



ACADEMIC  
PRESS

Available online at [www.sciencedirect.com](http://www.sciencedirect.com)

SCIENCE @ DIRECT®

Games and Economic Behavior 45 (2003) 2–14

GAMES and  
Economic  
Behavior

[www.elsevier.com/locate/geb](http://www.elsevier.com/locate/geb)

## Presidential address

Robert J. Aumann

*Center for Rationality and Interactive Decision Theory, and Departments of Mathematics and Economics,  
the Hebrew University of Jerusalem, 91904 Jerusalem, Israel*

Received 21 February 2002

---

### Abstract

The Presidential Address at the First International Congress of the Game Theory Society, held in Bilbao, Spain, in July of 2000. The address contains a discussion of the Congress, of the functions and activities of the Society, of the Logo of the Society, of past accomplishments of the discipline, and of some future directions for research. The address is preceded by an introduction by David Kreps.

© 2003 Elsevier Inc. All rights reserved.

---

**David Kreps:** Yesterday I asked Bob what he would like me to say; as he often does, he responded with a reminiscence, about how he used to listen to introductions at political conventions on the radio. And how mayors, senators, governors would be given flowery introductions explaining who they were, what they had done; but when it came time to introduce President Roosevelt, the formula was very simple: “Ladies and Gentlemen, the President of the United States.” Having said in as many words that simple and brief would be better than long and convoluted, Bob concluded that he hoped in this setting he would need no introduction.

Well, in fact, he needs no introduction. But, he does deserve something a *bit* longer. Now, I won’t trouble you with biographical data or a list of his honors; I don’t think I have the time. And I certainly won’t trouble you with a list or a recapitulation of all his contributions to our discipline, because I certainly don’t have time for that. Instead, I’ll simply share with you that aspect of Bob Aumann that has always struck me as most characteristic.

---

*E-mail address:* [raumann@math.huji.ac.il](mailto:raumann@math.huji.ac.il).

0899-8256/\$ – see front matter © 2003 Elsevier Inc. All rights reserved.  
doi:10.1016/S0899-8256(02)00545-6

I had the great fortune to have been raised intellectually at Stanford University in the early 1970s, when in the summers, the IMSSS<sup>1</sup> was in full flower. In the summers for six weeks or eight weeks, week after week, four times a week, all of economic theory would present itself for inspection in front of Ken Arrow, Frank Hahn, Bob Aumann, and the rest of the company. Sometime in the summer, it would be Bob's turn to get up and give a seminar. Whenever that was, he would always begin with his slightly diffident, slightly insistent tone to develop whatever subject he had that day. I seem to recall—I may have forgotten, but I seem to recall—that he would always begin with an example, usually a very small and simple example, that over the course of two hours would be rotated and inspected, developed and examined, and then by some trick of expositional magic it would turn into a very powerful and general theory. Aumann used this expositional technique, from simple example to more complex example to theory, because it was always extraordinarily effective at teaching the audience what he meant to teach.

And that, for me, is the essence of Bob Aumann. He is, of course, a brilliant and creative scholar. He is, of course, one of the pioneers of our subject. But first and foremost, he's a spectacularly gifted teacher, who never fails to engage his audience while teaching them something new and exciting, deep and powerful.

So, without further ado: Ladies and gentlemen, the Charter President of the Game Theory Society.

**Bob Aumann:** Dave, thank you very much.

OK, the plan for today is to talk a little bit about the Congress, then about the Game Theory Society, then about the logo of the Game Theory Society (Fig. 1), and then a little about future directions. So it's not going to be a presentation of research, or teaching, as Dave mentioned, but something more "presidential."

This congress is really a magnificent occasion. Many of us are absolutely overpowered by it. It's by a factor of three larger than the previous large game theory congress, which was held in Jerusalem five years ago; a little over two hundred people came to that one, two hundred twenty or two hundred thirty. Here we have over six hundred participants.

It's a magnificent place, a beautiful city. It's marvellous, and very exciting. One can't go to all the lectures, but at the ones to which one does go, the rooms are always filled to overflowing, and the ideas are magnificent. The whole thing is just very exciting.

This didn't come to be by itself. It took a lot of work, a lot of organizing, a lot of preparation, and I'd like to acknowledge some of those people who made it possible. Already last night at the Guggenheim Museum, I acknowledged the local organizing committee, under the chairmanship of Federico Valenciano, who did a magnificent job. The other members of the committee—María Paz Espinosa, Federico Grafe, Elena Iñarra, and José Manuel Zarzuelo—of course were also instrumental in making this possible.

And then there's the program committee. We have Vince Crawford and Joel Watson, who were responsible for economics, management, experimentation, evolution, a large part of the program. We have Tim Feddersen, who was responsible for political science applications and applications for the other social sciences. We have Yoav Shoham, who brought the very special computer science twist to this conference, which is playing a big

---

<sup>1</sup> Institute for Mathematical Studies in the Social Sciences–Economics.

starring role over here, and we're certainly grateful for his enthusiastic participation; and Sylvain Sorin, who did the theory, the mathematical part of it. Everything over here has been tremendously exciting, so we are very grateful to all these individuals.

And I'm grateful to you, all of you, who've really made this possible. The participants, the speakers, and all the "senior" people, who gave their stamp of approval, specifically by their participation: John Nash, Ken Arrow, Reinhard Selten, Lloyd Shapley. So, you've all made it a wonderful event. And, there's one individual who really pulled it together and made it possible, and that's the producer (OK, these are the credits) Ehud Kalai. (Applause)

Now let's talk about the purposes of the Game Theory Society. Why was the Society formed? There are a number of reasons. First, we have the ordinary purposes of any learned society. There are many learned societies in the world, and this is one of them, so we have the ordinary purposes, which include journals. The *International Journal of Game Theory* and *Games and Economic Behavior* are the two official journals of the Game Theory Society, and we are thinking of founding some other journals, or at least one other one for the moment. The new journal, perhaps electronic—probably electronic—would not compete with the *International Journal of Game Theory* and *Games and Economic Behavior*, but it would be something else; it would be for review articles and for research announcements, somewhat in the spirit of many of the journals in computer science, which have extended abstracts, which afterwards grow to journal articles.

So that is one thing, the journals, and of course we have meetings, which is another one of the ordinary purposes of a learned society. We have this Congress, which is our first meeting; we're planning another Congress in four years—we can't pull together something of this magnitude every year. It's better to wait four years, and we invite proposals from people and places who would like to host this Congress in four years. If we keep going at this rate, we have a factor of three, so we'll have about two thousand people in 2004. We are also envisioning smaller regional and smaller, better-defined disciplinary meetings. I'll say a little bit more about that right away. These are some of the ordinary activities of any learned society.

There are also other activities, that have to do with the special nature of game theory. And by the special nature, I mean the very broad, interdisciplinary sweep of this subject. There are very few subjects that have such a broad, interdisciplinary sweep. Let me just put over here some of the ordinary disciplines that are involved in game theory. We have mathematics, computer science, economics, biology, (national) political science, international relations, social psychology, management, business, accounting, law, philosophy, statistics. Even literary criticism; Steve Brams once wrote an article with game-theoretic analyses of various items of literature. One of them was "The Gift of the Magi," by O. Henry. This turns out to be very closely related to game theory; the main point of that was the battle of the sexes, and how people who do not maximize utility can get to one of the "bad" outcomes in that game. We have sports;<sup>2</sup> we have a recent analysis of championship tennis (Walker and Wooders, 2001); it turns out to be a very good verification of the minimax theorem.

---

<sup>2</sup> A special session on Game Theory in Sports was held at the Bilbao Congress. It included papers by Walker and Wooders ("Minimax play at Wimbledon"), Chiappori and Levitt ("Do soccer players randomize?"), Palacios-Huerta ("Game theory in the grass"), and Kirman and Hardle ("When to accept and when to refuse").

We have all these things; it's very highly interdisciplinary. In addition to that, we have the tools of analysis. We have the whole gamut of tools of analysis in science; we have theory, experiments, empirics, and most recently, and perhaps most importantly, engineering.

This sort of dictates a broader scope for the Society, and let me write down here some of the possibilities. We plan to have not only a specific learned society, but we plan to have chapters and affiliates, national chapters; we've already had expressions of interest, we expect to have the Russian Game Theory Society, the Italian Game Theory Society, the Indian Game Theory Society. We plan to have other national affiliates. We plan to have disciplinary affiliates, like the International Society for Dynamic Games; dynamic games are like differential games, and stochastic, and repeated games, and things like that, that's dynamic games. There is a society that is devoted specifically to those things, and they will be affiliated with the Game Theory Society.

We are interested in having educational programs to teach game theory in schools and things of that nature, because it's not a well-defined discipline like economics or mathematics, which has its own curriculum that's well-established, but it's a sort of language of discourse. That is one of the things that we need within game theory itself also, within the people who are over here; we have to develop a common language. We could, for example, have a society for computers and games, or games and computers.

So, there are things of this kind that are dictated by the interdisciplinary character of our "racket," and we can be active in those things.

I'd like to take this opportunity to thank the people who have been active, not only in organizing this Congress, but in running the Society. So, we have the officers of the Society and the editors of the journals. We have Ehud Kalai, who is the Executive Vice-President of the Society; he produces not only the Congress, he produces in fact the whole Society. He is also the person who instigated it, who initiated the Society. We have Eric van Damme, who is the Secretary-Treasurer of the Society; we have Adam Brandenburger, who is the Communications Vice President; specifically, he runs the webpage of the Society, he developed it. We have Dov Samet, who is the editor of the *International Journal of Game Theory*, and here's a new face, Ehud Kalai, who is the editor of *Games and Economic Behavior*, one of the two official journals of the Society.

And, we have a very distinguished advisory board. Here it is: Arrow, Debreu, Harsanyi, Nash, Selten, and Shapley. Their presence as the advisory board of the Society, their presence within the Society, is very important and vital in making this go.

Let's get to the next item, namely the logo (see Fig. 1). Inside the triangle there is a game tree; this tree represents the strategic—sometimes called non-cooperative or competitive— aspects of game theory. On the outside, there is a triangle, with three distinguished points. Many of you will know that this represents the unique symmetric stable set of a three-player majority game (von Neumann and Morgenstern, 1944), which historically, was the first solution concept in coalitional game theory; thus it is symbolic of the coalitional, cooperative approach to game theory. So, this is in some sense symbolic of coalitional game theory. The logo says that coalitional and strategic game theory are really two parts of the same whole; that both viewpoints are important, they complement each other, they're part of the same unity.



Fig. 1.

There is a beautiful word that was mentioned by Dave Kreps in the Don Jacobs Lecture yesterday. In this magnificent lecture, he used the word “co-opetition,” which had previously been used by Adam Brandenburger and Barry Nalebuff.<sup>3</sup> That says it all. It means that those two aspects of game theory are really not two separate disciplines, they are part of the same whole. The strategic and the coalitional viewpoints. But they are not only viewpoints; each one has its function, and they’re related to each other, they complement each other.

Let’s try to lay that out a little. Strategic game theory—competitive, non-cooperative game theory—is at its best when you have well-defined rules; you are able to get equilibria, you are able to compute things and get well-defined answers. This is becoming more and more important as we move into the Internet era, the electronic era, when a lot of these business situations will be defined in a very precise way, and we’ll be able to deal with them using the tools of strategic game theory.

Coalitional game theory is more suited to situations when the rules are not well-defined, and when you don’t say who moves first and who moves second; when it’s a matter of power relationships, things of that kind. There is the strength of coalitional game theory. That kind of situation will persist. Even in the era of the Internet, nobody’s going to tell you on the Internet which large firms merge with each other; nobody’s going to tell you how to form a government coalition on the Internet. You have very important insights that come from coalitional game theory, where the rules are not well-defined, the sequence of events is not all there but you can get general insight. It’s a little similar to the difference between looking at something close-up and from far away. Yesterday at the Guggenheim Museum, there was this exhibition of avant-garde Russian woman painters in the first two decades of this century. There was a young man with me who was not familiar with the history of modern art, and he was standing in front of this cubist painting and couldn’t make any sense of it. I said, “let’s move away a little, let’s back up five, ten meters, then look at it again.” When you look at it again from five or ten meters away, it becomes much clearer. So, one could perhaps say that the strategic theory is the “micro” theory, you’re

<sup>3</sup> In their book *Co-opetition*, Currency/Doubleday, New York, 1996. They write (pp. 4–5): “Business is War and Peace. But it’s not Tolstoy—endless cycles of war followed by peace followed by war. It’s simultaneously war and peace. As Ray Noorda, founder of the networking software company Novell, explains: ‘You have to compete and cooperate at the same time.’” After their book was published, quite a few people got in touch with them, claiming to have coined the term “co-opetition” independently of Noorda. Apparently, it had been used in the business community for some time.

looking closely at what's going on in a game; the coalitional theory is the "macro" theory, you back up a distance. And then some things, that are not clear when you're looking at it closely, become clear when you're looking at it from far away. Both are necessary; you need both.

There are important bridges, of course, between coalitional and strategic game theory, the cooperative and competitive situations. I'll mention just two. In repeated games, when you analyze a repeated game from a strategic viewpoint, you look at equilibria of repeated games. It turns out that the equilibria of the repeated game are closely related to cooperative, coalitional solutions of the one-shot games. In some sense, the strategies in a repeated game make explicit the kind of considerations that are implicit in a coalitional analysis.

Another important bridge is bargaining theory. We have, for example, the Nash two-person bargaining solution. The two-person bargaining problem can be realized strategically, so to speak, by an explicit bargaining game; this is one of the important contributions of Rubinstein (1982). There are other approaches of this kind in more general  $n$ -person games. So, we find important bridges, important relationships, important ways of expressing similar ideas in both repeated games and bargaining theory, and there are other bridges that tie the two aspects together. So, these two objects, the tree and the triangle here, symbolize two important aspects of the same thing.

Let's press on, and discuss some challenges of game theory, basically for the future. But, before we look at the future, note that people who try to predict the future, or discuss the future, rarely look at what discussions of this kind did in the past. Newspaper columnists always write analyses about what will happen in the future, but nobody ever reads yesterday's, or last year's, newspaper columns. Let's try to avoid that mistake.

You know, I have a friend, a businessman actually, not a game theorist, but he *is* a game theorist in some kind of practical sense. He does business in tens of millions of dollars, and he said to me once, "Johnny," he said—I have another name in addition to Bob—"Johnny, I just got a letter from one of the banks; an offer to send me their annual newsletter, and they want two thousand dollars for it. You're a game theorist, you pretend to be an economist, should I buy this?" I said, "you know what, Marcel, send them a check for a hundred and fifty dollars, and tell them to send you last year's newsletter." So, he said, "good idea," and he sent it to them; they returned the check, and they didn't send the newsletter.

I want to avoid that mistake. I do want to look back fifty years. We could call this the "Club of Rome syndrome." The Club of Rome was something that was popular several decades ago, predicting the future. They predicted what will happen in 2000. I don't remember what they predicted, but nobody's interested. Now I *am* interested in what was said fifty years ago about the future of Game Theory. At that time, the introduction to Volume I of the *Contributions to the Theory of Games* (Kuhn and Tucker, 1951) listed fourteen problems, or challenges, of game theory. The list was drawn up in April of 1950; it was the work of Harold Kuhn, who is with us here today, with a lot of input from Lloyd Shapley, who is also with us here today. Let's look at some of the problems, and see what became of them. We'll find that they were remarkably prophetic. Kuhn and Shapley did a good deal better than the Club of Rome or the bank's newsletter.

*The 1951 Kuhn–Shapley problems (partial list)*

1. Existence of stable sets.
2. Valuations.
3. Many players.
4. Dynamic game theory.
- ...

One of the problems was the existence of von Neumann–Morgenstern stable sets, to which we alluded before. That was not known at the time. This was a specific, well-defined mathematical problem; there were a number of others, and this was the most difficult. The problem—namely, do  $n$ -person, side-payment coalitional games have a stable set or not—was finally answered in the negative by Bill Lucas (1969); he provided the counterexample. It was solved twenty-five years after it was first posed in von Neumann and Morgenstern's book, twenty years after appearing in the Kuhn–Shapley list. So, this was solved.

Another problem listed there was the problem of valuations. Evaluating a game; this of course was solved spectacularly by the introduction of the Shapley value (Shapley, 1953) shortly thereafter, with its many applications, its many ramifications. So, this was another problem that was successfully dealt with. It's not only a question of successfully dealing with these things; this became a central item of interest in research from then until this very day. It's not some odd thing in which people are no longer interested; it's a major item.

Next is the matter of many players: to get significant asymptotic results, significant representations of games with many players. This was 1950; nothing was known about that at the time. This also became a major item of concern, of analysis, starting with the work of Shubik (1959) about markets with many players. And then it continued with Shapley, Milnor, Shapiro, Debreu, Scarf, your humble servant, Schmeidler, and many, many others. Inter alia, it indicated that something that von Neumann and Morgenstern had conjectured was indeed true: In markets with many players, game-theoretic solutions lead to the laws of economics: to competitive equilibrium, the law of supply and demand. This was verified in the decades following, the sixties and seventies.

And just today, just one hour ago, there was a magnificent presentation by a name that I haven't mentioned up to now, Ehud Kalai. He gave a beautiful presentation on the asymptotic properties of strategic games with many players (Kalai, 2000). The bottom line of that is that mixed strategies are not important when you're talking about many players; basically, in any game with sufficiently many players, the properties of mixed strategies are mimicked by pure strategies. So, the subject of many players has played a fundamental role in game theory, and in its applications to economics and also to political science, in the theory of elections with many players which has been very well-developed. This is another verification of the insight that went into these problems in the 1950 list.

Another one of the Kuhn–Shapley problems is the dynamic theory of games. That was a very difficult problem. Most of game theory until this day is concerned with equilibrium notions. It asks, when is the system at rest? But what about the dynamics of the system? It's very difficult to talk about that in game theory, because a rational agent is presumed to look ahead, and not just to respond to forces that act on him at some particular moment. He's expected to take into account what may happen in the future, and that makes any dynamic theory very difficult. And it was very difficult for many years, even after this problem

was suggested. And now, since the advent of evolutionary theory into game theory in the last decade or two or three, it turns out that game-theoretic equilibrium is just about the same thing as a population equilibrium in biology. That suggests, of course, that game-theoretic dynamics can be represented by population dynamics—a fruitful idea, which has in fact been followed up; evolutionary dynamics have been studied for a decade or so. This is a very promising, very important field. Evolutionary game theory gives us a door for studying dynamics in game theory. This is also a very important item, and at least we see the beginnings of a successful attack on this.

So Kuhn and Shapley look rather good at this point, in the sense that the problems they discussed really did turn out to be important, central problems. It's unlikely that I can do as well over here, but let me try. So, let's call this "some directions."

*Some directions for future research*

1. Stochastic games.
2. Evolution.
3. Coalitional games with incomplete information.
4. How much to compute.
5. Epistemology in imperfect information games.
6. Engineering:
  - Expert labor.
  - Auctions.
  - Elections.
  - Cost allocation.
  - Cake cutting.
  - Formal bargaining (like final offer arbitration).
  - Internet.
7. Endogenous tastes.
8. Consciousness.
9. Cryptography.

Now is the year 2000. It was in 1900 that David Hilbert proposed at the International Congress of Mathematicians his famous Hilbert's Problems. No doubt the problems I'm suggesting here do not have a similar degree of centrality to game theory, but they're important problems, or at least they are problems that interest me, and they might interest whoever's interested in coming along and working something out.

The first problem here is that of stochastic games. This, like the first Kuhn–Shapley problem, stable sets, is a well-defined mathematical problem. And it's the *only* one of my problems that's mathematically well-defined, i.e., prove or disprove; there's no question as to what it means, no conceptual component in the formulation. We don't have time to define stochastic games precisely, but roughly speaking, a stochastic game is a dynamic game, you play it again and again, and there are finitely many states; at each stage, you play a game, and this determines not only the payoff for that stage, but to what state you will go, what you will be doing at the next stage, what game you'll be playing next. Stochastic games were introduced by Shapley in the early fifties; with a discount factor, he proved that all two-person, zero-sum stochastic games have a value. This was extended by Bewley and

Kohlberg (1976) to some asymptotic<sup>4</sup> models without discounting. The final proof of the existence of a value in two-person, zero-sum undiscounted games was given by Mertens and Neyman (1981); and there the matter stood for a long time. The problem of two-person non-zero-sum games was very, very difficult, and remained open, in spite of intensive research, for many years. Then finally, Nicolas Vieille (2000) proved the existence of equilibria in undiscounted two-person non-zero-sum games. This is a spectacular result.

The problem of general,  $n$ -person, undiscounted, finite-state stochastic games remains open. It appears to be very difficult; there's no counterexample—where an equilibrium does not exist—and there's no proof. This is a mathematically well-defined problem; it's very important, the outstanding problem of game theory of that kind; that is, not something that's "develop this or that," but "prove or disprove."

Evolution is a direction for the future development of game theory, both social evolution and biological evolution; we've already said something about that. Also, equilibrium, and population dynamics in the theory of evolution.

An item that is important, and that has been developed somewhat but not enough, is the matter of coalitional games with incomplete information. I don't think we have, so to speak, the "right" answer, though there's been a good bit of work on it.<sup>5</sup> We're talking about coalitional games defined by a "characteristic function"  $v(S)$ , games of the kind discussed above, but where the players don't know what  $v(S)$  is; they have incomplete information. The matter of incomplete information was successfully formulated by Harsanyi for strategic games; here we're looking for some kind of parallel formulation for coalitional games. Again, the way to deal with this may well be bridges; in other words, let's look at repeated games with incomplete information; let's see what are the strong equilibria of repeated games with incomplete information, and maybe this will give us a handle on coalitional games with incomplete information.

A very important problem, on which some work has been done, but not very satisfactory, is how much to compute. A lot of our solutions depend on heavy computation. Sometimes one feels it is not worthwhile; the cost of computation overshadows the cost of what you're going to get as a result. On the other hand, if you don't compute, you won't know how much to compute. So, it's really a very puzzling conceptual problem; some work has been done on it, but it doesn't really solve the problem. Perhaps the way to go is some kind of evolutionary approach; but, it's not clear how to work it out. It's very puzzling, and it's very important: How to know how much to compute without computing?

My own "racket" in the last few years has been epistemology, knowledge theory, some of the things that Yoav Shoham was talking about yesterday. In particular, the area of perfect information games has been explored thoroughly; that exploration is not finished yet, we're still in the throes of it. But, extending it—seeing what notions like common knowledge of rationality say in non-perfect information games—is an important subject.

Engineering: One of the most significant aspects of game theory lately has been the engineering aspect. Here are some of the subjects: expert labor markets, which have been explored with great success by Al Roth and his associates (Roth and Sotomayor, 1992;

<sup>4</sup> E.g., they showed that the value of an  $n$ -stage stochastic game converges when  $n \rightarrow \infty$ .

<sup>5</sup> See, for example, Allen (1997). A more recent survey is Forges et al. (2003).

Roth and Peranson, 1999); this was mentioned in the Jacobs Lecture yesterday by David Kreps. Auctions, of course, is a matter of extreme importance in game-engineering. But there also are other matters. “Engineering” means when game theorists suggest practical solutions to real-world problems of some complexity. So, we have auctions (Wilson, 1992), we have elections (Brams, 1994), cost allocation (Young, 1994), cake cutting (Brams and Taylor, 1996). We all know what cake cutting means: division of resources, in some sense “fair” as well as optimal; like in a divorce. Steve Brams has worked a lot on that, that’s a typical engineering problem. We have formal bargaining: suggesting ways of overcoming bargaining impasses; for example, final-offer arbitration. There is the whole issue of building the Internet in a strategically optimal manner; Dave mentioned that yesterday, and he mentioned a number of other topics on the engineering front.

He also mentioned endogenous tastes; that is a very important problem. Maybe the right way to treat it is by evolution; why is it evolutionarily important for me to like coffee and for you to like tea, or hot chocolate? What is the evolutionary function of tastes, of utility? Rather than taking those things as given, one wants to account for them in some way.

Ahhh, consciousness. It’s the central problem, not only in game theory, but in all of science. To make sense of consciousness is the most important problem in science. When I was a child, people said, well, we understand the way the planets move, the way the stars move; we understand gravity, we understand geology and chemistry. But we don’t understand life. Life, we don’t understand. There’s no good scientific account of life.

Now, sixty years later, one can say that we do understand life. With the advent of DNA, there’s a good understanding of how life works, in some sense. It’s not a mystery anymore.

But consciousness remains a mystery. It’s a mystery, and it has important game-theoretic and evolutionary connections. So I’m willing to say that that’s the number one problem in science, and it’s also an important problem in game theory.

Let’s enlarge on that. The way we understand life, it’s a mechanistic thing. There are molecules, and they interact in a certain way, they fit together, they reproduce. How things got to be that way is explained by evolution; survival of the fittest, and so on. It’s like a giant, complicated, incredibly intricate machine; better, it *is* a giant, complicated, incredibly intricate machine. But in the end, it’s just that—a machine.

Very good. That explains almost everything about life. But not everything. It explains the life of trees, of flowers, of ants, of mice, and of men. To me, gentle listeners, it explains *your* lives. But not *my* life.

Why? Because there is one—and only one—thing that I know *for sure*; namely, that I am more than just a machine. I think, I see, I hear, I experience pleasure and pain. I *experience*. Descartes said, “*cogito, ergo sum.*” But it is not just “*cogito.*” It is the whole gamut of experience—*ergo, sum. I am conscious.*

Now I really *know this only* about myself. Presumably, gentle listeners, also you think, see, hear, experience pleasure and pain. You look and act more or less like me, so you probably experience like me. But I can’t be sure of that. I am *sure* only about myself.

There is no way that I can verify that you are indeed conscious in this sense, that you do experience. Everything that I observe about you can be explained by a mechanistic model. But the mechanistic model cannot explain the one observation that I make for sure—namely, that *I* am conscious. Paradoxically, this one incontrovertible fact is itself scientifically unverifiable.

Machines are not—cannot be—conscious. They cannot experience. Or perhaps they can; but that would be a giant leap forward into the unknown and uncomprehended. Indeed, I do not comprehend my own consciousness either, cannot account for it scientifically, even remotely. That is the next frontier of science, perhaps the final frontier.

Some people tell me that they don't understand the problem; why can't a machine be conscious? I can understand their question only by supposing that they themselves are not really conscious, so they don't understand what consciousness means.

So, it's a fascinating problem. But what does it have to do with game theory?

In his book, *The Growth of Biological Thought*, Ernst Mayr (1982) distinguishes between two kinds of explanation in biology, corresponding to the questions “how” and “why.” The question “how do we see” may be answered “we see with our eyes—with their lenses, retinas, neural connections to the brain, and so on.” On the other hand, the question “why do we see” may be answered “we see because it helps us enormously in getting along in the world.” In brief, “how” refers to the mechanism, “why” to the function.

The problem of consciousness set forth above is a “how” problem; it concerns the mechanism of consciousness. Game Theory comes in at the “why” end; it explains, or may explain, the (evolutionary) *function* of consciousness.

For evolutionary success, the organism needs food; it needs to reproduce; it needs protection from the elements and from predators; and it has other, secondary needs. All plants, as well as many animals (especially those that are “lower” on the evolutionary scale) are “programmed” to accomplish these ends. For example, a Venus flytrap detects, by means of a well-understood mechanism, when an insect enters; then it closes up and digests the insect. No volition or even true sensation is involved; neither is there any volition or sensation on the part of the plant in the process of photosynthesis, pollination, the drawing up of water from the roots, or the growing of thorns or poisons to defend against predators. It is all entirely mechanical. (Purely physical and chemical processes are included under the heading “mechanical.”)

When a human being eats, reproduces, dresses, takes shelter, or evades predators, the process is more complicated. Take the case of food. Before eating, one either feels hungry, or experiences a desire for food that “tastes good,” i.e., the eating of which causes pleasure. In one case, there is discomfort or pain (hunger); in the other, pleasure. To avoid or end the pain, or to achieve the pleasure, one eats.

But unlike with plants, with human beings eating is a complicated process. Primitive hunters need to select what to hunt, construct and use weapons, stalk the quarry, skin and prepare it, and so on; similarly with gathering. The more “advanced” the society, the more complicated the process. We have to go to a store, buy the food, undo the package, cook it, and so on. Even more complicated, we have to earn the money to buy the food and the stoves and the refrigerators and so on. Programming all that is perhaps not entirely impossible; as we all know, evolution has achieved astounding degrees of complexity, even in areas that are entirely “hard-wired.” Still, it sounds a little improbable.

So what I'd like to suggest is that consciousness—specifically, the sensations of pleasure and pain—serve as a kind of decoupling mechanism, like prices in economics.

The price mechanism allows economic agents to achieve full coalitional efficiency<sup>6</sup> by *individually* maximizing over budget sets rather than engaging in complex multi-sided barter operations; consciousness enables organisms to obtain food, say, by maximizing pleasure (and/or minimizing pain) rather than being directly programmed to do all the things that are necessary to obtain food. Similarly for sex; that is the function of pleasure in sex.

Of course, all this goes only towards answering the “why” question for consciousness. The “how” question remains a deep mystery.

The last problem or direction in our “Club of Rome” list is cryptography. It’s a little surprising that Shoham didn’t mention that today. It’s a very important element in game-theoretic analyses: ways of communicating with cryptographic methods; one-way communications, or communications between subsets of players. This is beginning to play more and more of a role<sup>7</sup> in game theory.

And that about ends this presentation of some directions in which to go. Perhaps some of these will turn out to be significant in the future, and let’s go to lunch!

## References

- Allen, B., 1997. Cooperative theory with incomplete information. In: Hart, S., Mas-Colell, A. (Eds.), *Cooperation: Game-Theoretic Approaches*, Proceedings of a NATO Advanced Study Institute held at Stony Brook, NY, July, 1994. Springer-Verlag, Berlin, pp. 51–65.
- Bewley, T., Kohlberg, E., 1976. The asymptotic theory of stochastic games. *Math. Oper. Res.* 1, 197–208.
- Brams, S., 1994. Voting procedures. In: Aumann, R., Hart, S. (Eds.), *Handbook of Game Theory with Economic Applications*, Vol. 2. Elsevier, Amsterdam, pp. 1055–1089.
- Brams, S., Taylor, A., 1996. *Fair Division*. Cambridge Univ. Press, Cambridge.
- Debreu, G., Scarf, H., 1963. A limit theorem on the core of an economy. *Int. Econ. Rev.* 4, 235–246.
- Forges, F., Minelli, E., Vohra, R., 2003. Incentives and the core of an exchange economy: a survey. *J. Math. Econ.* 38, 1–41.
- Kalai, E., 2000. Private information in large games, DP 1312, Center for Mathematical Studies in Economics and Management Science, Northwestern University, November.
- Kuhn, H.W., Tucker, A.W. (Eds.), 1951. *Contributions to the Theory of Games*, *Ann. of Math. Stud.*, Vol. 24. Princeton Univ. Press, Princeton.
- Linial, N., 1994. Game-theoretic aspects of computing. In: Aumann, R., Hart, S. (Eds.), *Handbook of Game Theory with Economic Applications*, Vol. 2. Elsevier, Amsterdam, pp. 1339–1395.
- Lucas, B., 1969. The proof that a game may not have a solution. *Trans. Amer. Math. Soc.* 137, 219–229.
- Mayr, E., 1982. *The Growth of Biological Thought*. The Belknap Press, Cambridge, MA.
- Mertens, J.F., Neyman, A., 1981. Stochastic games. *Int. J. Game Theory* 10, 53–66.
- von Neumann, J., Morgenstern, O., 1944. *Theory of Games and Economic Behavior*. Princeton Univ. Press, Princeton.
- Roth, A., Peranson, E., 1999. The redesign of the matching market for American physicians: some engineering aspects of economic design. *Amer. Econ. Rev.* 89, 748–780.
- Roth, A., Sotomayor, M., 1992. Two-sided matching. In: Aumann, R., Hart, S. (Eds.), *Handbook of Game Theory with Economic Applications*, Vol. 1. Elsevier, Amsterdam, pp. 485–541.

<sup>6</sup> I.e., a point in the core. This is the easy part of the “Equivalence Theorem;” see, e.g., Debreu and Scarf (1963).

<sup>7</sup> For a nice survey up to 1994, see Section 3.2 (pp. 1360–1368) of Linial (1994). A more recent contribution is Urbano and Vila (2003).

- Rubinstein, A., 1982. Perfect equilibrium in a bargaining model. *Econometrica* 50, 107–109.
- Shapley, L.S., 1953. A value for  $n$ -person games. In: Kuhn, H.W., Tucker, A.W. (Eds.), *Contributions to the Theory of Games II*, Ann. of Math. Stud., Vol. 28. Princeton Univ. Press, Princeton, pp. 305–317.
- Shubik, M., 1959. Edgeworth market games. In: Luce, R.D., Tucker, A.W. (Eds.), *Contributions to the Theory of Games IV*, Ann. of Math. Stud., Vol. 40. Princeton Univ. Press, Princeton, pp. 267–278.
- Urbano, A., Vila, J., 2003. Computational complexity and communication: coordination in two-player games. *Econometrica* 70, 1893–1927.
- Vieille, N., 2000. Two-person stochastic games. *Israel J. Math.* 119, 55–126.
- Walker, M., Wooders, J., 2001. Minimax play at Wimbledon. *Amer. Econ. Rev.* 91, 1521–1538.
- Wilson, R., 1992. Strategic analysis of auctions. In: Aumann, R., Hart, S. (Eds.), *Handbook of Game Theory with Economic Applications*, Vol. 1. Elsevier, Amsterdam, pp. 227–279.
- Young, H.P., 1994. Cost allocation. In: Aumann, R., Hart, S. (Eds.), *Handbook of Game Theory with Economic Applications*, Vol. 2. Elsevier, Amsterdam, pp. 1193–1235.