AN INTERVIEW WITH ROBERT AUMANN [*]

Interviewed by Sergiu Hart [**]

Who is Robert Aumann? Is he an economist or a mathematician? A rational scientist or a deeply religious man? A deep thinker or an easygoing person?

These seemingly disparate qualities can all be found in Aumann; all are essential facets of his personality. A pure mathematician who is a renowned economist, he has been a central figure in developing game theory and establishing its key role in modern economics. He has shaped the field through his fundamental and pioneering work, work that is conceptually profound, and much of it mathematically deep. He has greatly influenced and inspired many people: his students, collaborators, colleagues, and anyone who has been excited by reading his papers or listening to his talks.

Aumann promotes a unified view of rational behavior, in many different disciplines: chiefly economics, but also political science, biology, computer science, and more. He has broken new ground in many areas, the most notable being perfect competition, repeated games, correlated equilibrium, interactive knowledge and rationality, and coalitions and cooperation.

But Aumann is not just a theoretical scholar, closed in his ivory tower. He is interested in real-life phenomena and issues, to which he applies insights from his research. He is a devoutly religious man; and he is one of the founding fathers—and a central and most active member—of the multidisciplinary Center for the Study of Rationality at the Hebrew University in Jerusalem.

Aumann enjoys skiing, mountain climbing, and cooking—no less than working out a complex economic question or proving a deep theorem. He is a family man, a very warm and gracious person—of an extremely subtle and sharp mind.

This interview catches a few glimpses of Robert Aumann's fascinating world. It was held in Jerusalem on three consecutive days in September 2004. I hope the reader will learn from it and enjoy it as much as we two did.

SH, Jerusalem, January 2005

---

[*] This is a shortened version of the interview. The full version is available at http://www.ma.huji.ac.il/hart/abs/aumann.html, and will appear in *Macroeconomic Dynamics*.

[**] Center for the Study of Rationality, Department of Economics, and Department of Mathematics, The Hebrew University of Jerusalem, Feldman Building, Givat Ram Campus, 91904 Jerusalem, Israel.

E-mail: hart@huji.ac.il *Web page*: http://www.ma.huji.ac.il/hart

Sergiu HART: Good morning, Professor Aumann. Let's start with your scientific biography, namely, what were the milestones on your scientific route?

Robert AUMANN: I did an undergraduate degree at City College in New York in mathematics, then on to MIT, where I did a doctorate with George Whitehead in algebraic topology, then on to a post-doc at Princeton with an operations research group affiliated with the math department. There I got interested in game theory. From there I went to the Hebrew University in Jerusalem, where I've been ever since. That's the broad outline.

Now to fill that in a little bit. My interest in mathematics actually started in high school—the Rabbi Jacob Joseph Yeshiva (Hebrew Day School) on the lower east side of New York City. There was a marvelous teacher of mathematics there, by the name of Joseph Gansler. The classes were very small; the high school had just started operating. He used to gather the students around his desk. What really turned me on was geometry, theorems and proofs. So all the credit belongs to Joey Gansler.

Then I went on to City College. Actually I did a bit of soul-searching when finishing high school, on whether to become a Talmudic scholar, or study secular subjects at a university. For a while I did both. But after one semester it became too much for me and I made the hard decision to quit the yeshiva and study mathematics.

At City College, there was a very active group of mathematics students. A lot of socializing went on. There was a table in the cafeteria called the mathematics table. Between classes we would sit there and have ice cream and—

H: Discuss the topology of bagels?

A: Right, that kind of thing. A lot of chess playing, a lot of math talk. We ran our own seminars, had a math club. Some very prominent mathematicians came out of there—Jack Schwartz of Dunford–Schwartz fame, Leon Ehrenpreis, Alan Shields, Leo Flatto, Martin Davis, D. J. Newman. That was a very intense experience. From there I went on to graduate work at MIT, where I did a doctorate in algebraic topology with George Whitehead.

Let me tell you something very moving relating to my thesis. As an undergraduate, I read a lot of analytic and algebraic number theory. What is fascinating about number theory is that it uses very deep methods to attack problems that are in some sense very "natural" and also simple to formulate. A schoolchild can understand Fermat's last theorem, but it took extremely deep methods to prove it. Another interesting aspect of number theory was that it was absolutely useless—pure mathematics at its purest.

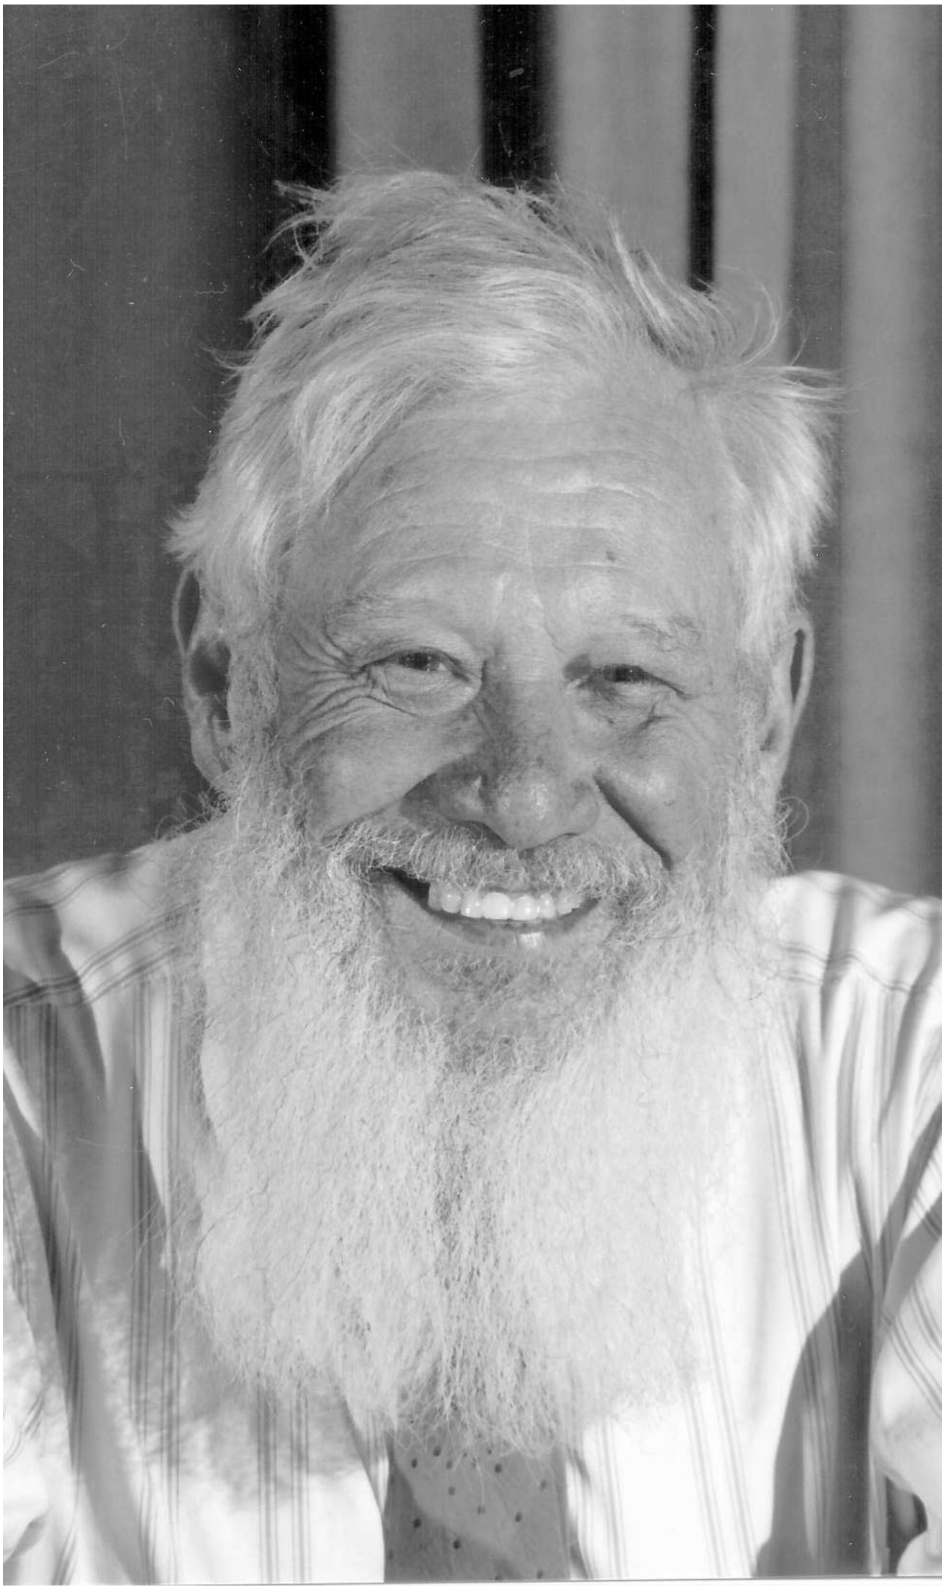

Picture 1. Bob Aumann, circa 2000

In graduate school, I heard George Whitehead's excellent lectures on algebraic topology. Whitehead did not talk much about knots, but I had heard about them, and they fascinated me. Knots are like number theory: the problems are very simple to formulate, a schoolchild can understand them; and they are very natural, they have a simplicity and immediacy that is even greater than that of Fermat's last theorem. But it is very difficult to prove anything at all about them; it requires really deep methods of algebraic topology. And, like number theory, knot theory was totally, totally useless.

So, I was attracted to knots. I went to Whitehead and said, I want to do a PhD with you, please give me a problem. But not just any problem; please, give me an open problem in knot theory. And he did; he gave me a famous, very difficult problem—the "asphericity" of knots—that had been open for twenty-five years and had defied the most concerted attempts to solve.

Though I did not solve that problem, I did solve a special case. The complete statement of my result is not easy to formulate for a layman, but it does have an interesting implication that even a schoolchild can understand and that had not been known before my work: alternating knots do not "come apart," cannot be separated.

So, I had accomplished my objective—done something that i) is the answer to a "natural" question, ii) is easy to formulate, iii) has a deep, difficult proof, and iv) is absolutely useless, the purest of pure mathematics.

It was in the fall of 1954 that I got the crucial idea that was the key to proving my result. The thesis was published in the *Annals of Mathematics* in 1956; but the proof was essentially in place in the fall of 1954.

That's Act I of the story. And now, the curtain rises on Act II—fifty years later, almost to the day. It's 10 p.m., and the phone rings in my home. My grandson Yakov Rosen is on the line. Yakov is in his second year of medical school. "Grandpa," he says, "can I pick your brain? We are studying knots. I don't understand the material, and think that our lecturer doesn't understand it either. For example, could you explain to me what, exactly, are 'linking numbers'?" "Why are you studying knots?" I ask; "what do knots have to do with medicine?" "Well," says Yakov, "sometimes the DNA in a cell gets knotted up. Depending on the characteristics of the knot, this may lead to cancer. So, we have to understand knots."

I was completely bowled over. Fifty years later, the "absolutely useless"—the "purest of the pure"—is taught in the second year of medical school, and my grandson is studying it. I invited Yakov to come over, and told him about knots, and linking numbers, and my thesis.

* * *

H: Okay, now that we are all tied up in knots, let's untangle them and go on. You did your PhD at MIT in algebraic topology, and then what?

A: Then for my post-doc, I joined an operations research group at Princeton. This was a rather sharp turn because algebraic topology is just about the purest of pure mathematics and operations research is very applied. It was a small group of about ten people at the Forrestal Research Center, which is attached to Princeton University.

H: In those days operations research and game theory were quite connected. I guess that's how you—

A: —became interested in game theory, exactly. There was a problem about defending a city from a squadron of aircraft most of which are decoys—do not carry any weapons—but a small percentage do carry nuclear weapons. The project was sponsored by Bell Labs, who were developing a defense missile.

At MIT I had met John Nash, who came there in '53 after doing his doctorate at Princeton. I was a senior graduate student and he was a Moore instructor, which was a prestigious instructorship for young mathematicians. So he was a little older than me, scientifically and also chronologically. We got to know each other fairly well and I heard from him about game theory. One of the problems that we kicked around was that of dueling—silent duels, noisy duels, and so on. So when I came to Princeton, although I didn't know much about game theory at all, I had heard about it; and when we were given this problem by Bell Labs, I was able to say, this sounds a little bit like what Nash was telling us; let's examine it from that point of view. So I started studying game theory; the rest is history, as they say.

In '56 I came to the Hebrew University. Then, in '60–'61, I was on sabbatical at Princeton, with Oskar Morgenstern's outfit, the Econometric Research Program. This was associated with the economics department, but I also spent quite a bit of time in Fine Hall, in the mathematics department.

It was great, and meeting Morgenstern and working with him was a tremendous experience, a tremendous privilege. He was extraordinary. You know, sometimes people make disparaging remarks about Morgenstern, in particular about his contributions to game theory. One of these disparaging jokes is that Morgenstern's greatest contribution to game theory is von Neumann. So let me say, maybe that's true—but that is a tremendous contribution. Morgenstern's ability to identify people, the potential in people, was enormous and magnificent, was wonderful. He identified the economic significance in the work of people like von Neumann and Abraham Wald, and succeeded in getting them actively involved. He identified the potential in many others; just in the year I was in his outfit, Clive Granger, Sidney Afriat, and Reinhard Selten were also there.

Morgenstern had his own ideas and his own opinions and his own important research in game theory, part of which was the von Neumann–Morgenstern solution to cooperative games. And, he understood the importance of the minimax theorem to economics. One of his greatnesses was that even though he could disagree with people on a scientific issue, he didn't let that interfere with promoting them and bringing them into the circle.

Picture 2. Sergiu Hart, Mike Maschler, Bob Aumann, Bob Wilson, and Oskar Morgenstern, at the 1994 Morgenstern Lecture, Jerusalem

For example, he did not like the idea of perfect competition and he did not like the idea of the core; he thought that perfect competition is a mirage, that when there are many players, perfect competition need *not* result. And indeed, if you apply the von Neumann–Morgenstern solution, it does not lead to perfect competition in markets with many people—that was part of your doctoral thesis, Sergiu. So even though he thought that things like core equivalence were wrongheaded, he still was happy and eager to support people who worked in this direction.

At Princeton I also got to know Frank Anscombe—

H: —with whom you wrote a well-known and influential paper—

A: —that was born then. At that time the accepted definition of subjective probability was Savage's. Anscombe was giving a course on the foundations of probability; he gave a lot of prominence to Savage's theory, which was quite new at the time. Savage's book had been published in '54; it was only six years old. As a result of this course, Anscombe and I worked out this alternative definition, which was published in 1963.

H: You also met Shapley at that time?

A: Well, being in game theory, one got to know the name; but personally I got to know Shapley only later. At the end of my year at Princeton, in the fall of ’61, there was a conference on “Recent Developments in Game Theory,” chaired by Morgenstern and Harold Kuhn. The outcome was the famous orange book, which is very difficult to obtain nowadays. I was the office boy, who did a lot of the practical work in preparing the conference. Shapley was an invited lecturer, so that is the first time I met him.

Another person about whom the readers of this interview may have heard, and who gave an invited lecture at that conference, was Henry Kissinger, who later became the Secretary of State of the United States.

At the conference, Kissinger spoke about game-theoretic thinking in Cold War diplomacy, Cold War international relations. It is difficult to imagine now how serious the Cold War was. People were really afraid that the world was coming to an end, and indeed there were moments when it did seem that things were hanging in the balance. One of the most vivid was the Cuban Missile Crisis in 1963. In his handling of that crisis, Kennedy was influenced by the game-theoretic school in international relations, which was quite prominent at the time. Kissinger and Herman Kahn were the main figures in that. Kennedy is now praised for his handling of that crisis; indeed, the proof of the pudding is in the eating of it—it came out well. But at that time it seemed extremely hairy, and it really looked as if the world might come to an end at any moment—not only during the Cuban Missile Crisis, but also before and after.

The late fifties and early sixties were the acme of the Cold War. There was a time around ’60 or ’61 when there was this craze of building nuclear fallout shelters. The game theorists pointed out that this could be seen by the Russians as an extremely aggressive move. Now it takes a little bit of game-theoretic thinking to understand why building a shelter can be seen as aggressive. But the reasoning is quite simple. Why would you build shelters? Because you are afraid of a nuclear attack. Why are you afraid of a nuclear attack? Well, one good reason to be afraid is that if you are going to attack the other side, then you will be concerned about retaliation. If you do not build shelters, you leave yourself open. This is seen as conciliatory because then you say, I am not concerned about being attacked because I am not going to attack you. So building shelters was seen as very aggressive and it was something very real at the time.

H: In short, when you build shelters, your cost from a nuclear war goes down, so your incentive to start a war goes up.

A: By the way, people often make the mistake of saying that war is irrational, strikes are irrational, racial discrimination is irrational. We take all the ills of the world and dismiss them by calling them irrational. They are not necessarily irrational. Though it hurts, they may be rational. Saying that war is irrational may be a big mistake. If it is rational, once we understand that it is, we can at least somehow address the problem. If we simply dismiss it as irrational we can’t address the problem.

* * *

H: Since you started talking about these topics, let's perhaps move to Mathematica, the United States Arms Control and Disarmament Agency (ACDA), and repeated games. Tell us about your famous work on repeated games. But first, what are repeated games?

A: It's when a single game is repeated many times. How exactly you model "many" may be important, but qualitatively speaking, it usually doesn't matter too much.

H: Why are these models important?

A: They model ongoing interactions. In the real world we often respond to a given game situation, not so much because of the outcome of that particular game, as because our behavior in a particular situation may affect the outcome of future situations in which a similar game is played. For example, let's say somebody promises something and we respond to that promise and then he doesn't keep it—he double-crosses us. He may turn out a winner in the short term, but a loser in the long term: if I meet up with him again, I won't trust him. Whether he is rational, whether we are both rational, is reflected not only in the outcome of the particular situation in which we are involved today, but also in how it affects future situations.

Another example is revenge, which in the short term may seem irrational; but in the long term, it may be rational, because if you take revenge, then the next time you meet that person, he will not kick you in the stomach. Altruistic behavior, revengeful behavior, any of those things, make sense when viewed from the perspective of a repeated game, but not from the perspective of a one-shot game. So, a repeated game is often more realistic than a one-shot game: it models ongoing relationships.

In 1959 I published a paper on repeated games (*ContribGameTh* IV). The brunt of that paper is that cooperative behavior in the one-shot game corresponds to equilibrium or egotistic behavior in the repeated game. This is to put it very simplistically.

H: There is the famous "Folk Theorem." In the seventies you named it, in your survey of repeated games. The name has stuck.

A: The Folk Theorem is quite similar to my '59 paper, but a good deal simpler, less deep. I called it the Folk Theorem because its authorship is not clear, like folk music, folk songs. It was in the air in the late fifties and early sixties.

H: Yours was the first full formal statement and proof of something like this. Even Luce and Raiffa, in their very influential '57 book, *Games and Decisions*, don't have the Folk Theorem.

A: The first people explicitly to consider repeated non-zero-sum games of the kind treated in my '59 paper were Luce and Raiffa. But as you say, they didn't have the Folk Theorem. Shubik's book *Strategy and Market Structure*, published in '59, has a special case of the Folk Theorem, with a proof that has the germ of the general proof.

At that time people did not necessarily publish everything they knew—in fact, they published only a small proportion of what they knew, only really deep results or something really interesting and nontrivial in the mathematical sense of the word—which is not a good sense. Some very important things would be considered trivial by a mathematician.

For example, take the Cantor diagonal method. Perhaps it really is "trivial." But it is extremely important; inter alia, Gödel's famous incompleteness theorem is based on it.

So, even within pure mathematics the trivial may be important. But certainly outside of it, there are interesting observations that are mathematically trivial—like the Folk Theorem. I knew about the Folk Theorem in the late fifties, but was too young to recognize its importance. I wanted something deeper, and that is what I did in fact publish. That's my '59 paper. It's a nice paper—my first published paper in game theory proper. But the Folk Theorem, although much easier, is more important. So it's important for a person to realize what's important. At that time I didn't have the maturity for this.

Quite possibly, other people knew about it. People were thinking about long-term interaction. There are Shapley's stochastic games, Everett's recursive games, the work of Gillette, and so on. I wasn't the only person thinking about repeated games. Anybody who thinks a little about repeated games, especially if he is a mathematician, will very soon hit on the Folk Theorem. It is not deep.

H: That's '59; let's move forward.

A: In the early sixties Morgenstern and Kuhn founded a consulting firm called Mathematica, based in Princeton, not to be confused with the software that goes by that name today. In '65 they started working with the United States Arms Control and Disarmament Agency on a project that had to do with the Geneva disarmament negotiations: a series of negotiations with the Soviet Union, on arms control and disarmament. The people on this project included Kuhn, Gerard Debreu, Herb Scarf, Reinhard Selten, John Harsanyi, Jim Mayberry, Mike Maschler, Dick Stearns (who came in a little later), and me. What struck Maschler and me was that these negotiations were taking place again and again; a good way of modeling this is a repeated game. The only thing that distinguished it from the theory of the late fifties that we discussed before is that these were repeated games of incomplete information. We did not know how many weapons the Russians held, and the Russians did not know how many weapons we held. What we—the United States—proposed to put into the agreements might influence what the Russians thought or knew that we had, and this would affect what they would do in later rounds.

H: What you do reveals something about your private information. For example, taking an action that is optimal in the short run may reveal to the other side exactly what your situation is, and then in the long run you may be worse off.

A: Right. This informational aspect is absent from the previous work, where everything was open and above board, and the issues are how one's behavior affects future interaction. Here the question is how one's *behavior* affects the other player's *knowledge*. So Maschler and I, and later Stearns, developed a theory of repeated games of incomplete information. This theory was set forth in a series of research reports between '66 and '68, which for many years were unavailable.

H: Except to the aficionados, who were passing bootlegged copies from mimeograph machines. They were extremely hard to find.

A: Eventually they were published by MIT Press in '95, together with extensive postscripts describing what has happened since the late sixties—a tremendous amount of work. The mathematically deepest started in the early seventies in Belgium, at CORE, and in Israel, mostly by my students and then by their students. Later it spread to France, Russia, and elsewhere. The area is still active.

H: What is the big insight?

A: It is always misleading to sum it up in a few words, but here goes: in the long run, you cannot use information without revealing it; you can use information only to the extent that you are willing to reveal it. A player with private information must choose between not making use of that information—and then he doesn't have to reveal it—or making use of it, and then taking the consequences of the other side finding it out. That's the big picture.

H: In addition, in a non-zero-sum situation, you may *want* to pass information to the other side; it may be mutually advantageous to reveal your information. The question is how to do it so that you can be trusted, or in technical terms, in a way that is incentive-compatible.

A: The bottom line remains similar. In that case you can use the information, not only if you are willing to reveal it, but also if you actually *want* to reveal it. It may actually have positive value to reveal the information. Then you use it *and* reveal it.

* * *

H: Let's move to another major work of yours, "Markets with a Continuum of Traders" (*Econometrica* 1964): modeling perfect competition by a continuum.

A: At Princeton in '60–'61, the Milnor–Shapley paper "Oceanic Games" caught my fancy. It treats games with an ocean—nowadays we call it a continuum—of small players, and a small number of large players, whom they called atoms. Then in the fall of '61, at the conference at which Kissinger and Lloyd Shapley were present, Herb Scarf gave a talk about large markets. He had a countable infinity of players. Before that, in '59, Martin Shubik had published a paper called "Edgeworth Market Games," in which he made a connection between the core of a large market game and the competitive equilibrium. Scarf's model somehow wasn't very satisfactory, and Herb realized that himself; afterwards, he and Debreu proved a much more satisfactory version, in their *IER* 1963 paper. The bottom line was that, under certain assumptions, the core of a large economy is close to the competitive solution, the solution to which one is led from the law of supply and demand. I heard Scarf's talk, and, as I said, the formulation was not very satisfactory. I put it together with the result of Milnor and Shapley about oceanic games, and realized that *that* has to be the right way of treating this situation: a continuum, not the countable infinity that Scarf was using. It took a while longer to put all this together, but eventually I did get a very general theorem with a continuum of traders. It has very few assumptions, and it is not a limit result. It simply says that the core of a large market is the *same* as the set of competitive outcomes.

H: Indeed, the introduction of the continuum idea to economic theory has proved indispensable to the advancement of the discipline. In the same way as in most of the natural sciences, it enables a precise and rigorous analysis, which otherwise would have been very hard or even impossible.

A: The continuum is an approximation to the "true" situation, in which the number of traders is large but finite. The purpose of the continuous approximation is to make available the powerful and elegant methods of the branch of mathematics called "analysis," in a situation where treatment by finite methods would be much more difficult or even hopeless—think of trying to do fluid mechanics by solving n-body problems for large n.

H: The continuum is the best way to start understanding what's going on. Once you have that, you can do approximations and get limit results.

A: Yes, these approximations by finite markets became a hot topic in the late sixties and early seventies. The '64 paper was followed by the *Econometrica* '66 paper on existence of competitive equilibria in continuum markets; in '75 came the paper on values of such markets, also in *Econometrica*. Then there came later papers using a continuum, by me with or without coauthors, by Werner Hildenbrand and his school, and by many, many others.

H: Before the '75 paper, you developed, together with Shapley, the theory of values of non-atomic games; this generated a huge literature. Many of your students worked on that. What are non-atomic games?

A: They are like Milnor and Shapley's oceanic games, except that in the oceanic games there were atoms—"large" players—and in non-atomic games there are no large players at all. There are *only* small players. But unlike in Milnor–Shapley, the small players may be of different kinds; the ocean is not homogeneous. The basic property is that no player by himself makes any significant contribution. An example of a non-atomic game is a large economy, consisting of small consumers and small businesses only, without large corporations or government interference. Another example is an election, modeled as a situation where no individual can affect the outcome. In a non-atomic game, large coalitions can affect the outcome, but individual players cannot.

H: And values?

A: The game theory concept of value is an a priori evaluation of what a player, or group of players, can expect to get out of the game. Lloyd Shapley's 1953 formalization is the most prominent. Sometimes, as in voting situations, value is presented as an index of power (Shapley and Shubik 1954). I have already mentioned the 1975 result about values of large economies being the same as the competitive outcomes of a market. This result had several precursors, the first of which was a '64 RAND Memorandum of Shapley.

H: Values of non-atomic games and their application in economic models led to a huge literature.

* * *

Another one of your well-known contributions is the concept of correlated equilibrium (*JMathEcon* '74). How did it come about?

A: Correlated equilibria are like mixed Nash equilibria, except that the players' randomizations need not be independent. Frankly, I'm not really sure how this business began. It's probably related to repeated games, and, indirectly, to Harsanyi and Selten's equilibrium selection. These ideas were floating around in the late sixties, especially at the very intense meetings of the Mathematica ACDA team. In the Battle of the Sexes, for example, if you're going to select *one* equilibrium, it has to be the mixed one, which is worse for *both* players than *either* of the two pure ones. So you say, hey, let's toss a coin to decide on one of the two pure equilibria. Once the coin is tossed, it's to the advantage of both players to adhere to the chosen equilibrium; the whole process, including the coin toss, is in equilibrium. This equilibrium is a lot better than the unique mixed strategy equilibrium, because it guarantees that the boy and the girl will definitely meet—either at the boxing match or at the ballet—whereas with the mixed strategy equilibrium, they may well go to different places.

With repeated games, one gets a similar result by alternating: one evening boxing, the next ballet. Of course, that way one only gets to the convex hull of the Nash equilibria.

This is pretty straightforward. The next step is less so. It is to go to three-person games, where two of the three players gang up on the third—correlate "against" him, so to speak. This leads *outside* the convex hull of Nash equilibria. In writing this formally, I realized that the same definitions apply also to two-person games; also there, they may lead outside the convex hull of the Nash equilibria.

H: So, correlated equilibria arise when the players get signals that need not be independent. Talking about signals and information—how about common knowledge and the "Agreeing to Disagree" paper?

A: The original paper on correlated equilibrium also discussed "subjective equilibrium," where different players have different probabilities for the same event. Differences in probabilities can arise from differences in information; but then, if a player knows that another player's probability is different from his, he might wish to revise his own probability. It's not clear whether this process of revision necessarily leads to the same probabilities. This question was raised—and left open—in the '74 paper. Indeed, even the formulation of the question was murky.

I discussed this with Arrow and Frank Hahn during an IMSSS summer in the early seventies. I remember the moment vividly. We were sitting in Frank Hahn's small office on the fourth floor of Stanford's Encina Hall, where the economics department was located. I was trying to get my head around the problem—not its solution, but simply its formulation. Discussing it with them—describing the issue to them—somehow sharpened and clarified it. I went back to my office, sat down, and continued thinking. Suddenly the whole thing came to me in a flash—the definition of common knowledge, the characterization in terms of information partitions, and the agreement theorem: roughly, that if the probabilities of two people for an event are commonly known by both, then they *must* be equal. It took a couple of days more to get a coherent proof and to write it down. The proof seemed quite straightforward. The whole thing—definition, formulation, proof—came to less than a page.

Indeed, it looked so straightforward that it seemed hardly worth publishing. I went back and told Arrow and Hahn about it. At first Arrow wouldn't believe it, but became convinced when he saw the proof. I expressed to him my doubts about publication. He strongly urged me to publish it—so I did (*AnnStat* '76). It became one of my two most widely cited papers.

Six or seven years later I learned that the philosopher David Lewis had defined the concept of common knowledge already in 1969, and, surprisingly, had used the same name for it. Of course, there is no question that Lewis has priority. He did not, however, have the agreement theorem.

H: The agreement theorem is surprising—and important. But your simple and elegant formalization of common knowledge is even more important. It pioneered the area known as "interactive epistemology": knowledge about others' knowledge. It generated a huge literature—in game theory, economics, and beyond: computer science, philosophy, logic. It enabled the rigorous analysis of very deep and complex issues, such as what is rationality, and what is needed for equilibrium. Interestingly, it led you in particular back to correlated equilibrium.

A: Yes. That's my *Econometrica* '87 paper. The idea of common knowledge really enables the "right" formulation of correlated equilibrium. It's not some kind of esoteric extension of Nash equilibrium. Rather, it says that if people simply respond optimally to their information—and this is commonly known—then you get correlated equilibrium. The "equilibrium" part of this is not the point. Correlated equilibrium is nothing more than just common knowledge of rationality, together with common priors.

* * *

H: Let's make a big jump. In 1991, the Center for Rationality was established at the Hebrew University.

A: Yoram Ben-Porath, who was rector of the university, asked Menahem Yaari, Itamar Pitowsky, Motty Perry, and me to make a proposal for establishing an interdisciplinary center. What came out was the Center for Rationality, which you, Sergiu, directed for its first eight critical years; it was you who really got it going and gave it its oomph. The Center is really unique in the whole world in that it brings together very many disciplines. Throughout the world there are several research centers in areas connected with game theory. Usually they are associated with departments of economics: the Cowles Foundation at Yale, CORE in Louvain, the late Institute for Mathematical Studies in the Social Sciences at Stanford. The Center for Rationality at the Hebrew University is quite different, in that it is much broader. The basic idea is "rationality": behavior that advances one's own interests. This appears in many different contexts, represented by many academic disciplines. The Center has members from mathematics, economics, computer science, evolutionary biology, general philosophy, philosophy of science, psychology, law, statistics, the business school, and education. There is nothing in the world even approaching the breadth of coverage of the Center for Rationality.

It is broad but nevertheless focused. There would seem to be a contradiction between breadth and focus, but our Center has both—breadth and focus. The breadth is in the number and range of different disciplines that are represented at the Center. The focus is, in all these disciplines, on rational, self-interested behavior—or the lack of it. We take all these different disciplines, and we look at a certain segment of each one, and at how these various segments from this great number of disciplines fit together.

H: Can you give a few examples? Readers may be surprised to hear about some of these connections.

A: I'll try; let's go through some applications. In computer science we have distributed computing, in which there are many different processors. The problem is to coordinate the work of these processors, which may number in the hundreds of thousands, each doing its own work.

H: That is, how processors that work in a decentralized way reach a coordinated goal.

A: Exactly. Another application is protecting computers against hackers who are trying to break down the computer. This is a very grim game, but it is a game. Still another kind comes from computers that solve games, play games, and design games—like auctions—particularly on the Web.

Biology is another example where one might think that games don't seem particularly relevant. But they are! There is a book by Richard Dawkins called *The Selfish Gene*. This book discusses how evolution makes organisms operate as if they were promoting their self-interest, acting rationally. What drives this is the survival of the fittest. If the genes that organisms have developed in the course of evolution are not optimal, are not doing as well as other genes, then they will not survive. There is a tremendous range of applications of game-theoretic and rationalistic reasoning in evolutionary biology.

Economics is of course the main area of application of game theory. The book by von Neumann and Morgenstern that started game theory rolling is called *The Theory of Games and Economic Behavior*. Psychology—that of decision-making—has close ties to game theory; whether behavior is rational or irrational—the *subject* is still rationality.

There is much political application of game theory in international relations, which we already discussed in connection with Kissinger. There also are national politics, like various electoral systems. Another aspect is forming a government coalition: if it is too small—a minimal winning coalition—is will be unstable; if too large, the prime minister will have too little influence. What is the right balance?

Law: more and more, we have law and economics, law and game theory. There are studies of how laws affect the behavior of people, the behavior of criminals, the behavior of the police. All these things are about self-interested, rational behavior.

* * *

H: So that's the Center for Rationality. I know this doesn't belong, but I'll ask it here. You are a deeply religious man. How does it fit in with a rational view of the world? How do you fit together science and religion?

A: As you say, it really doesn't belong here, but I'll respond anyway. Before responding directly, let me say that the scientific view of the world is really just in our minds. When you look at it carefully, it is not something that is out there in the real world. For example, take the statement "the earth is round." It sounds like a very simple statement that is either true or false. Either the earth is round or it isn't; maybe it is square, or elliptical, or whatever. Roundness means that there is a point—the "center" of the earth—such that any point on the surface of the earth is at the same distance from that center as any other point on the surface of the earth. Now that already sounds a little complex. But the complexity only begins there. What exactly do we mean by equal distance? For that you need the concept of a distance between two points. The concept of distance between two points is something that is fairly complex even if we are talking about a ball that we can hold in our hands; it involves taking a ruler and measuring the distance between two points. But when we are talking about the earth, it is even more complex, because there is no way that we are going to measure the distance between the center of the earth and the surface of the earth with a ruler. One problem is that we can't get to the center. Even if we could find it we wouldn't be able to get there. We certainly wouldn't be able to find a ruler that is big enough. So we have to use some kind of complex theory in order to give that a practical meaning. Even when we have four points and we say the distance from A to B is the same as the distance from C to D, that is fairly complex already. Maybe the ruler changes. We are using a whole big theory, a whole big collection of ideas, in order to give meaning to this very, very simple statement that the earth is round.

Don't get me wrong. We all agree that the earth is round. What I am saying is that the roundness of the earth is a concept that is in our minds. It's a product of a very complex set of ideas, and ideas are in people's minds. So the way I think of science, and even of fairly simple things, is as being in our minds; all the more so for things like gravitation, the energy that is emitted by a star, or even the concept of a "species." Yes, we are both members of the species homo sapiens. What does that mean? Obviously we are different. My beard is much longer than yours. What exactly does species mean? What exactly does it even mean to say "Bob Aumann" is sitting here? Is it the same Bob Aumann as five minutes ago? These are very complex ideas. Identity, all those things that we think of trivially on a day-to-day basis, are really complex ideas that are in our minds; they are not really out there. Science is built to satisfy certain needs in our minds. It describes *us*. It does have a relationship with the real world, but this relationship is very, very complex.

Having said that, I'll get to your question. Religion is very different from science. The main part of religion is not about the way that we model the real world. I am purposely using the word "model." Religion is an experience—mainly an emotional and aesthetic one. It is not about whether the earth is 5,765 years old. When you play the piano, when you climb a mountain, does this contradict your scientific endeavors? Obviously not. The two things are almost—though not quite—orthogonal. Hiking, skiing, dancing, bringing up your children—you do all kinds of things that are almost orthogonal to your scientific endeavor. That's the case with religion also. It doesn't contradict; it is orthogonal. Belief is an important part of religion, certainly; but in science we have certain ways of thinking about the world, and in religion we have different ways of thinking about the world. Those two things coexist side by side without conflict.

H: A world populated by rational players—is it consistent with the religious view?

A: Yes. Religion places a lot of emphasis on coliving with your fellow man. A large part of religion is, be nice to other people. We can understand this in the religious context for what it is and we can understand it scientifically in the sense of repeated games that we discussed before, and we can understand it from the evolutionary viewpoint. These are different ways of understanding the phenomenon; there is no contradiction there. Fully rational players could be deeply religious; religion reflects other drives.

H: This applies to person-to-person interaction. But isn't there, in a sense, an extra player, which would be G-d?

A: Certainly, G-d has His objectives and desires. These are set forth in the Torah, which is interpreted—given practical meaning—by human beings.

It is this practical meaning that lies behind the figure of G-d—call it a model, a way of thinking, a way of living. It is similar, broadly speaking, to the earth being round. G-d is a way of thinking of our lives; translated into practical terms, it tells us how to live.

H: Let me try to summarize. On the one hand there is an emotional and aesthetic experience, to which I can very clearly relate, like going to a concert or seeing something beautiful. On the other hand, religion dictates certain rules of behavior. These rules, first of all, are not well defined. They are interpreted by human beings. Second, these rules may be justified in a rational way. Like in your *JET* '85 paper with Michael Maschler, where you gave a game-theoretic interpretation of a passage from the Talmud that nobody could understand, and suddenly everything became crystal clear. So you are saying that there are rules, which are good rules. And they are good not just because G-d gave them to us. We may not understand the reasons, but if we go deep enough and start analyzing, we may find good reasons for them. Moreover, if people are following these rules it leads perhaps to a better society—a Pareto improvement. Is that correct?

A: Well, it is your way of putting it. Let me enlarge on it. The observance of the Sabbath is extremely beautiful, and is impossible without being religious. It is not even a question of improving society—it is about improving one's own quality of life. For example, let's say I'm taking a trip a couple of hours after the Sabbath. Any other person would spend the day packing, going to the office, making final arrangements, final phone calls, this and that. For me it's out of the question. I do it on Friday. The Sabbath is *there*. The world stops.

So, we have this one day in the week when nothing can come in the way and we are shut off from the world. We don't answer the phone, we don't operate electricity, we don't drive cars.

H: It is a self-committing device, if you translate it into rational terms.

A: Exactly, it's a self-committing device.

Here is another example. There was a period fifteen, twenty years ago when stealing software was considered okay by many people, including many academics. There was an item of software that I needed, and I was wondering whether to "steal" it—make a copy of which the developers of the software disapprove. Then I said to myself, why do you have to wonder about this? You are a religious person. Go to your rabbi and ask him. So I went to my rabbi. I figured, whatever he'll say, I'll do. He said, "It's absolutely forbidden to do this, absolutely forbidden."

In short, you can be a moral person, but morals are often equivocal. In the eighties, copying software was considered moral by many people. The point I am making is that religion—at least my religion—is a sort of force, a way of making a commitment to conduct yourself in a certain way, which is good for the individual and good for society.

H: But then, in a world where everybody follows these rules, there is perhaps no reason for game theory. Of course, there is a problem in the details; the rules of conduct may not be enough to tell you exactly what to do in every situation. But in principle, in a world populated by religious people, do we need game theory?

A: Certainly. The rules cover only the moral or ethical issues. There is a lot of room within these rules for strategic behavior. For example, the rules tell you that if you made an offer and it was accepted, then you can't renege. But they don't tell you how much to offer. The rules tell you that you must bargain in good faith, but they don't tell you whether to be tough, or compromising, or whatever. The rules tell you, "You may not steal software"; but they don't tell you how much to pay for the software, when to buy it and when not. The rules tell you to give a lot to charity, but not how much. There was a study made in the United States of income tax deductions to charity. It turned out that orthodox Jews were among the largest contributors to charities. It's a religious command.

The religious community, by the way, is very close. The matter of *khessed*, of helping your fellow man, is very strong in religious communities; it is a commandment, like eating kosher and keeping the Sabbath.

Up to now we have been discussing the normative side of game theory—advising individuals how to act—but there are also other sides. One is "public normative." The religion will not tell you how to conduct elections, or when to cut the discount rate, or how to form a government. It will not tell you how to build a distributed computer, or how to run a spectrum auction, or how to assign interns to hospitals.

Besides, one needs game theory to explain the ethical and moral rules themselves. *Why* not steal software? *Why* have accurate weights and measures? Why love one's neighbor as oneself? How did it come about, what function does it serve, what keeps it together? All these are game-theoretic questions.

Finally, let's not forget that the world is very far from being—to use your phrase—populated by religious people only.

In short, the Bible and the Talmud are fascinating documents, and they cover a lot of ground, but there still is a lot of room for game theory—and for all of science.

H: So, to summarize this point: game theory definitely has a place in a religious world. In the "micro," the rules of conduct are principles that cover only certain issues, and there is "freedom of decision." In the "macro," the structures that arise, and the rules of conduct themselves, are subject to game-theoretic analysis: how and why did they come about?

Rational people can very well exist in this religious world. You have reconciled that very nicely. That was very interesting.

A: I haven't reconciled. I tried not to reconcile, but to say, these are different things.

H: Reconciled in the sense that those things can coexist.

* * *

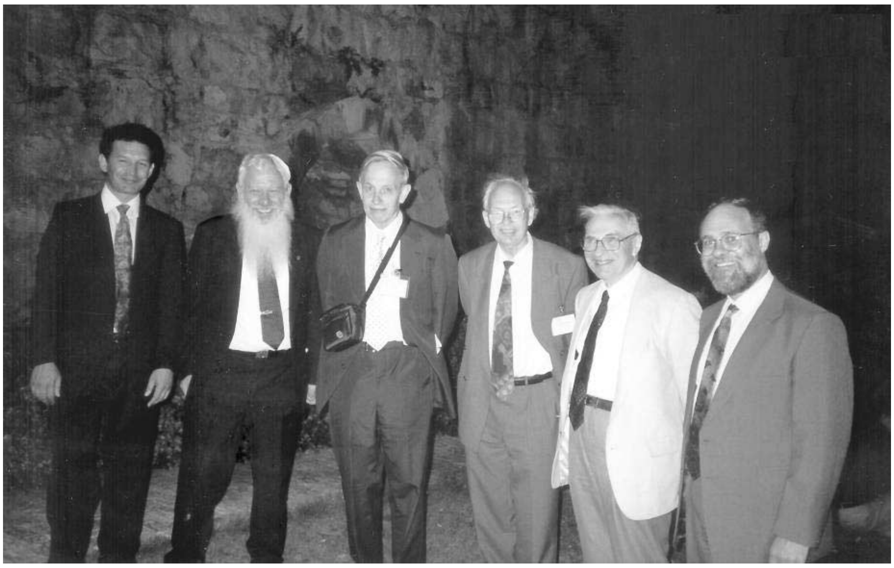

Picture 6. At the GAMES 1995 Conference in Honor of Aumann's 65th birthday, Jerusalem, June 1995: Abraham Neyman, Bob Aumann, John Nash, Reinhard Selten, Ken Arrow, Sergiu Hart (from left to right)

Let's move now to your personal biography.

A: I was born in 1930 in Frankfurt, Germany, to an orthodox Jewish family, the second of two boys. My father was a wholesale textile merchant—a fine, upright man, a loving, warm father. My mother was extraordinary. She got a bachelor's degree in England in 1914, at a time when that was very unusual for women. She was a medal-winning long-distance swimmer, sang Shubert lieder while accompanying herself on the piano, introduced us children to nature, music, reading. We would walk the streets and she would teach us the names of the trees. At night we looked at the sky and she taught us the names of the constellations. When I was about twelve, we started reading Dickens's *A Tale of Two Cities* together—until the book gripped me and I raced ahead alone. From then on, I read voraciously. She even introduced me to interactive epistemology; look at the "folk ditty" in *GEB* '96. She always encouraged, always pushed us along, gently, unobtrusively, always allowed us to make our own decisions.

We got away in 1938. Actually we had planned to leave already when Hitler came to power in 1933, but for one reason or another we didn't.

People convinced my parents that it wasn't so bad; it will be okay, this thing will blow over. The German people will not allow such a madman to take over, etc., etc. A well-known story. But it illustrates that when one is in the middle of things it is very, very difficult to see the future. Things seem clear in hindsight, but in the middle of the crisis they are very murky.

H: Especially when it is a slow-moving process, rather than a dramatic change: every time it is just a little more and you say, that's not much, but when you look at the integral of all this, suddenly it is a big change.

A: That is one thing. But even more basically, it is just difficult to see. Let me jump forward from 1933 to 1967. I was in Israel and there was the crisis preceding the Six-Day War. In hindsight it was "clear" that Israel would come out on top of that conflict. But at the time it wasn't at all clear, not at all. I vividly remember the weeks leading up to the Six-Day War, the crisis in which Nasser closed the Tiran Straits and massed troops on Israel's border; it wasn't at all clear that Israel would survive. Not only to me, but to anybody in the general population. Maybe our generals were confident, but I don't think so, because our government certainly was not confident. Prime Minister Eshkol was very worried. He made a broadcast in which he stuttered and his concern was very evident, very real. Nobody knew what was going to happen; people were very worried, and I, too, was very worried. I had a wife and three children and we all had American papers. So I said to myself, Johnny, don't make the mistake your father made by staying in Germany. Pick yourself up, get on a plane and leave, and save your skin and that of your family; because there is a very good chance that Israel will be destroyed and the inhabitants of Israel will be wiped out totally, killed, in the next two or three weeks. Pick yourself up and GO.

I made a conscious decision not to do that. I said, I am staying. Herb Scarf was here during the crisis. When he left, about two weeks before the war, we said good-bye, and it was clear to both of us that we might never see each other again.

This illustrates that it is very difficult to judge a situation from the middle of it. When you're swimming in a big lake, it's difficult to see the shore, because you are low, you are inside it. One should not blame the German Jews or the European Jews for not leaving Europe in the thirties, because it was difficult to assess the situation.

We did get away in time, in 1938. We left Germany, and made our way to the United States. In this passage, my parents lost all their money. They had to work extremely hard in the United States to make ends meet, but nevertheless they gave their two children, my brother and myself, a good Jewish and a good secular education.

When the State of Israel was created in 1948, I made a determination eventually to come to Israel, but that didn't actually happen until 1956. In 1954 I met an Israeli girl, Esther Schlesinger, who was visiting the United States. We fell in love, got engaged, and got married. We had five children; the oldest, Shlomo, was killed in action in Lebanon in 1982. My other children are all happily married. Shlomo's widow also remarried and she is like a daughter to us. Shlomo had two children, the second one born after he was killed. Altogether I now have seventeen grandchildren and one great-grandchild. We have a very good family relationship, do a lot of things together. One of the things we like best is skiing. Every year I go with a different part of the family. Once in four or five years, all thirty of us go together.

H: I can attest from my personal knowledge that the Aumann family is really an outstanding, warm, unusually close-knit family. It is really great to be with them.

A: My wife Esther died six years ago, of cancer, after being ill for about a year and a half. She was an extraordinary person. After elementary school she entered the Bezalel School of Art—she had a great talent for art. At Bezalel she learned silversmithing, and she also drew well. She was wonderful with her hands and also with people. When about fifty, she went to work for the Frankforter Center, an old-age day activities center; she ran the crafts workshop, where the elderly worked with their hands: appliqué, knitting, embroidery, carpets, and so on. This enabled Esther to combine her two favorite activities: her artistic ability, and dealing with people and helping them, each one with his individual troubles.

When she went to school, Bezalel was a rather Bohemian place. It probably still is, but at that time it was less fashionable to be Bohemian, more special. Her parents were very much opposed to this. In an orthodox Jewish family, a young girl going to this place was really unheard of. But Esther had her own will. She was a mild-mannered person, but when she wanted something, you bet your life she got it, both with her parents and with me. She definitely did want to go to that school, and she went.

* * *

H: There is a nice story about your decision to come to Israel in '56.

A: In '56 I had just finished two years of a post-doc at Princeton, and was wondering how to continue my life. As mentioned, I had made up my mind to come to Israel eventually. One of the places where I applied was the Hebrew University in Jerusalem. I also applied to other places, because one doesn't put all one's eggs in one basket, and got several offers. One was from Bell Telephone Laboratories in Murray Hill; one from Jerusalem; and there were others. Thinking things over very hard and agonizing over this decision, I finally decided to accept the position at Bell Labs, and told them that. We started looking around for a place to live on that very same day.

When we came home in the evening, I knew I had made the wrong decision. I had agonized over it for three weeks or more, but once it had been made, it was clear to me that it was wrong. Before it had been made, nothing was clear. Now, I realized that I wanted to go to Israel immediately, that there is no point in putting it off, no point in trying to earn some money to finance the trip to Israel; we'll just get stuck in the United States. If we are going to go at all we should go right away. I called up the Bell Labs people and said, "I changed my mind. I said I'll come, so I'll come, but you should know that I'm leaving in one year." They said, "Aumann, you're off the hook. You don't have to come if you don't want to." I said, "Okay, but now it's June. I am not leaving until October. Could I work until then at Bell Labs?" They said, "Sure, we'll be glad to have you." That was very nice of them.

This story illustrates the difficulties with practical decision-making. It's much more complex than our models. In practical decision-making, you don't know the right decision until after you've made it.

H: This, at least to my mind, is a good example of some of your views on experiments and empirics. Do you want to expand on that?

A: Yes. I have grave doubts about what's *called* "behavioral economics," but isn't really behavioral. The term implies that that is how people actually behave, as distinguished from what the theory says. But that's not what behavioral economics is about. On the contrary, most of behavioral economics deals with artificial laboratory setups, at best. At worst, it deals with polls, questionnaires. One type of so-called behavioral economics is when people are asked, what would you do if you were faced with such and such a situation. Then they have to imagine that they are in this situation and they have to give an answer.

H: Your example of Bell Labs versus the Hebrew University shows that you really can give the wrong answer when you are asked such a question.

A: Polls and questionnaires are worse than that; they are at a double remove from reality. In the Bell Labs case, I actually was faced with the problem of which job to take. Even then I took a decision that was not the final one, in spite of the setup being real. In "behavioral economics," people ask, "What would you do if …"; it is not even a real setup.

Behavioral economists also do experiments with monetary rewards. But these are usually very small. More importantly, the decisions that people face are not ones that they usually take. The whole setup is artificial. It is not a decision that really affects them and to which they are used.

An example is the famous "probability matching" experiment. A light periodically flashes, three quarters of the time green, one quarter red, at random. The subject has to guess the color beforehand, and gets rewarded if he guesses correctly. This experiment has been repeated hundreds of times; by far the largest number of subjects guess green three quarters of the time and red one quarter of the time.

That is not optimal; you should always guess green. If you get a dollar each time you guess correctly, and you probability-match—three quarters, one quarter—then your expected payoff is five eighths of a dollar. If you guess green all the time you get an average of three quarters of a dollar. Nevertheless, people probability-match. The point is that the setting is artificial: people don't usually sit in front of flashing lights. They don't know how to react, so they do what they think is expected of them, which becomes probability-matching.

In real situations people don't act that way. An example is driving to work in the morning. Many people have a choice of routes, and each route has a certain probability of taking less time. It is random, because one can't know where there will be an accident, a traffic jam. Let's say that there are two routes; one is quicker three quarters of the time and the other, one quarter of the time. Most people will settle down and take the same route every day, although some days it will be the longer one; and that is the correct solution.

In short, I have serious doubts about behavioral economics as it is practiced. Now, *true* behavioral economics does in fact exist; it is called empirical economics. This really *is* behavioral economics. In empirical economics, you go and see how people behave in real life, in situations to which they are used. Things they do every day.

There is a wonderful publication called the *NBER Reporter*. NBER is the National Bureau of Economic Research, an American organization. They put out a monthly newsletter of four to six pages, in which they give brief summaries of research memoranda published during that month. It is all empirical. There is nothing theoretical there. Sometimes they give theoretical background, but all these works are empirical works that say how people actually behave. It is amazing to see how well the actual behavior of people fits economic theory.

H: If I may summarize your views on this, empirical economics is a good way of finding out what people actually decide. On the other hand, much of what is done in experimental work is artificial and people may not behave there as they do in real life.

A: Yes. Let me expand on that a little bit. The thesis that behavioral economics attacks is that people behave rationally in a conscious way—that they consciously calculate and make an optimal decision based, in each case, on rational calculations. Perhaps behavioral economists are right that that is not so. Because their experiments or polls show that people, when faced with certain kinds of decisions, do not make the rational decision. However, nobody ever claimed that; they are attacking a straw man, a dead horse. What *is* claimed is that economic agents behave in a way that could be described as derived from rationality considerations; not that they actually are derived that way, that they actually go through a process of optimization each time they make a decision. Milton Friedman already had this idea, that people behave *as if* they were rational.

* * *

H: This brings us to the matter of "rule rationality," which you have been promoting consistently at least since the nineties.

A: Rule rationality means that people evolve rules of behavior by which they usually act, and they optimize these *rules*. They don't optimize each single decision. One very good example is the ultimatum game, an experiment performed by Werner Güth and associates in the early eighties.

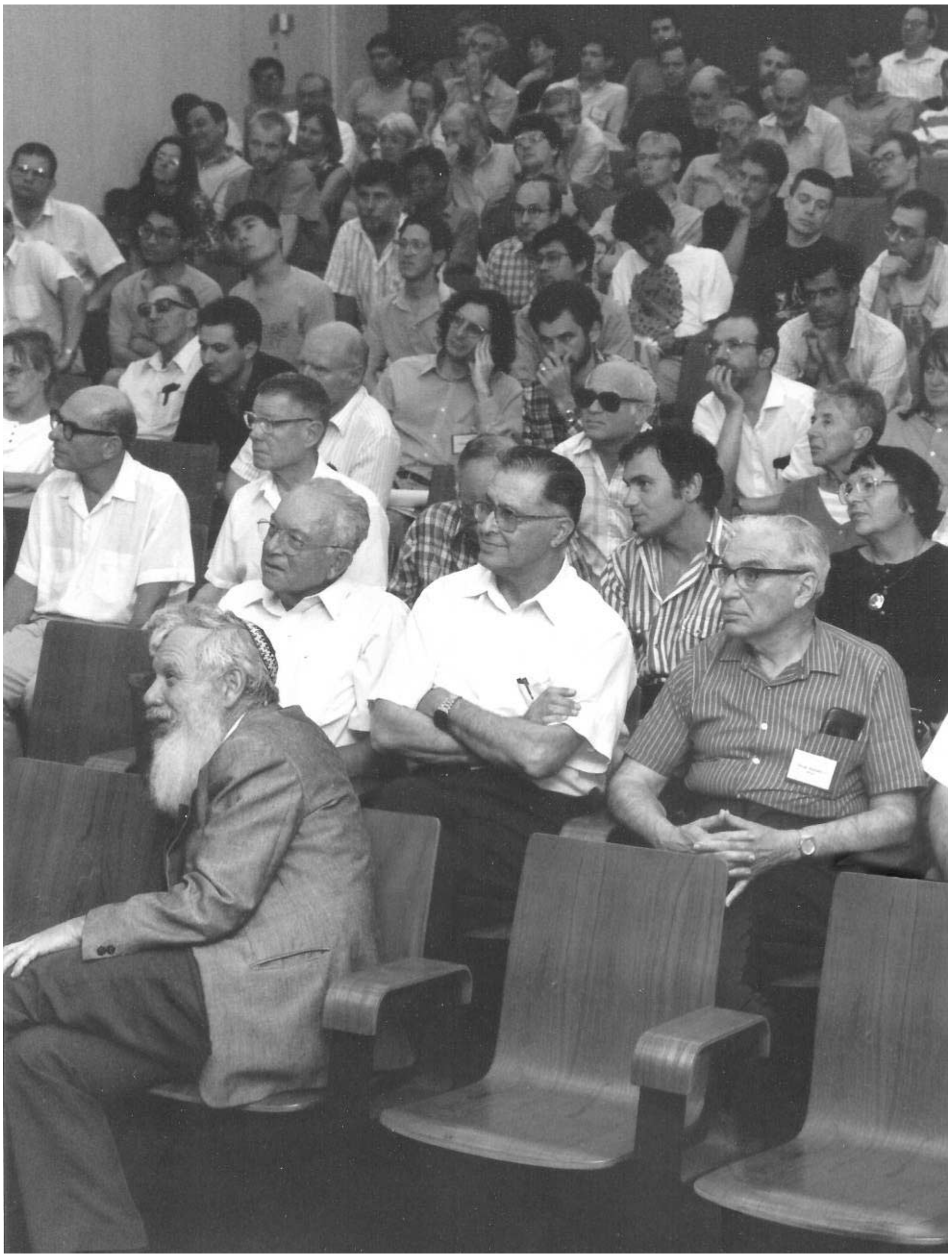

Picture 7. At the 1994 Morgenstern Lecture, Jerusalem:
Bob Aumann (front row), Don Patinkin, Mike Maschler, Ken Arrow (second row, left to right), Tom Schelling (third row, second from left); also Marshall Sarnat, Jonathan Shalev, Michael Beenstock, Dieter Balkenborg, Eytan Sheshinski, Edna Ullmann-Margalit, Maya Bar-Hillel, Gershon Ben-Shakhar, Benjamin Weiss, Reuben Gronau, Motty Perry, Menahem Yaari, Zur Shapira, David Budescu, Gary Bornstein

H: And then replicated in many forms by other people. It is a famous experiment.

A: This experiment was done in various forms and with various parameters. Here is one form. Two subjects are offered 100 Deutsch Marks, which in the early eighties was equivalent to 150–200 Euros of today—a highly non-negligible amount. They are offered this amount to split in whatever way they choose, as long as they agree how. If they cannot agree, then both get nothing. The subjects do not speak with each other face to face; rather, each one sits at a computer console. One is the offerer and the other, the responder. The offerer offers a split and the responder must say yes or no. If he says yes, that's the outcome. If he says no, no one gets anything.

This experiment was done separately for many pairs of people. Each pair played only once; they entered and left the building by different entrances and exits, and never got to know each other—remained entirely anonymous. The perfect equilibrium of this game is that the offerer offers a minimum amount that still gives the responder something. Let's say a split of 99 for the offerer and 1 for the responder.

H: The idea being that the responder would not leave even one DM on the table by saying no.

A: That is what one might expect from rationality considerations. I say "might," because it is not what game theory necessarily predicts; the game has many other equilibria. But rationality considerations might lead to the 99-1 split.

In fact, what happened was that most of the offers were in the area of 65-35. Those that were considerably less—let's say 80-20—were actually rejected. Not always, but that was the big picture. In many cases a subject was willing to walk away from as much as 20 DM; and the offerer usually anticipated this and therefore offered him more.

Walking away from 20 DM appears to be a clear violation of rationality. It *is* a violation—of *act* rationality. How does theory account for this?

The answer is that people do not maximize on an act-by-act basis. Rather, they develop *rules* of behavior. One good rule is, do not let other people insult you. Do not let other people kick you in the stomach. Do not be a sucker. If somebody does something like that to you, respond by kicking back. This is a good rule in situations that are not anonymous. If you get a reputation for accepting twenty or ten or one when 100 DM are on the table, you will come out on the short end of many bargaining situations. Therefore, the rule of behavior is to fight back and punish the person who does this to you, and then he won't do it again.

This does not apply in the current situation, because it is entirely anonymous. Nobody will be told that you did this. Therefore, there are no reputational effects, and this rule that you've developed does not apply. But you have not developed the rule consciously. You have not worked it out. You have learned it because it works in general. Therefore you apply it even in situations where, rationally speaking, it does not apply. It is very important to test economic theories in contexts that are familiar to people, in contexts in which people really engage on a regular basis. Not in artificial contexts. In artificial contexts, other things apply.

Another example of rule rationality is trying to please. It is a good idea to please the people with whom you deal. Even this can be entirely subconscious or unconscious. Most people know that voting in elections is considered a positive thing to do. So if you are asked, “Did you vote?”, there is a very strong tendency to say yes, even if you didn’t vote. Camil Fuchs, one of the important polltakers in Israel, gave a lecture at the Center for Rationality, in which he reported this: in the last election in Israel, people were asked several hours after the polls closed, did you vote? Ninety percent of the people in the sample said yes; in fact, only sixty-eight percent of the electorate voted.

H: It calls into question what we learn from polls.

A: It sheds a tremendous amount of doubt; and it shows something even more basic. Namely, that when people answer questions in a poll, they try to guess what it is that the questioner wants to hear. They give that answer rather than the true answer; and again, this is not something that they do consciously.

I am not saying that people do this because there is something in it for them. They do it because they have a general rule: try to please the people to whom you are talking; usually they can help you. If you are unpleasant to them it is usually not to your good. So people subconsciously develop tools to be pleasant and being pleasant means giving the answer that’s expected.

H: What you are saying is that one should evaluate actions not on a decision-by-decision basis, but over the long run.

Kahneman and Tversky say that there are a lot of heuristics that people use, and biases, and that these biases are not random, but systematic. You say, yes, systematic biases occur because if you look at the level of the rule, rules indeed are systematic; they lead to biases since they are not optimal for each individual act. Systematic biases fit rule rationality very well.

A: That’s a good way of putting it. If you look at those systematic biases carefully you may well find that they are rule optimal. In most situations that people encounter, those systematic biases are a short way of doing the right thing.

* * *

H: Game theory has really grown—from a few dozen people in the fifties and sixties, to more than six hundred at the last game theory congress in Marseille—which is only the tip of the iceberg. The Game Theory Society was established in ’99. You were the first, founding president. So, what is game theory?

A: Game theory is the study of interactions from a rational viewpoint. Even though the rationality does not have to be conscious, it is still there in the background. So we are interpreting what we see in the world from a rational viewpoint.

In other words, we ask, what is best for people to do when there are other people, other decision-makers, other entities who also optimize their decisions? Game theory is optimal decision-making in the presence of others with different objectives.

H: And where everyone's decision influences everyone's outcomes. One takes into account that everyone is doing his own optimization and everyone is trying to advance his own objectives.

This is a good point to discuss the universality of game theory. In the Preface to the first volume of the *Handbook of Game Theory* we wrote that game theory may be viewed as a sort of umbrella or unified field theory.

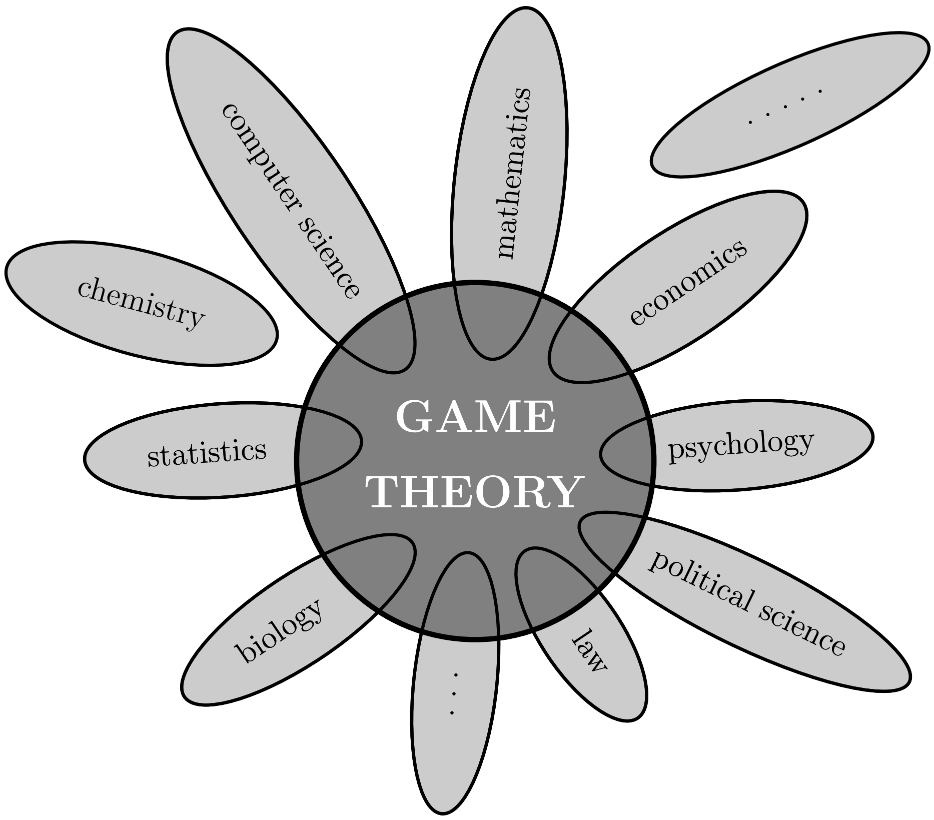

Figure 1. The blooming of game theory

A: It's a way of talking about many sciences, many disparate disciplines. Unlike other approaches to disciplines like economics or political science, game theory does not use different, ad-hoc constructs to deal with various specific issues, such as perfect competition, monopoly, oligopoly, international trade, taxation, voting, deterrence, animal behavior, and so on. Rather, it develops methodologies that apply in principle to all interactive situations, then sees where these methodologies lead in each specific application.

But rather than being an umbrella for all those disciplines, it's perhaps better to think of it as a way of thinking about a *certain aspect* of each—the interactively rational aspect. There are many things in these disciplines that have nothing to do with this aspect. In law, in computer science, in mathematics, in economics, in politics, there are many things that have nothing to do with game theory. It is not like a unified field theory, which would cover *all* of gravitation, magnetism, and electricity. Figure 1 is a stylized representation.

* * *

H: You have had an enormous impact on the profession by influencing many people. Foremost, your thirteen doctoral students. Twelve of them are by now professors who are well recognized in the field and also in related fields.

A: These are my doctoral students: Bezalel Peleg, David Schmeidler, Shmuel Zamir, Binyamin Shitovitz, Zvi Artstein, Elon Kohlberg, Sergiu Hart, Eugene Wesley, Abraham Neyman, Yair Tauman, Dov Samet, Ehud Lehrer, and Yossi Feinberg. Also, there are about thirty or forty masters students.

Each student is different. They are all great. In all cases I refused to do what some people do, and that is to write a doctoral thesis for the student. The student had to go and work it out by himself. In some cases I gave very difficult problems. Sometimes I had to backtrack and suggest different problems, because the student wasn't making progress. There were one or two cases where a student didn't make it—started working and didn't make progress for a year or two and I saw that he wasn't going to be able to make it with me. I informed him and he left. I always had a policy of taking only those students who seemed very, very capable. I usually worked quite closely with them, meeting once a week or so at least, hearing about progress, making suggestions, asking questions.

But it is really more interesting to hear from the students, so, Professor Hart, what do you think?

H: Most doctoral students want to finish their thesis and get out as soon as possible. Aumann's students usually want to continue—up to a point, of course. This was one of the best periods in my life—being immersed in research and bouncing ideas back and forth with Professor Aumann; it was a very exciting period. It was very educating for my whole life. Having a good doctoral advisor is a great investment for life. There are many stories among your students, who are still very close to one another.

* * *

H: Any closing "words of wisdom"?

A: Just one: Game theory is ethically neutral. That is, game theorists don't necessarily advocate carrying out the normative prescriptions of game theory. Bacteriologists do not advocate disease, they study it; similarly, studying self-interested behavior is different from advocating it. Game theory says nothing about whether the "rational" way is morally or ethically right. It just says what rational—self-interested—entities will do; not what they "should" do, ethically speaking. If we want a better world, we had better pay attention to where rational incentives lead.

H: That's a very good conclusion to this fascinating interview. Thank you.

A: And thank *you*, Sergiu, for your part in this wonderful interview.