# 1

# Robert Aumann


## Professor of Mathematics

Center for the Study of Rationality

and Institute of Mathematics

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 Aumann, how and why were you initially drawn to game theory?

A: Well, the story is this: I graduated from MIT in 1955 with a Ph.D. in pure mathematics. In fact, I had already left MIT in 1954 to work for the Forrestal Research Center, which was connected to the mathematics department at Princeton, where I did my post doc. At Forrestal they were doing operations research; very, very practical operations research, not theoretical stuff at all. There I came to work on a problem about defending a city from air attack. This was being done for Bell Telephone Laboratories, which was developing a missile under contract with the Defense Department. The problem had to do with a squadron of aircraft, a small percentage of which were carrying nuclear weapons; that is, most of them were decoys. How would one respond to this? What would be the optimal strategy to deal with that kind of situation? They sent this problem over to us, the Analytical Research Group at Princeton, and the problem was assigned to me. At that time I didn't know very much about game theory at all. I had met John Nash at MIT, when he was a young instructor and I was a senior graduate student. There we became quite friendly, and naturally he told his friends about game theory: duelling and that kind of thing. I wasn't particularly attracted to it, but I listened, and some of the problems were quite interesting. When I was assigned the aircraft problem at Princeton, I remembered these conversations with Nash and then realized that it was a game theory problem. I read a little bit about game theory, thought a little bit about it, and did what I could with that problem. In fact, I wrote a report. It may be in my files someplace, but I don't know, it is more than fifty years ago. That started me thinking about game theory, and from there I went on to more theoretical things. At that time Princeton was the center of the world in game theory, so naturally I started attending some of the game theory conferences there and mixing with the game theory people and became more and more interested. Also, I read the book of Luce and Raiffa, which got me interested in repeated games. One thing led to another. By the time I went to the Hebrew University, I was already identified as a person whose major interest is game theory. So I started giving courses in game theory, and when you teach something you get interested in it yourself. That is the story: it wasn't something that I sought out, but something on which I sort of stumbled on the path.

Q: Were you working on the Folk Theorem before you went to the Hebrew University in 1956?

A: I came to Israel in October of 1956, but I was not aware of the Folk Theorem then. In fact, I was not even working on repeated games before coming to Israel. In the summer of 1957 I was at the National Bureau of Standards, which is now in Bethesda, but which at that time was still in Washington, inside the district line. Over there I was working hard on the paper on repeated games, and I was aware of the Folk Theorem by that time. I think the major part of my work on this, the 1959 paper on acceptable points, was generated at the Hebrew University.

Q: You did much groundbreaking work on formalizing the Folk Theorem. What other examples from your work illustrate the use of game theory for foundational studies and applications?

A: That's a big question; all of game theory is about either applications or foundational studies; that's all there is, right? Everything, the whole theory, is about that. But I could single out the work on interactive epistemology, and, certainly, the work on the equivalence theorem. You might look on them as being either applicational or foundational. The equivalence theorem on the competitive equilibrium of the market uses very many different concepts from game theory: the core, the Shapley value, even the bargaining set. It is foundational if you're looking at it from the point of view of an economist. There's also the idea of correlated equilibrium. That's certainly foundational.

Q: I was just talking to Sergiu Hart a few minutes ago, and he, like many others, mentioned your work on the Folk Theorem and the agreement theorem, but in particular, he emphasized those parts of your work that most beautifully couple conceptual ideas with mathematical proofs.

A: Sergiu has made a wide range of very important contributions. Most recently, he's been working on heuristic methods for reaching equilibria; in other words, not simply algorithms, but dynamic processes leading to equilibria: Nash equilibrium, correlated equilibrium, etc. However, he's also done very important work on characterizing the value, in particular, distinguishing between the Harsanyi value and the Shapley value in non-transferable utility games. He introduced the idea of potential to characterize Shapley values and the core. In fact, some of his early work was about applying the von Neumann-Morgenstern stable set idea to the formation of oligopolistic markets. The von Neumann-Morgenstern solution is fundamentally different from other solution concepts in that it does not, in the case of markets, predict competitive equilibrium. It predicts the formation of large cartels; that is, large groups of people who cooperate with each other in order to compete with other groups. Basically, it says that a large number of small traders does not guarantee free competition. If you read the *Theory of Games and Economic Behavior* of von Neumann and Morgenstern – and I think these parts of the book were written by Morgenstern – you'll find that they did not really believe in the idea of free competition, of perfect competition. They felt that the "right" solution concept, so to speak, or at least the only one that they had for cooperative games, was the von Neumann–Morgenstern solution – the stable set idea – and this really does not lead to price competition.

Q: In the introduction to the first volume of the *Handbook of Game Theory* you and Sergiu Hart discuss game theory as an umbrella, or as a unified ...

A: ... a unified field theory. Yes, I think game theory really is unique in this respect. Other disciplines, like economics, take various real-life situations, like oligopoly, monopoly, large markets, international trade, and each one of those situations becomes a problem unto itself. Economists devise methods for dealing with this or that problem, but there is no overall methodology. Game theory, on the other hand, has the advantage of defining things very, very broadly, so that the concepts that apply to game theory apply to any game. Well, in principle they apply to any game. The core, for example, might be empty in certain games, but still it applies. In fact, any interactive situation – whether monopoly, duopoly, large markets, international trade, or, to look beyond economics, elections, bio-systems of various kinds, political systems, international relations – any one of these interactive situations is a game to which you can apply the methods of game theory, the same methods. Take the nucleolus. If you apply it to a market situation, it will yield free competition, it will yield price competition, it will give you prices. But you can also apply it to elections and it will give you minimal winning coalitions and things like that. So the same ideas apply to very different contexts. In that way, it's like one theory that gives you magnetism, gives you electricity, and gives you gravitation—it's an umbrella, that's the idea.

Q: In 2005 you were awarded the Nobel Prize in Economics together with Thomas C. Schelling. In 2002 Daniel Kahneman and Vernon Smith were likewise awarded the Nobel Prize. Some people in the press took the 2002 prize as support for a transition from classical economics to experimental economics. How do you view this issue?

A: It is interesting that the 2002 Nobel Prize in Economics was awarded to Kahneman and Smith. The behavioral economist Smith became famous for verifying the claims of economic theory with economic experiments. The Nobel Committee was not making a definite statement, as some have thought, that the assumptions and conclusions of classical economic theory are incorrect. Rather, they were saying, let's at least recognize experimental methods in this. Now, one kind of interesting conclusion was reached by the psychologist Kahneman, and another kind of conclusion – the opposite kind of conclusion, really – was reached by Smith. It's true that in the ensuing years behavioral economics gained some drive, but I'm not sure that this is going to continue. A lot of voices are challenging behavioral economics. One problem is that the conclusions of behavioral economics are based to a large extent on questionnaires and polls. These are notoriously unreliable sources. What people *say* they'll do is not what they *do*. Another problem is that the conclusions are based partly on experiments with money. Though there is some incentive given, this is usually small and people don't really pay much attention to it. True, some of the conclusions are also based on market behavior. But here again, we are not talking about large effects, things that are important to people. We are talking about whether they buy something at the checkout counter that they didn't really want to buy—and who cares about that?

I don't think that behavioral economics is going to last, though I think that it's an interesting idea. For instance, I agree with the behavioral economists that people don't think about the decisions they make. Maybe they don't think about them at all, or maybe they think about them very little. But I have this idea of rule rationality as opposed to act rationality. People act on the basis of what they have gotten used to, and what has worked in the past in the kind of situation that they are in. But this, on the whole, is usually rational. That is, people evolve rules for behavior that work on the whole and then sometimes, just sometimes, they apply them when it is inappropriate to apply them. So there is some validity, some truth, in what behavioral economists say. It is important that we face the challenges that they pose, but on the whole I'm not convinced, and I'm not the only one. Ariel Rubinstein, for example, ask *him* about behavioral economics. And many other prominent people are, let's not use such a strong word as "rejecting," but many other prominent people are sceptical, very sceptical, that behavioral economics will survive the test of time.

Q: And just to make sure, empirical economics is something quite different?

A: Yes, empirical economics is very important. Empirical economics existed before behavioral economics, and will continue to exist afterwards. It's natural that we look for something that has a connection with the real world; and indeed, empirical economics often fits very well with received economic theory.

Q: If some topics have received too much attention at times, then what do you consider the most neglected topics in late twentieth-century game theory?

A: I don't have any strong feelings about a topic that's neglected. If I feel it's neglected I will work on it. Still, the coalitional theory is more important than many people think it is. There is some kind of myth that this is useless stuff; but, most of the insights of game theory in the past have been products of the coalitional theory. I'll give you just two examples. There is the work on the

Shapley value as applied to elections, various governing bodies, or systems of government like the UN, the US Congress, and so on. This has yielded a lot of insight. Now, of course, the Shapley value also yields a price equilibrium when you apply it to economic systems. So for these applications you have all the equivalence theorems, all the results that relate game-theoretic concepts to price equilibria, and that's all cooperative game theory. The same goes for the Gale-Shapley work on matching markets, which has been applied by Roth and Sotomayor and others to many real-life contexts. That's also cooperative game theory. In fact, I think in applications you have more cooperative theory than non-cooperative theory. On the other hand, there are auctions. That's an important area and auctions are non-cooperative theory. So I don't want to be misunderstood as saying that cooperative theory is the only or the most important branch. There are people that take sides. They say "I'm a non-cooperative game theorist," or "I'm a cooperative game theorist"; and "if I'm non-cooperative, I'm going to say that the cooperative stuff is useless ...," or vice versa. I don't have such a feeling of patriotism with regard to a particular branch of the theory. I work on both. Both are important. I think that maybe – if you were to say that it was being neglected – the cooperative theory deserves a little more attention. Still, we saw several presentations on cooperative theory here at the conference. It's not so bad off that I would say "neglected." Cooperative theory is alive and kicking.

Q: So no real neglected topics?

A: No, I have no complaints about the way the theory should go or has gone so far. We are looking forward to a lot more interaction between game theory and computer science and various computer systems, but I wouldn't call that a neglected topic. It is something that has to be developed, that will be developed, and is already being developed.

Q: What do you think are the most important open problems in game theory?

A: I'll mention just one that was touched on in the conference here, and that is the problem of computational costs. This is not necessarily just a game-theoretic problem. It is a decision-theoretic problem. Let me explain. When you have a problem that involves some computation, how much should you invest in the computation? If you had some idea of how long it will take to do the computation, then you'd be okay—but you don't! Trying to figure that out may involve a computation that's more difficult than the one you're talking about. On the other hand, although it's not easy to say when you're playing a game of chess how much time you should spend on a specific move, we all somehow succeed in solving the problem at hand. So how much thinking should you invest in a given problem?

This is not an easy problem. To make progress, we have to treat it as a conceptual problem. Though you may figure out how long it will take to do the computation, just figuring that out may not be worthwhile. You have to try to approach it from some new angle. It is not clear what to do in that kind of situation. It's not your classical optimization problem. In classical optimization you just say, well, we're going to solve this problem, and then you optimize the amount of gas to put in the tank or whatever, without caring how long it takes to compute. But here you are talking not about how much gas to put in the tank, but rather how much time to compute. You *don't* want to spend more time computing how much to compute. So there is a conceptual problem here and I'm not even sure what *kind* of solution we will get. It might be some kind of evolutionary thing, without precise answers. It is a difficult conceptual problem that has been around for many years, and I don't know what to do with it.