9

Sergiu Hart

Kusiel and Vorreuter University Professor

Institute of Mathematics, Department of Economics, and

Center for the Study of Rationality

The Hebrew University of Jerusalem, Israel

Interviewed at the 17th International Conference on Game Theory at Stony Brook University, July 13, 2006, by Pelle Guldborg Hansen.

Q: Professor Hart, why were you initially drawn to game theory?

A: The short answer is "Aumann." Bob Aumann came to Tel Aviv University to give a course on game theory when I was a second-year math student there. At the time I didn't know anything about game theory, but it sounded interesting. I took the course, and became quite excited. What I found so exciting about game theory, and still do, is that it is such a varied topic—in terms of questions, answers, methodologies, and so on. In mathematics one specializes in one area, say probability theory, and then does probability theory. But in game theory one can do probability theory one day and combinatorics the next, as well as logic, computer science, and biology. Let's just say there is never a dull moment. It's very varied, and for me that was extremely attractive. Of course, so was Aumann's personality. He knows how to fire people up. So that's how I started, and I don't regret it.

Q: What was your first impression of Aumann? Did you know that he was going to lure you into game theory?

A: No, I was just a student in his class. Then in my third year of studies he conducted a seminar. We were about ten students, and Aumann took us through the work of Schmeidler and Kohlberg on the nucleolus, which was an important breakthrough in those years. Every week Aumann would give us another set of problems, and those were tough problems. The following week, those who had managed to solve the problems would present their solutions. I remember that there were three of us who were competing neck and neck. We really had to work very hard and get into deep issues; the seminar was great! That's how I decided to do my M.Sc. with him, and later on, my Ph.D.

Q: What examples from your work best illustrate the use of game theory for foundational studies and/or applications?

A: I'm originally a mathematician, so most of my work is foundational. My work hasn't revolved around applications, even though they are extremely important. One cannot do game theory in a completely abstract way, without any roots in economics, biology, computer science, political science, and so on. If one wants to do something interesting, it is important for that something to be *relevant*! But I'm mainly on the foundational or theoretical side of game theory. My Ph.D.-thesis and some of my work throughout the years has been in cooperative game theory: axiomatizations, values, coalition formation. But I have also worked in noncooperative game theory, for instance, on repeated games. Since the nineties I have done much work with Andreu Mas-Colell on dynamic models. This includes adaptive heuristics and various dynamics, like regret matching, that lead to correlated equilibria and Nash equilibria.

Dynamic models are now quite an exciting area in which many people are working, so let me explain the general setup. Early game theory concepts were static concepts; Nash equilibrium, for example. Though there certainly is a dynamic intuition in the background, the definition is static, in the sense that this is a rest point. But that always leaves one asking how those equilibria are reached. If the players start at an equilibrium point, they will most probably stay there. But if they don't start at an equilibrium point, and they are reacting (say, best replying or better replying), or if it is an evolutionary process, then where does it all lead? Will it converge to this or that equilibrium, or not? These are difficult questions. While static analysis is relatively simple—not that it is simple, but it is *simpler than* dynamic analysis—the mathematics of dynamic systems is very complex.

On top of this, the class of interesting dynamics is huge. First, there are highly rational dynamics, where people observe what happens, calculate (say, the posterior probabilities), and optimize accordingly. This requires a lot of rationality and a lot of computation.

Second, there are dynamics that are essentially mechanistic and automatic, like evolutionary dynamics. In such setups there is no conscious computation or optimization by the participants. Instead, some process like natural selection makes the frequency of successful strategies increase. In addition, there are mutations – a kind of small random "noise" – that introduce new strategies. What happens is that those mutations that are successful keep expanding in the population.

And third, there are dynamics, like adaptive heuristics, where there is just a little rationality, that is, highly bounded rationality. The players act in simple and myopic ways that seem to be going in a "good" direction. Though the players are far from fully optimizing, in the long run their behavior may nevertheless yield the same outcomes that fully rational players would have achieved; take, for example, the simple adaptive heuristics of Hart–Mas-Colell that lead to correlated equilibria. So, to get back to your question, some of my recent work is on dynamics. But I have worked also on various other topics. For instance, I just gave a talk at the conference here in Stony Brook on the sure-thing principle and agreement theorems.

Q: For those who missed your talk, would you elaborate a little on this work?

A: It starts with joint work with Bob Aumann and Motty Perry. In decision theory, the famous sure-thing principle of Savage says the following. If I decide something when I know A, and I decide the same thing when I know B, then I should decide the same also when it is either A or B but I do not know which one. Now, that sounds exactly like the standard sure-thing principle of logic, right? In logic: if "A implies C" and "B implies C," then " 'A or B' implies C". *But it is not* the same thing, because in logic it does not matter whether A and B are compatible events or not, whereas in decision theory it turns out that it is essential that these events are *not compatible*, that is, that they are *disjoint*. If they are not disjoint, then the sure-thing principle of decision theory need not apply, and one may get into trouble using it. This is something that people don't realize, and it took the three of us quite some time and many arguments to put our finger on what exactly matters for the sure-thing principle, and why. And that came as a surprise. In fact, I gave a simple example to a full room here at my talk. At first everybody agreed that it was right, but then they saw that it was completely wrong. It takes people by surprise.

Aumann and I then went on to try to understand how to extend to decisions his famous agreement theorem that people with a common prior probability cannot disagree on their posterior probabilities when these are commonly known. Obtaining a decision-theoretic version of this result turns out to be conceptually difficult, though mathematically easy. One needs to understand what conditions are needed for it to apply. This, like much of my other work, is very foundational.

Q: Are there other examples of your work that you would like to mention?

A: Well, there is my work with Andreu Mas-Colell on the connection between strategic approaches and coalitional approaches. Let me explain. Game theory has two main branches: noncooperative and cooperative. The essential distinction is that noncooperative game theory deals with the strategic approach and the resulting strategic equilibria, whereas cooperative game theory asks questions such as what agreements players should reach. For instance, assume that there is a pile of money to split, and the players have the possibility of signing binding agreements—how will they split the money? To determine that, they will of course take into account what their other options are, in various coalitions. That's the focus of cooperative game theory: it is about coalitions and what outcomes the players should agree upon. Now an important issue here is how to make the connection between the noncooperative and cooperative branches. This is the classical "Nash program" started by John Nash; John Harsanyi is also instrumental here. It is about providing strategic foundations to cooperative game theory. Coalitions form, operate, decide, and divide the proceeds through strategic interactions between the players. The problem is that the strategic and bargaining procedures are not well defined; there are many possibilities to take into account, such as who proposes first and who second, what the rules for acceptance are, and how everything is conducted. Andreu Mas-Colell and I are trying to find procedures that are very general, yet give nice insights into what is happening. That's the direction our work took in the nineties, and we have just revived our interest in it (Andreu gave a talk about this here). It is an important direction, and I think that we have some new nice insights.

Q: What do you consider the proper role of game theory in relation to other disciplines?

A: Game theory is universal. It is relevant to all the disciplines in which people make decisions. It is also relevant to other disciplines, like evolutionary biology. Genes don't make decisions, but when you model evolution formally—the model of genes interacting, and *interacting* is the important issue here—then you see that it is a game-theoretic model. The insights and understandings of game theory become important here. But it also works the other way around: we game theorists learn from biologists. They develop something and we say, hey, that's a good idea. In fact, it has been a very fruitful connection in both directions. For instance, when the biologist John Maynard Smith introduced game theory into evolutionary biology in the seventies, it was quite rudimentary, but it caught on and flourished; then ideas from biology, like replicator dynamics and evolutionary approaches in general, came back to us and developed into what is now a very big area of game theory. Today the insights, tools, and concepts of evolutionary biology are used in game theory, and in economics too. Open *Econometrica* and you'll see not only theoretical papers that use evolutionary models, but, more generally, that the evolutionary paradigm is important and useful in economics as well.

Computer science is another relevant discipline. Computers make decisions too, and networks of computers interact, coordinate, or fight over resources; in short, game-theoretic problems need to be solved. The question of how to design the rules, the "protocols," in order to obtain desirable outcomes belongs to the game-theoretic area called "mechanism design." Nowadays, with all the electronic commerce, computer science is getting very heavily into game theory.

And one could go on. There is political science. There is philosophy. We have already touched upon logic and interactive epistemology. Basically, there are many, many areas to which game theory is relevant. It's like mathematics in the sense that a physicist uses mathematics to formulate what he is trying to model and explain. In the same way, an economist, an evolutionary biologist, or a political scientist uses game theory to formalize his insights and ideas. Game theory provides the tool for analyzing interactive situations, in which, unlike in simple optimization problems, one person alone cannot determine the outcome. What makes this a "game" is that what I do influences what happens to you, and what you do influences what happens to me. I thus have to take into account your rationality and your optimizing, and you have to take into account my rationality and my optimizing, and so on. At this point it all seems like one big mess, but game theory succeeds in cutting this Gordian knot; for example, by pointing to this or that kind of equilibrium. Game theory is thus a methodology that is applicable to the social sciences in general, as well as to other sciences like biology and computer science. I don't know if it is used in physics (but I recall some physicists in Jerusalem who presented a paper on a game-theoretic question). It is probably not used in any real sense in chemistry; I don't think that atoms are playing games—but one never knows. But in the social sciences, law, philosophy, biology, computer science, clearly game theory is an important tool.

I've compared game theory to mathematics, but it needs to be emphasized that game theory is not just a branch of mathematics. It is an applicable science, and as such relevant ideas and insights are essential. We had here at Stony Brook a "Nobel Session" two days ago. Bob Aumann and Tom Schelling, the Nobel Prize winners of last year, were both there. I gave an introduction to Aumann's work, and then Dick Zeckhauser gave an introduction to Schelling's work. At some point Zeckhauser quoted an article of Avinash Dixit on Schelling in the *Scandinavian Journal of Economics*, where he gives advice to the "budding economic theorist": do not just try "relaxing the condition of semi-strict quasi-concavity to hemi-demi-proper pseudo-concavity," but rather "obtain your primary motivation from life." I buy this statement completely. Taking existence theorems and merely improving the conditions here and there is not enough. (Of course, that doesn't mean that that should not be done. It is important when it allows us to expand the range of models and attack new problems.) You are not going to make a successful career out of doing *just that*. Ideas and concepts are essential. Intuition and understanding are essential. That was the point of Dixit and Zeckhauser.

But I want to emphasize that this is not the end of the story. It is *also essential to formalize* these ideas. Only by doing so do you realize that although your insights may sound convincing and look good, they do not give the full picture. If you just think of it conceptually you may miss many of those things. It is only when you try to prove formally what you think is the result of your model that you realize that it doesn't work; it might not work, for instance, because additional conditions are needed. Having beautiful and important insights and then being able to establish them formally is the right way to proceed. Game theory is not a "verbal" science. We are trying to be precise, because it's important to formalize the ideas. Aumann, for example, is a mathematician, but his great success does not come from just proving theorems. His success comes from taking a new concept, a new idea, and first of all being able to formalize it and strip it of all that is irrelevant, getting it to its barest. And then, when this is done, the results, theorems, and proofs usually become simple and clear. A significant part of Aumann's work, but not all, is like that; for example, the agreement theorem that we talked about. It's very important that you be able to formalize and to see whether your intuition is indeed correct. Now, if you have no intuition and you're just extending existing results, you're not headed for a successful career. It is the combination of intuition and formalization that is essential. That's an important lesson to learn. Now Schelling is not a mathematician and his contributions are mainly conceptual. But there is only one Schelling. And many people went on to formalize his beautiful insights by building the models and doing the math. As for Aumann, his forte is both on the conceptual side and on the mathematical side, and that's really the winning combination.

Q: That brings us back to the second question of good examples of game-theoretic work. That would mean that work like Aumann's not only uses mathematical tools but rests on new ideas and insights.

A: Yes, mathematics is a *tool*. You need an idea that is interesting and important and relevant, and you need to formalize it, which is the only way that you can verify that your intuition is correct. Verbal arguments are nice, but they are only the beginning; they cannot be the end. If you really want to verify that your idea is correct and that you're not missing something very important, then you have to be able to prove it formally.

Q: Would you touch on the famous quotation from the *Handbook of Game Theory*, where Aumann and you talk about game theory as a unified theory of the social sciences?

A: If I remember correctly, it says that game theory is like a unified theory for the *rational side* of the social sciences. It doesn't say "for the social sciences." For example, there is no claim that psychology is a subset of game theory. There is much in psychology that is not game theory, but at the same time there are areas of psychology, like rational decision-making, where game theory is of much relevance.

Q: But how far do these rational sides extend? That is, how far do the insights and ideas of game theory extend into the social sciences?

A: When you study rationality, you also study bounded rationality and irrationality. In addition to asking what happens in an ideal world where the decision maker is fully rational, you ask what happens when he is bounded—in his computations, in how much effort he makes, in his optimization. Beyond the ideal case of *homo rationalis*, there is a large part of game theory that deals with models of bounded rationality. It would be preposterous to say that everybody is fully rational. One should understand rationality and the limits of rationality. For example, what Danny Kahneman and Amos Tversky say is very relevant: it is important to understand under what circumstances we have biases and make errors in our decisions. Clearly we have biases; I know I do, and I try to correct them, but it's not easy. When I hear about a fifty percent discount, I tend to head for the store, but then I stop and think, wait a minute! I'm going to gain perhaps 10 dollars, or 50 dollars. Compare that amount to a one percent reduction in the price of the last house I bought. People will spend more effort on the former fifty percent than on the latter one percent, despite the fact that that one percent on the price of a house may be a huge amount, perhaps a few months' income. But they won't spend enough effort in trying to get that one percent reduction, by shopping around and bargaining. It doesn't come easily and naturally to us; we have to think about it. Also, we cannot always be the kind of rational person who computes everything. If you try to do that you'll be run over by the first car on the street. It's not a good survival strategy!

Q: What are the limits of rationality in behavioral economics?

A: People are definitely making mistakes when making decisions. They have biases and they make lots of errors. However, when it counts, when it *really* matters, we make far fewer mistakes. John List gave a beautiful example in an *Econometrica* paper a couple of years ago. He conducted his experiments on the floor of a sportscard market, where he found clear evidence of the so-called "endowment effect." This refers to valuing an object more when you own it than when you do not own it. Specifically, when people are given the choice of getting, say, a mug or a chocolate bar, about half the people take the cup and half the people take the bar. But if you give them, say, the mug, and ask them if they want to exchange it for the chocolate bar, very few agree; and the same if you give them the chocolate bar. Once they own it, they do not want to trade it for something else worth about the same. They value what they have more than the thing they don't have, while they are indifferent between the two objects when they do not own either of them. One needs to be careful in interpreting the results here, because they are all in a region where people are more or less indifferent. You might offer me a cup which is worth a few dollars or you might offer me a chocolate bar which is worth the same. I'll take the cup or the chocolate bar, I don't really care which; I just have to make a decision. And then keeping what I have when I am indifferent is a good and simple rule: "taking the path of least resistance." Also, in general I have more information on what I own than on what I don't own, and so risk aversion can explain this effect.

As I said, List found a very strong endowment effect among the people coming to the sportscard market. But then he also ran his experiment with the professional traders in that market—and, as it turned out, they had no endowment effect! The endowment effect vanishes perhaps because those traders who have it lose money and drop out of the market, or perhaps because they learn the hard way that that's not the way to trade. If you want to make money you had better make sure that you don't have biases that can hurt you ...

Q: ... when the stakes are high?

A: When the stakes are *real*, when they matter. We make errors all the time, we have many biases. But when it really matters, when you depend on it, you make far fewer errors.

Q: What do you consider the most neglected topics in late twentieth-century game theory?

A: Game theory is a dynamic subject. It evolves, and so it is hard to say that there are "neglected topics." There are so many people working in so many directions and new things are opening up all the time. There is no neglected topic in the sense that people should work on it and no one does. It's like evolution: if a topic is interesting, someone will work on it. Well, I will if no one else does. Of course, some topics are extremely difficult, we lack the tools to analyze others, and we don't know how to approach still others. It may take a lot of time until one starts developing the machinery.

For example, Harsanyi had a beautiful idea on how to deal with games of incomplete information. Everybody wanted to analyze this clearly important topic, but nobody knew how. Then Harsanyi came up with the idea of transforming a game of incomplete information into a game of imperfect information by using "types." This allowed, inter alia, the development of a whole new field, that of (repeated) games of incomplete information, which subsequently bloomed. It's not that this field was neglected before; it was just too tough. It takes time until somebody finds an opening.

I'm not sure that I can point to a neglected topic. The boundary of the field keeps being pushed out. Now sometimes there are biases in science. Sometimes a topic that lots of people are working on may turn out not to be that important. For example, too much effort was put into equilibrium refinements, and they became quite esoteric at some point. It hasn't died out, but it has reached a certain maturity and is no longer as active as before. Evolution works beautifully in science. There are always many ideas, and the good ones spread and develop; "mutations" produce new ideas, and the good ones catch on, the bad ones die out.

Q: What are the most important problems in game theory today?

A: Dynamics is a crucial topic to understand. Most of our interactions are not static, one-time interactions, but rather repeated interactions. We interact day after day, whether with the same people or with different people. There may be repeated interactions, or things may change from one day to the next: the game, the people I play with, the environment, etc. All in all, dynamics is an important topic. A lot of work is being done now on dynamic models, and it is advancing very nicely.

Another area that is picking up quite significantly is the interface of game theory and computer science. With the Internet and all the electronic commerce, people realize that they have to understand game-theoretic notions, like mechanism design and auctions. Then there is also "algorithmic game theory," which deals with the problems of computing and finding equilibria, particularly in large games, like those involving big networks.

Speaking of applications, the latest *MIT Technology Review* selected the ten most significant emerging technologies of the year; "cognitive radio" was one of them. Think of all the cell phones, pagers, laptops, and wireless devices that everyone carries around nowadays. How can they use the bandwidth efficiently? Centralized allocation protocols are not practical; it has to be decentralized. Computer scientists have started looking into cognitive radio, programming each device to use a set of behavior rules that are based on game-theoretic ideas. Interestingly, the day after I found out about cognitive radio, I got an e-mail from another group of researchers working on something similar for cellular communication, who were trying to apply regret-matching algorithms from my work with Andreu Mas-Colell (but don't blame me if your cell phone stops working ...).

A further area of application has to do with congestion games, which arise in transportation problems. More and more cars are equipped nowadays with GPS navigation devices. Today these are mostly one-way machines: they get information on your location from satellites, and then compute your route. But there is no reason that it shouldn't be a two-way machine, where your GPS transmits your location. When a route becomes clogged with too many cars, the individual GPS devices could start routing cars to a different road. There are clearly problems due to the huge number of nodes—cars, intersections, routes—and the fact that they have to be solved online in real time. In all these areas game-theoretic dynamic approaches are very relevant. You could use adaptive dynamics that will lead the system to an equilibrium, or perhaps close to one. These are just a few of the many areas where game theory is very applicable.