician must choose the best among an available set of models: a classical approach, a nonparametric approach, a Bayesian approach, etc. An applied game theorist must choose between a coalitional model, a strategic model, or a hybrid of several models.

But thinking of game theory as similar to probability and statistics makes the relationship to other sciences clear. First, like probability theory, a well-developed game theory is foundational in any subject that deals with interaction. Second, in practical applications that involve strategic interaction, there is no way to avoid using game theory. However, following the practice of statistics, it may be necessary to have several different game-theory models, with the user choosing the appropriate model for the application.

As we have discussed, the interaction of game theory with economics over the previous century has concentrated mostly on theoretical applications. And indeed, in a similar manner to the use of probability theory, the use of game theory has become unavoidable in essentially all rigorous studies of strategic economic phenomena. We see that political science is going through a similar progression, even at a more fundamental level. Unlike economics, were much formal modeling was done prior to the arrival of game theory (through equilibrium models of supply and demand, for example), the initial formal modeling of political systems had to start with the use of game theory.

The interaction with evolutionary biology and computer science is interesting because there is a reciprocal fertilization between game theory and these other subjects. Since evolutionary biology studies interaction among species, it is natural to apply game theory there. But the reverse is also true. Evolution theories describe how the behavior of species evolves, without resorting to rationality but using concepts such as imitation, survival of the fittest, etc. An important finding there is that despite no reliance on rationality, species' behavior converges to what is predicted by Nash equilibrium. This is an important connection. It makes both theories more robust, and it is the subject of much current research.

Similar reciprocal fertilization is present in the interaction of game theory and computer science. Both game theory and computer science share a common goal: the mathematization of rational choices and behavior. Historically, however, computer science concentrated mostly on algorithms that generate rational choices, subject to complexity constraints. For a long time it ignored strategic and interactive aspects in systems that involved more than a single decision maker. Game theory, on the other hand, ignored issues of computation and complexity and concentrated mostly on the strategic interactive aspects. In recent years, we have seen a growing interaction and cross-fertilization between the two fields.

Questions 4 & 5. What do you consider the most neglected topics and/or contributions in late 20th century game theory? What are the most important open problems in game theory and what are the prospects for progress?

What questions should a theory address? Questions should be sufficiently simple so that answers can be obtained, yet they should be sufficiently advanced so that the answers are meaningful. An important issue is the level of detail that is assumed as input to a theoretical model. In the case of game theory, where rational players are part of the model, the details must include the level of complexity that can be handled by the players. Game theorists refer to a model that assumes limited information or limited ability to handle complexity as one with bounded rationality.

A beautiful example of a theory of bounded rationality is the older subject of probability. Imagine going to a scientist before probability theory was invented, to ask for his advice on how to bet (H or T) on a coin. The fully rational answer is as follows. Find out the initial position of the coin, the rotational velocity, and the height to which the coin will be flipped, and the hardness of the material on which it will land. Then apply this data to the differential equations from physics that describe the trajectory of the coin. If you know all of these parameters and you are able to do the computation, you should know what side the coin will land on, and you should bet accordingly.

Clearly, the differential equations of physics described above are not the right tool to deal with the issue of betting on a coin. Both difficulties prevail: it is impossible to get all the necessary information, and it is impossible to do the computation, even if the information were available. Viewed in this light, probability is a wonderful theory of bounded rationality. Rather than attempting to answer the question of what side the coin will show, it addresses the question of what are the frequencies of H and T, if the coin is flipped many times. This latter question can be answered by experimentation, which leads to a reasonable recommendation for our decision maker.

In many respects, game theory is like decision theory before the introduction of probability. We are still wresting with the identification of questions to answer, and how to go about the rational modeling of bounded rationality. For many applications, our models require the knowledge of too many parameters and assume unrealistic computational ability. For these reasons, we are often restricted to what social scientist refer to as stylized, or toy, models.

For older successful sciences, the identification of good models was guided by real-life applications. I think that game theory should move more in this direction. There are two related guiding questions: how to construct a game and how to play a game. The "how to play" question should go beyond 2 by 2 Prisoners' Dilemma games, and deal with more complex and realistic problems. (I do not mean to say that understanding the Prisoners' Dilemma game is unimportant – just as it was important for Newton to explain the effect of the law of gravity on an apple falling from a tree before moving to the movement of the stars. But the move from Prisoners' Dilemma to real life games may require substantial new approaches.) To outsiders who are not familiar with the communities of the researchers, it seems strange that the chess-playing program that defeated the best human player is of no interest to game theorists.

The chess example points to an area that needs addressing in game theory, which may be called "macro game theory." In chess, one uses ad hoc functions to rank board positions. For example, a position where a player controls more squares on the board is superior to one in which he controls fewer. It is better to have more pieces than the opponent. While not true in every board positions, global considerations like these still guide the decision of chess players and programs.

Global considerations, like the ones in chess, are also used in our everyday decisions. For example, you are better off majoring in a subject that leads to more career opportunities. You are better off choosing actions that make friends, rather than enemies. While considerations like these are a part of our every-day real-life interactions, they are practically absent in strategic game theory.

Another important step for future game theory may be the development of additional sub areas and specializations, in a similar way to the development of statistics as a different subject from probability. While highly related, these subjects require different knowledge, tools and skills, and separation and specialization proves to be useful. A useful analogy is obtained by thinking of the construction of and flying of airplanes. It takes physicists to develop the underlying basic theory, aeronautic engineers to design a plane, and pilots to fly it.

While the engineers have basic knowledge of the physics involved, they have additional practical knowledge obtained by experimenting in wind tunnels and by learning about earlier plane designs. The pilots have basic knowledge of the physics and engineering, but have additionally needed skills.

The design and play of auctions may require decomposition similar to the one above. Game theorists offer the underlying theory. In addition to basic game theory, auction designers should have knowledge of behavioral theories, obtained in the lab and by studying the play of earlier auctions. Successful bidders understand the general principles and the rules of the auction, but have additional abilities to assess specific environments and opponents.

Recent criticism of game theory is that knowledge of game theory does not make one a better game player. While this may or may not be true, the airplane analogy is useful. It seems unlikely that a physicist would be good at flying a plane (or that a pilot could publish papers in physics), yet it does not mean that physics is useless in the construction of planes. Purely on knowledge of game theory, one is not likely to do well bidding in an auction. But this knowledge is useful for both the auction designer and the bidders.

It seems that game theory is starting to evolve in this direction. The development of the behavioral areas may constitute the first step in creating more practical knowledge, similar to the one of an aeronautical engineer. But it is important to recognize the need for these specializations and encourage, rather than resist them.

My hope is that this next century will lead to the development of more useful game-theory. The problems addressed will continue to come from the social and biological sciences, but also from concrete real-life (as opposed to theoretical) applications. The methodology will rely on tools, concepts and ideas from other areas such as computer science, evolutionary biology, psychology, and others.